

Cloudera Edge Management Installation

Date published: 2019-04-15

Date modified: 2025-02-28



Legal Notice

© Cloudera Inc. 2025. 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

Installing standalone Cloudera Edge Management.....	4
Before you begin installing Edge Flow Manager.....	4
System requirements for Edge Flow Manager.....	4
Installing JDK.....	4
Installing Edge Flow Manager.....	5
Installing databases for Edge Flow Manager.....	5
Installing the Edge Flow Manager server.....	9
Configuring the Edge Flow Manager server.....	11
Setting the encryption password for Edge Flow Manager.....	11
Open network ports for Edge Flow Manager.....	11
Starting Edge Flow Manager.....	12
 Installing Cloudera Edge Management on Cloudera Base on premises.....	 12
Preparing your database and Cloudera Base on premises cluster.....	13
Installing JDK.....	13
Installing and configuring a database for Edge Flow Manager.....	13
Installing Cloudera Manager and Cloudera Base on premises.....	18
Installing Cloudera Edge Management parcel and CSD files.....	19
Installing Edge Flow Manager on your Base cluster.....	21

Installing standalone Cloudera Edge Management

Get started with Edge Flow Manager by installing, configuring, and securing it. Learn about the system requirements, and how to obtain the Edge Flow Manager software package, install the application, configure its settings, and ensure that security measures are in place.

Before you begin installing Edge Flow Manager

Before you start the Edge Flow Manager installation, review the system requirements. Understand the supported operating systems, Java Development Kit (JDK) versions, and browser specifications, and learn how to install the required JDK to optimize Edge Flow Manager functionality.

System requirements for Edge Flow Manager

Operating system support

Operating System	Version
RHEL/CentOS	8.x, 9.x
Ubuntu	20.04, 22.04, 24.04

JDK support

JDK	Version
OpenJDK	JDK 17
Oracle JDK	JDK 17

Supported databases

Database	Version
PostgreSQL	12.x, 13.x, 14.x, 15.x, 16.x
MySQL	8.0.x
MariaDB	10.4, 10.5, 10.6, 10.11

Browser support

Browser	Version
Chrome	The latest version at the time of the last Cloudera Edge Management release, and the two previous major versions
Firefox	The latest version at the time of the last Cloudera Edge Management release, and the two previous major versions

Installing JDK

About this task

To ensure that Cloudera Edge Management functions correctly, you must install JDK on each machine where the Edge Flow Manager server will be installed.



Note: Only Java 17 is supported. Older or newer versions are not guaranteed to work and may cause startup failures.

Procedure

1. Install Java 17.

- Download the JDK package from the appropriate website.
- Install JDK on your system by running the following command.

For RHEL/CentOS: `yum install java-17-openjdk`

2. Set the JAVA_HOME environment variable.

In environments with multiple Java versions, it is important to set the correct JAVA_HOME to avoid runtime issues.

- To set JAVA_HOME permanently, add the following to your shell profile.

(For example: `~/.bashrc` or `/etc/profile.d/java.sh`)

```
export JAVA_HOME=/usr/lib/jvm/java-17-openjdk
export PATH=$JAVA_HOME/bin:$PATH
```

- Reload the environment or restart the shell session.

```
source ~/.bashrc
```

3. Verify the Java version to ensure that the correct Java version is in use before starting Edge Flow Manager.

```
java -version
```

The output should show Java 17.

What to do next

After installing the JDK, ensure that you have the required database installed and properly configured for Cloudera Edge Management.

Related Information

[OpenJDK Download](#)

[Oracle JDK Download](#)

Installing Edge Flow Manager

Learn how to install the Edge Flow Manager. You need to install a database first, and then install and configure the Edge Flow Manager server.

Installing databases for Edge Flow Manager

Learn how to install a database and configure it to seamlessly integrate with your Edge Flow Manager for efficient management, control, and monitoring of dataflows.

The Edge Flow Manager server requires a relational database. A default H2 database is embedded in Edge Flow Manager, which is suitable for development and testing, it is recommended to use an external database for production environments. You can use either MySQL, MariaDB, or PostgreSQL for your external database solution.



