

Cloudera Edge Management Installation

Date published: 2019-04-15

Date modified: 2023-04-18



Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Before you begin installing EFM.....	4
System requirements for EFM.....	4
Install Java for EFM.....	4
Installing the EFM server.....	5
Installing databases for EFM.....	5
Installing and configuring MySQL.....	5
Installing and configuring PostgreSQL.....	6
Installing and configuring MariaDB.....	8
Installing EFM server.....	9
Installing EFM as an operating system service.....	10
Configuring the EFM server.....	11
Set the encryption password for EFM.....	11
Open network ports for EFM.....	11
Starting the EFM server.....	11

Before you begin installing EFM

To start using Edge Flow Manager (EFM), you need to install it. Learn how to obtain the EFM software bits, install EFM, configure EFM, and configure security. Also learn the system requirements to do so.

System requirements for EFM

Before you begin your installation of Edge Flow Manager (EFM), carefully review the system requirements for EFM to understand operating system, database, browser, and JDK support.

Operating system support

Operating System	Version
RHEL/CentOS	7.x, 8.x
Debian	9
Ubuntu	18.04, 20.04

JDK support

JDK	Version
OpenJDK	JDK8, JDK11
Oracle JDK	JDK8, JDK11

Supported databases

Database	Version
PostgreSQL	10.x, 11.x, 12.x, 13.x
MySQL	8.0.x
MariaDB	10.2, 10.3, 10.4, 10.5, 10.6

Browser support

Browser	Version
Chrome	>=76.0 (latest version 80.0 at time of release)
Firefox	>=68.0 (latest version 72.0 at time of release)

Other support

Package	Version
Stunnel	5.x

Install Java for EFM

You should install Java on the machine where you will install the EFM server.

Procedure

1. Download JDK from the appropriate website.

2. Run the installation command appropriate for your operating system:

For RHEL/CentOS:

```
yum install java-1.8.0-openjdk
```

For Debian and Ubuntu:

```
apt-get install openjdk-8-jre
```

Related Information

[OpenJDK Download](#)

[Oracle JDK Download](#)

Installing the EFM server

Learn how to install the Edge Flow Manager (EFM) server. You need to install database, and then install and configure the EFM server.

Installing databases for EFM

Learn how to install a database and configure it with your Edge Flow Manager (EFM) to manage, control, and monitor dataflows.

The EFM server requires a relational database. An H2 database is bundled with EFM that is used by default. While this is fine for development and test environments, for production environments, an external database is recommended.

For an external database, you can use either MySQL, MariaDB, or PostgreSQL. These topics describe how to install and configure MySQL, PostgreSQL, and MariaDB.



Note:

You should install either PostgreSQL or MySQL or MariaDB; all are not necessary.

Related Information

[System requirements for EFM](#)

Installing and configuring MySQL

Learn how to install and configure MySQL for Edge Flow Manager (EFM).

For supported database versions, see *System Requirements for EFM*.

Install MySQL

If you want to use PostgreSQL or MariaDB instead of MySQL, you may skip these steps. See the instructions for PostgreSQL and MariaDB in the respective sections.

1. Log in to the machine on which you want to install MySQL to use for the EFM server.
2. Install MySQL and the MySQL community server, and start the MySQL service. For installation instructions, check <https://dev.mysql.com/doc/refman/5.7/en/installing.html>.
3. Obtain the randomly generated MySQL root password.

```
grep 'A temporary password is generated for root@localhost' \  
/var/log/mysqld.log |tail -1
```

4. Reset the MySQL root password. Enter the following command. You are prompted for the password you obtained in the previous step. MySQL then asks you to change the password.

```
/usr/bin/mysql_secure_installation
```

5. Download the MySQL JDBC Connector and place it in the EFM lib directory: /path/to/efm-1.4.0/lib/.



Note: The folder is created when EFM is being installed.

Download the MySQL database driver from <https://dev.mysql.com/downloads/connector/j/>.



Note: The above link redirects you to the latest version of the driver. You need to choose the right driver version that fits your DB setup.

It is recommended to select the Platform Independent offering, download one of the archives, extract, and then copy the connector JAR to the lib directory.

Configure MySQL for use by EFM

1. Launch the MySQL shell:

```
mysql -u root -p
```

2. Create the database for the EFM service to use:

```
CREATE DATABASE efm CHARACTER SET latin1;
```

3. Create the efm user account, replacing the final IDENTIFIED BY string with your password:

```
CREATE USER 'efm'@'%' IDENTIFIED BY 'efmPassword';
```

4. Assign privileges to the efm account:

```
GRANT ALL PRIVILEGES ON efm.* TO 'efm'@'%' ;
```

5. Commit the operation:

```
FLUSH PRIVILEGES;
```

Configure the EFM database properties

1. Configure the database properties in the efm.properties file:

```
efm.db.url=jdbc:mysql://localhost:3306/efm
efm.db.driverClass=com.mysql.cj.jdbc.Driver
efm.db.username=efm
efm.db.password=efmPassword
```

The URL should match the host and port of the machine running MySQL. Username and password at the DB side and in the efm.properties file must match.