Note:

You should install either PostgreSQL, MySQL, or MariaDB. Installing all three is not necessary.

Installing and configuring MySQL

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

For supported database versions, see *System Requirements for Edge Flow Manager*.

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 Edge Flow Manager server.
2. Install MySQL and the MySQL community server, and start the MySQL service. For installation instructions, check [Installing and Upgrading MySQL](#).
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 Edge Flow Manager lib directory: /path/to/efm-1.4.0/lib/.



Note: The folder is created when Edge Flow Manager is being installed.

Download the MySQL database driver from [MySQL Community Downloads](#).



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 Edge Flow Manager

1. Launch the MySQL shell:

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

2. Create the database for the Edge Flow Manager 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 Edge Flow Manager 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.

For supported database versions, see *System Requirements for Edge Flow Manager*.

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 Edge Flow Manager. If you do not perform these steps in advance of installing the Edge Flow Manager 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 Edge Flow Manager

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 Edge Flow Manager 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. MariaDB should act as a drop-in binary for MySQL, so its configuration should be very similar.

For more information about supported database versions, see *System Requirements for Edge Flow Manager*.

Install MariaDB

If you are using MySQL or PostgreSQL instead of MariaDB, you can skip these steps and use the MySQL and PostgreSQL instructions in their respective sections.

1. Log in to the machine where you want to install MariaDB for the Edge Flow Manager server.
2. Install MariaDB.

- a. Install MariaDB and the MariaDB server.
- b. Start the MariaDB service.

For detailed installation instructions, see [Getting, Installing, and Upgrading MariaDB](#).

3. Reset the MariaDB root password using the following command.

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

MariaDB prompts you to set a new password.

4. Download the MariaDB JDBC connector driver from the [MariaDB Connector/J documentation](#).



Note: The documentation link redirects to the latest version. Choose the driver version compatible with your database setup.

Cloudera recommends selecting the platform-independent archive. Extract the archive and copy the connector JAR to the Edge Flow Manager lib directory.

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



Note: The lib directory is created during Edge Flow Manager installation.

Configure MariaDB for use by Edge Flow Manager

1. Launch the MySQL shell.

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

2. Create the database for the Edge Flow Manager 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 changes.

```
FLUSH PRIVILEGES;
```

Configure the Edge Flow Manager 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 the Edge Flow Manager server

Learn how to install the Edge Flow Manager server.

Procedure

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

For download locations, see *Edge Flow Manager download locations*.

2. Extract the Edge Flow Manager tarball in the desired installation directory:

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



Important: You should not install Edge Flow Manager 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

[Edge Flow Manager download locations](#)

Installing Edge Flow Manager as an operating system service

The Edge Flow Manager 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 Edge Flow Manager server from the `efm.sh` executable included in the Edge Flow Manager bin directory.

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

Install Edge Flow Manager as an `init.d` service

To install Edge Flow Manager 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 Edge Flow Manager to start automatically on system boot, use `update-rc.d`. See `man update-rc.d` for information on using this utility.



Note: The Edge Flow Manager 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 Edge Flow Manager. 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 Edge Flow Manager installation. The rule of setting minimal access permissions applies. All files in the Edge Flow Manager installation should only be accessible to the Edge Flow Manager 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 Edge Flow Manager install location should be readable and writable to the Edge Flow Manager user (for example, `chmod 600 <dir>`).

Install Edge Flow Manager 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 Edge Flow Manager using `systemd` as a service configuration.

To install Edge Flow Manager 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 Edge Flow Manager to start automatically on system boot, use systemctl. See man systemctl for information on using this utility.

Configuring the Edge Flow Manager server

Once the Edge Flow Manager server is installed, you can configure it by editing the `efm.properties` file. At minimum, you should edit the Edge Flow Manager 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 Edge Flow Manager home directory is the root directory where you installed the Edge Flow Manager binary.

2. Configure the Edge Flow Manager 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.

Setting the encryption password for Edge Flow Manager

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 server.