Installing and configuring PostgreSQL

Learn how to install and configure PostgreSQL for Edge Flow Manager (EFM).

For supported database versions, see *System Requirements for EFM*.

Install PostgreSQL

If you are using MySQL or MariaDB instead of PostgreSQL, you may skip these steps. See the instructions for MySQL and MariaDB in the respective sections.

1. Install Red Hat Package Manager (RPM) according to the requirements of your operating system:

```
sudo yum install -y https://download.postgresql.org/pub/repos/yum/reporpms/EL-7-x86_64/pgdg-redhat-repo-latest.noarch.rpm
```

2. Install PostgreSQL version 9.6:

```
yum install postgresql96-server postgresql96-contrib postgresql96
```

3. Initialize the database:

For example, if you are using CentOS 7, use the following syntax:

```
/usr/pgsql-9.6/bin/postgresql96-setup initdb
```

4. Start PostgreSQL.

For example, if you are using CentOS 7, use the following syntax:

```
systemctl enable postgresql-9.6.service  
systemctl start postgresql-9.6.service
```

5. Verify that you can log in:

```
sudo su postgres  
psql
```

Configure PostgreSQL to allow remote connections

It is critical that you configure PostgreSQL to allow remote connections before you deploy EFM. If you do not perform these steps in advance of installing the EFM server, the installation fails.

1. Open `/var/lib/pgsql/9.6/data/pg_hba.conf` and update to the following:

```
# "local" is for Unix domain socket connections only  
local all all trust  
  
# IPv4 local connections:  
host all all 0.0.0.0/0 trust  
  
# IPv6 local connections:  
host all all ::/0 trust
```

2. Open `/var/lib/pgsql/9.6/data/postgresql.conf` and update to the following:

```
listen_addresses = '*'
```

3. Restart PostgreSQL.

```
systemctl stop postgresql-9.6.service  
systemctl start postgresql-9.6.service
```

Configure PostgreSQL for use by EFM

1. Log in to PostgreSQL:

```
sudo su postgres
```

```
psql
```

2. Create a database for the EFM service to use:

```
create database efm;
CREATE USER efm WITH PASSWORD 'efmPassword';
GRANT ALL PRIVILEGES ON DATABASE "efm" to efm;
```

Configure the EFM database properties

1. Configure the database properties in the efm.properties file:

```
efm.db.url=jdbc:postgresql://localhost:5432/efm
efm.db.driverClass=org.postgresql.Driver
efm.db.username=efm
efm.db.password=efmPassword
```

The URL should match the host and port of the machine running PostgreSQL. The password should match the value that you set using the following command:

```
CREATE USER efm WITH PASSWORD 'efmPassword';
```

Installing and configuring MariaDB

Learn how to install and configure MariaDB for Edge Flow Manager (EFM). MariaDB should act as a drop-in binary for MySQL so configuration should be very similar.

For supported database versions, see *System Requirements for EFM*.

Install MariaDB

If you are using MySQL or PostgreSQL instead of MariaDB, you may skip these steps. See the instructions for MySQL and PostgreSQL in the respective sections.

1. Log in to the machine on which you want to install MariaDB to use for the EFM server.
2. Install MariaDB and the MariaDB server, and start the MariaDB service.

For install instructions, check <https://mariadb.com/kb/en/getting-installing-and-upgrading-mariadb/>.

3. Use the following command to reset the MariaDB root password:

```
/usr/bin/mysql_secure_installation
```

MariaDB then asks you to change the password.

4. Download the MariaDB JDBC connector and place it in the EFM lib directory:

```
/path/to/efm-1.4.0/lib/
```



Note: The folder is created when EFM is being installed.

5. Download the MariaDB database driver from <https://mariadb.com/kb/en/about-mariadb-connector-j/>.



Note: The above link redirects you to the latest version of the driver. You need to choose the right driver version that fits your DB setup.



Note: Cloudera recommends to select the platform independent offering, download one of the archives, extract, and then copy the connector JAR to the lib directory.

Configure MariaDB for use by EFM

1. Launch the MySQL shell:

```
mysql -u root -p
```

2. Create the database for the EFM service to use:

```
CREATE DATABASE efm CHARACTER SET latin1;
```

3. Create the efm user account, replacing the final IDENTIFIED BY string with your password:

```
CREATE USER 'efm'@'%' IDENTIFIED BY 'efmPassword';
```

4. Assign privileges to the efm account:

```
GRANT ALL PRIVILEGES ON efm.* TO 'efm'@'%' ;
```

5. Commit the operation:

```
FLUSH PRIVILEGES;
```

Configure the EFM database properties

1. Configure the database properties in the efm.properties file:

```
efm.db.url=jdbc:mariadb://localhost:3306/efm?useMysqlMetadata=true  
efm.db.driverClass=org.mariadb.jdbc.Driver  
efm.db.username=efm  
efm.db.password=efmPassword
```

The URL should match the host and port of the machine running MariaDB. The password should match the value that you set using the following command:

```
CREATE USER 'efm'@'%' IDENTIFIED BY 'efmPassword';
```

Installing EFM server

Learn how to install the Edge Flow Manager (EFM) server.