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

By default, the Edge Flow Manager 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 Edge Flow Manager 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 Edge Flow Manager

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

Component	Port number
Edge Flow Manager Server HTTP	10090

Starting Edge Flow Manager

Learn how you can start the Edge Flow Manager server after after you installed and configured it.

Procedure

1. From the Edge Flow Manager server home directory, run the following command.

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

2. Access the UI at the following location.

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

Installing Cloudera Edge Management on Cloudera Base on premises

Before you start installing Cloudera Edge Management on a Cloudera Base on premises cluster, it is helpful to understand the steps of the installation process.

The Cloudera Edge Management installation workflow consists of preparing your Cloudera Base on premises cluster and then installing Edge Flow Manager on the Base cluster.

Installing Edge Flow Manager on your Base cluster is a simplified deployment during which you create only one cluster. You can use this set-up if you are implementing trial deployments or in simple production scenarios.



Important:

Running multiple Edge Flow Manager instances on the same node is not recommended by Cloudera.

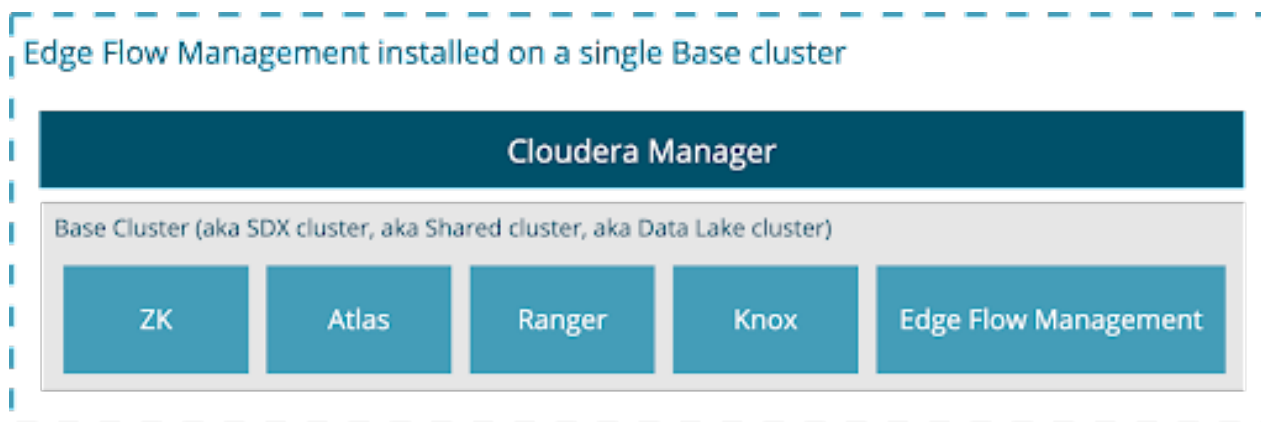
A Base cluster is a large cluster that contains SDX services like Ranger and Atlas, along with compute and storage services like Edge Flow Manager, Impala, and Hive. It is sometimes called a Regular cluster in Cloudera Manager. Additional synonymous terms are SDX cluster, Shared cluster, or Data Lake cluster.



Important:

Cloudera recommends that you always install Edge Flow Manager on the Base cluster.

After a successful installation on your Base cluster, your Cloudera Edge Management cluster layout will look similar to the following:



Preparing your database and Cloudera Base on premises cluster

Before you start installing Edge Flow Manager, you need to prepare your Cloudera Base on premises cluster and a database for the service.

Installing JDK

About this task

To ensure that Cloudera Edge Management functions correctly, you must install JDK on each machine where the Edge Flow Manager server will be installed.



Note: Only Java 17 is supported. Older or newer versions are not guaranteed to work and may cause startup failures.

Procedure

1. Install Java 17.

- a) Download the JDK package from the appropriate website.
- b) Install JDK on your system by running the following command.

For RHEL/CentOS: `yum install java-17-openjdk`

2. Set the JAVA_HOME environment variable.

In environments with multiple Java versions, it is important to set the correct JAVA_HOME to avoid runtime issues.