Procedure

1. Move the efm-version-bin.tar.gz tarball to the desired host and installation directory.

For download locations, see *EFM download locations*.

2. Extract the EFM tarball in the desired installation directory:

```
tar -xzf efm-version-bin.tar.gz
```



Important: You should not install EFM server on HDF nodes, more precisely, where net.ipv4.tcp_tw_recycle=1 is set, as it is not compatible with NAT and Load Balancers.

Related Information

[EFM download locations](#)

Installing EFM as an operating system service

The Edge Flow Manager (EFM) executable supports installation as a service on most Linux distributions. This is an optional installation step that is not required if you prefer to start the EFM server from the `efm.sh` executable included in the EFM bin directory.

You can start the application as a service by using either `init.d` or `systemd`.

Install EFM as an `init.d` service

To install EFM as an `init.d` service, symlink `bin/efm.sh` to `init.d`.

```
$ sudo ln -s /path/to/efm/bin/efm.sh /etc/init.d/efm
```

Once installed, you can start and stop the service as you would other OS services. For example:

```
$ service efm start
```

To configure EFM to start automatically on system boot, use `update-rc.d`. See `man update-rc.d` for information on using this utility.



Note: The EFM application runs as the user who owns the `efm.sh` launch script. It is recommended to never run as root. The recommended best practice is to create a specific user for running `efm`. Then use `chown` to make that user the owner of `efm.sh`. For example:

```
$ chown efm:efm /path/to/efm/bin/efm.sh
```

It is also recommended to use Unix or Linux filesystem permissions in order to secure the EFM installation. The rule of setting minimal access permissions applies. All files in the EFM installation should only be accessible to the EFM run-as user. Configuration files should be made read-only (for example, `chmod 400 <file>`). Executable files, such as those in the `bin` directory, should be made read and executable only (for example, `chmod 500 <file>`). Directories in the EFM install location should be readable and writable to the EFM user (for example, `chmod 600 <dir>`).

Install EFM as a `systemd` service

Most modern Linux distributions now use `systemd` as the successor to `init.d` (System V). In many cases you can continue to use `init.d`, but it is also possible to launch EFM using `systemd` as a service configuration.

To install EFM as a `systemd` service, create a file named `efm.service` in the `/etc/systemd/system` directory. For example:

```
[Unit]
Description=efm
After=syslog.target

[Service]
User=efm
ExecStart=/path/to/efm/bin/efm.sh
SuccessExitStatus=143

[Install]
WantedBy=multi-user.target
```



Note: When using `systemd`, the run-as user, the PID file, and the console log file are managed by `systemd` and therefore must be configured by using appropriate fields in the service script. Consult the service unit configuration man page for more details.

To configure EFM to start automatically on system boot, use `systemctl`. See `man systemctl` for information on using this utility.

Configuring the EFM server

Once the Edge Flow Manager (EFM) server is installed, you can configure it by editing the `efm.properties` file. At minimum, you should edit the EFM server address and configure the connection to your database.

Procedure

1. Open the `efm.properties` file located in `$EFM_HOME/conf/efm.properties`.



Note: The EFM home directory is the root directory where you installed the EFM binary.

2. Configure the EFM server address. Change `efm.server.address=localhost` to `efm.server.address={EFM_IP_OR_HOSTNAME}`, or use `0.0.0.0` to listen on all network interfaces.

Set the encryption password for EFM

You need to set the `efm.encryption.password` property which specifies a master password used for encrypting sensitive data saved to the Edge Flow Manager (EFM) server.

You can set it through the `efm.properties` file, a command line argument, or an OS environment variable.

By default, the EFM application uses AES encryption. The encryption key used is deterministically derived from an encryption password that the admin user must provide to the application at runtime. The property that is read for the encryption password is `efm.encryption.password`. You can set the value for this property in following ways:

- As a command line argument: `./bin/efm.sh --efm.encryption.password=myEfmPassword`
- As a Java System Property: `-Defm.encryption.password=myEfmPassword`
- As an OS environment variable: `export EFM_ENCRYPTION_PASSWORD=myEfmPassword`
- As a key/value pair in the `efm.properties` file: `efm.encryption.password=myEfmPassword`



Note: The master encryption password must be at least 12 characters long. It must be the same for all EFM instances.

The derived encryption key length is determined by your Java Runtime Environment encryption strength profiles.

- Unlimited Strength Encryption active: AES 256-bit key
- Unlimited Strength Encryption inactive: AES 128-bit key

It is strongly recommended that you enable Unlimited Strength Encryption in your Java Runtime Environment.

Open network ports for EFM

You should ensure that the required ports are available for the Edge Flow Manager (EFM) server and its components.

Component	Port number
EFM Server HTTP	10090
EFM Server Constraint application protocol (CoAP)	8989

Starting the EFM server

After you install and configure the Edge Flow Manager (EFM) server, you can start it. Learn how to start the EFM server.

Procedure

1. From the EFM server home directory, run:

```
bin/efm.sh start
```

2. Access the UI by browsing to the following location:

```
http://{EFM_HOST_OR_IP}:10090/efm/
```