- a) To set JAVA_HOME permanently, add the following to your shell profile.

(For example: `~/bashrc` or `/etc/profile.d/java.sh`)

```
export JAVA_HOME=/usr/lib/jvm/java-17-openjdk
export PATH=$JAVA_HOME/bin:$PATH
```

- b) Reload the environment or restart the shell session.

```
source ~/.bashrc
```

3. Verify the Java version to ensure that the correct Java version is in use before starting Edge Flow Manager.

```
java -version
```

The output should show Java 17.

What to do next

After installing the JDK, ensure that you have the required database installed and properly configured for Cloudera Edge Management.

Related Information

[OpenJDK Download](#)

[Oracle JDK Download](#)

Installing and configuring a database for Edge Flow Manager

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

By default, Edge Flow Manager is preconfigured with an embedded H2 database. But when setting up Edge Flow Management, you have the option to install and configure an external MySQL or PostgreSQL database to use with Edge Flow Manager.

The H2 database that is preconfigured in the default `efm.properties` file. The contents of the embedded H2 database are stored in a file on your local file system. The H2 database location is specified by the JDBC URL property: Cloudera Edge Management JDBC URL (`efm.db.url`) – `jdbc:h2:./database/efm`

If you plan to use the embedded H2 database, you can skip the steps for installing and configuring an external database. Alternatively, you can install an external MySQL or PostgreSQL database. To proceed with external database installation, review the following steps.

For a list of supported MySQL or PostgreSQL databases, see *System requirements for Edge Flow Manager*.

Related Information

[System requirements for Edge Flow Manager](#)

Installing MySQL

Learn how to install a MySQL external database to support Cloudera Edge Management. If you have already installed a PostgreSQL or an Oracle database, you can skip these steps, as only one of the databases is required.

Before you begin

Ensure that you have installed the JDK. For installation instructions, see [Installing the JDK](#).

Procedure

1. Log in to the node where you intend to install Edge Flow Manager.
2. Install MySQL and the MySQL community server, and start the MySQL service.

```
yum localinstall https://dev.mysql.com/get/mysql[***MYSQL_VERSION***]-community-release-el[***RHEL_VERSION***]-[***PACKAGE_VERSION***].noarch.rpm
yum install mysql-community-server
systemctl start mysqld.service
```

Replace the placeholders with the appropriate values for your specific use case:

- `[***MYSQL_VERSION***]`: Selected MySQL version
- `[***RHEL_VERSION***]`: Red Hat Enterprise Linux (RHEL) version you are using
- `[***PACKAGE_VERSION***]`: Specific package version

For a list of supported MySQL databases in Cloudera Edge Management, *System requirements for Edge Flow Manager*.

For example, if you wanted to install the MySQL 8.0 Community Edition repository configuration package on a RHEL 8 system, you would use:

```
yum localinstall https://dev.mysql.com/get/mysql80-community-release-el8-8.noarch.rpm
```

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.

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

You are prompted for the password you obtained in the previous step. MySQL will request you to change the password.

What to do next

Once you have completed the MySQL installation, configure it for use with Edge Flow Manager. For instructions, see [Configuring MySQL](#).

Related Information

[System requirements for Edge Flow Manager](#)

[Configuring MySQL](#)

Configuring MySQL

MySQL offers the option to use an externally located database that supports high availability. Learn how to set up a MySQL database for use with the Edge Flow Manager.

Before you begin

Ensure that you have installed a MySQL database. For instructions, see *Installing MySQL*.

Procedure

1. Download the MySQL JDBC driver and save it in an accessible location for later use.

```
/path/to/drivers/mysql-connector-java-8.0.16.jar
```

2. Create a database in MySQL.

a) Access the MySQL shell using the following command: `mysql -u root -p`

b) Run the following command to create a database:

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

3. Create a database user and grant privileges.

For remote hosts, specify the appropriate IP address using `nifireg'@'<IP-ADDRESS>` or `nifireg'@'%'`

```
CREATE USER 'efm'@'%' IDENTIFIED BY 'efmPassword';
GRANT ALL PRIVILEGES ON efm.* TO 'efm'@'%' ;
FLUSH PRIVILEGES;
```

4. After the Cloudera Edge Managementservice is installed, copy the MySQL driver JAR file saved in Step 1 to `/opt/cloudera/parcels/CEM/EFM/lib`
5. Configure the following database properties in Cloudera Manager.
 - Cloudera Edge Management JDBC URL (efm.db.url) – `jdbc:mysql://<MYSQL-HOSTNAME>:<MYSQL-PORT>/efm`
 - Cloudera Edge Management JDBC Driver (efm.db.driverClass) – `com.mysql.cj.jdbc.Driver`
 - Cloudera Edge Management Database Username (efm.db.username) – `efm`
 - Cloudera Edge Management Password (efm.db.password) – `efmPassword`

What to do next

When you have completed the Cloudera Edge Management database configuration, proceed to installing Cloudera Manager and your Cloudera Base on premises cluster.

Related Information

[Installing MySQL](#)

Installing PostgreSQL

Learn how to install a PostgreSQL external database to support Cloudera Edge Management. If you have already installed a MySQL or an Oracle database, you can skip these steps. Only one of the databases is required.

Before you begin

Ensure that you have installed the JDK. For installation instructions, see [Installing the JDK](#).

Procedure

1. Install Red Hat Package Manager (RPM) according to your operating system's requirements.

```
yum install https://yum.postgresql.org/[***POSTGRESQL_VERSION***]/redhat/rhel-[***RHEL_VERSION***]-x86_64/[***REPOSITORY_PACKAGE_NAME***]
```

Replace the placeholders with the appropriate values for your specific use case:

- [***POSTGRESQL_VERSION***]: Desired PostgreSQL version
- [***RHEL_VERSION***]: Red Hat Enterprise Linux (RHEL) version you are using
- [***REPOSITORY_PACKAGE_NAME***]: Name of the repository configuration package for the specified PostgreSQL version and RHEL version.

For a list of supported PostgreSQL databases in Cloudera Edge Management, *System requirements for Edge Flow Manager*.

For example, to add the PostgreSQL 12 repository to a RHEL 8 system, you would use:

```
yum install https://yum.postgresql.org/12/redhat/rhel-8-x86_64/pgdg-redhat12-12-4.noarch.rpm
```

2. Install the PostgreSQL version of your choice.

```
yum install postgresql[***POSTGRESQL_VERSION***]-server postgresql[***POSTGRESQL_VERSION***]-contrib postgresql[***POSTGRESQL_VERSION***]
```

For example, if you wanted to add the PostgreSQL 12, you would use:

```
yum install postgresql12-server postgresql12-contrib postgresql12
```

3. Initialize the database.

For CentOS 7, use the following syntax:

```
[***POSTGRESQL_BIN_PATH***]/postgresql[***POSTGRESQL_VERSION***]-setup initdb
```

For example, to initialize PostgreSQL 12, you would use:

```
/usr/pgsql-12/bin/postgresql12-setup initdb
```

4. Enable and start PostgreSQL.

```
systemctl enable postgresql-[***POSTGRESQL_VERSION***].service systemctl start postgresql-[***POSTGRESQL_VERSION***].service
```

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

```
systemctl enable postgresql-12.service systemctl start postgresql-12.service
```

5. Verify that you can log in.

```
sudo su postgres  
psql
```


What to do next

When PostgreSQL is installed, configure the database for use with Cloudera Edge Management. For instructions, see [Configuring PostgreSQL](#).

Related Information

[System requirements for Edge Flow Manager](#)

[Configuring PostgreSQL](#)

Configuring PostgreSQL

PostgreSQL offers the option to use an externally located database that supports high availability. Learn how to set up a MySQL database for use with the Edge Flow Manager.

Before you begin

Ensure that you have installed a PostgreSQL database. For instructions, see *Installing PostgreSQL*.

Procedure

1. Download the Postgres JDBC driver and save it in an accessible location for later use.

```
/path/to/drivers/postgresql-42.2.2.jar
```

2. Create a database in PostgreSQL using createdb efm.
3. Create a database user and grant privileges.

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

4. After Cloudera Edge Management service is installed, copy the PostgreSQL driver JAR file saved in Step 1 to /opt/cloudera/parcels/CEM/EFM/lib
5. Configure the following database properties in Cloudera Manager.
 - Cloudera Edge Management JDBC URL (efm.db.url) – jdbc:mysql://<POSTGRES-HOSTNAME>:<POSTGRES-PORT>/efm
 - Cloudera Edge Management JDBC Driver (efm.db.driverClass) – org.postgresql.Driver
 - Cloudera Edge Management Database Username (efm.db.username) – efm
 - Cloudera Edge Management Password (efm.db.password) – efmPassword

What to do next

When you have completed the Cloudera Edge Management database configuration, proceed to installing Cloudera Manager and your Cloudera Base on premises cluster.

Related Information

[Installing PostgreSQL](#)

Installing Oracle

Learn how to install an external Oracle database to support Cloudera Edge Management.

Follow the Oracle documentation to install the database. This step is required only if you do not have the Oracle database installed in your data center or environment. If you are using an existing Oracle database, you can skip this step. If you have already installed a PostgreSQL or a MySQL database, you can also skip these steps. Only one of the databases is required.

For a list of supported Oracle databases in Cloudera Edge Management, see *System requirements for Edge Flow Manager*.

What to do next

Once you have installed Oracle, configure the database for use with Cloudera Edge Management.

Related Information

[System requirements for Edge Flow Manager](#)

[Configuring Oracle](#)

Configuring Oracle

Learn how to configure a Oracle database to integrate with Edge Flow Manager.

About this task

Oracle provides the option to use an externally located database that supports high availability.

Before you begin

You have installed an Oracle database. For details, see *Installing Oracle*.

Procedure

1. Download the [Oracle JDBC connector](#) and save it somewhere accessible for later use.

```
/path/to/drivers/ojdbc11.jar
```

2. Open Oracle SQL*Plus or SQLcl using OS-level authentication.

```
sqlplus / as sysdba
```

3. Create a database user and grant privileges.

```
CREATE USER efm IDENTIFIED BY efmPassword;  
GRANT CONNECT, RESOURCE, CREATE VIEW TO efm;  
COMMIT;
```

4. After the Cloudera Edge Management service is installed, copy the Oracle driver jar file saved in Step 1 to /opt/cloudera/parcels/CEM/EFM/lib.
5. Configure the database properties in Cloudera Manager.
 - Cloudera Edge Management JDBC Url (efm.db.url) – jdbc:oracle:thin://<ORACLE-HOSTNAME>:<ORACLE-PORT>/efm
 - Cloudera Edge Management JDBC Driver (efm.db.driverClass) – oracle.jdbc.OracleDriver
 - Cloudera Edge Management Database Username (efm.db.username) – efm
 - Cloudera Edge Management Password (efm.db.password) – efmPassword

What to do next

When you have completed the Cloudera Edge Management database configuration, proceed to installing Cloudera Manager and your Cloudera Base on premises cluster.

Related Information

[Installing Oracle](#)

Installing Cloudera Manager and Cloudera Base on premises

Before installing Cloudera Edge Management, you need to set up Cloudera Base on premises with Cloudera Manager.

Before you begin

- Make sure you have installed a JDK.

- If you intend to use an external database, ensure you have installed and configured it.

About this task

See the *Cloudera Base on premises Installation Guide* for comprehensive instructions on installing Cloudera Manager and a Cloudera Base on premises cluster. Ensure to perform the following minimum steps.

Procedure

1. Install Cloudera Base on premises.
2. Enable Auto-TLS and Kerberos.



Important: Unsecured Cloudera Edge Management is not supported.

For details on security options and recommendations, see the *Cloudera Edge Management Security* documentation.



Note:

- Wildcard certificates are not supported.
- If Cloudera Edge Management operates behind Knox, do not use wildcard certificates for Knox.

For more information on certificate requirements, see *TLS certificate requirements and recommendations*.

As Cloudera Edge Management has no dependency on any other Cloudera Base on premises service, you do not need to install additional services.

What to do next

When the Cloudera Base on premises cluster installation is ready, add the Cloudera Edge Management parcel and CSD files.

Related Information

[Cloudera Edge Management security overview](#)

Installing Cloudera Edge Management parcel and CSD files

To set up Cloudera Edge Management on a Cloudera Base on premises cluster, you need to add the Cloudera Edge Management parcel to your cluster using Cloudera Manager, making the Cloudera Edge Management service available for installation. Additionally, you also need to upload the Cloudera Edge Management Custom Service Descriptor (CSD) files to the default CSD directory. These files contain crucial configuration information required to describe and manage the new service effectively.

Before you begin

- You have installed a JDK. For more information, see *Installing the JDK*.
- You have installed and configured a database for use with Edge Flow Management. For more information, see *Installing and configuring a database for Cloudera Edge Management*.
- You have installed a Cloudera Base on premises cluster. For more information, see the *Cloudera Base on premises Installation Guide*.

Procedure

1. Log into Cloudera Manager.
2. Navigate to the Parcels page by selecting Hosts Parcels .
3. Click Parcel Repository & Network Settings.

4. Add the Repository URL.

- a) On the Parcel Repository & Network Settings page, click + to add an additional row in the Remote Parcel Repository URLs list.
- b) Add the URL of the Cloudera Edge Management parcel repository where the parcel file that you intend to install is hosted.

**Note:**

When providing the URL, make sure that the actual parcel file name is not included in the string you add. Cloudera Manager will search for any parcel files located within the specified directory.

RHEL/CentOS 8

```
https://archive.cloudera.com/p/CEM/2.3.1.0/redhat8/yum/tars/parcel/
```

RHEL/CentOS 9

```
https://archive.cloudera.com/p/CEM/2.3.1.0/redhat9/yum/tars/parcel/
```

Ubuntu 20

```
https://archive.cloudera.com/p/CEM/2.3.1.0/ubuntu20/apt/tars/parcel/
```

See *Download locations* for a complete list of parcel links.

5. Click Save & Verify Configuration.

The new Cloudera Edge Management parcel is displayed on the Parcels page.

- 6.** Click Download next to Cloudera Edge Management to download the parcel to the local repository.
- 7.** When the download is completed, click Distribute to distribute the parcel to all clusters.
- 8.** Download the CSD files from the Cloudera Edge Management repository and place them in the default CSD directory: `/opt/cloudera/csd`

For more information, see *Download locations*.

**Note:**

Ensure that the CSD version you download is the same as the Cloudera Edge Management version. Different CSD and Cloudera Edge Management versions can cause problems when configuring and managing Cloudera Edge Management through Cloudera Manager.

9. Change the ownership of the CSD files.

In `/opt/cloudera/csd`, enter:

```
chown cloudera-scm:cloudera-scm /opt/cloudera/csd/*
chmod 644 /opt/cloudera/csd/*
```

10. Restart the Cloudera Manager Server.

```
sudo service cloudera-scm-server restart
```

- 11.** Go back to Cloudera Manager, click **Clusters Cloudera Management Service** and click **Actions Restart** to restart the service .
- 12.** Navigate to the Parcels page by selecting **Hosts Parcels** .
- 13.** Click **Activate** next to Cloudera Edge Management in the parcel list to activate the Cloudera Edge Management parcel.
- 14.** Click **OK** when confirmation is required.

What to do next

You have completed your Cloudera Base on premises installation. You can now proceed with installing, securing, and configuring Cloudera Edge Management.

Installing Edge Flow Manager on your Base cluster

Learn how to add and configure the Cloudera Edge Management service on Cloudera Base on premises.


Before you begin

Ensure that you have successfully installed a Cloudera Base on premises cluster.

Procedure

1. Navigate to the Home Status tab and click Actions Add Service .

Install one service at a time.



2. Select the Edge Flow Management service and click Continue to proceed to the Add Service wizard.
3. Configure the Cloudera Edge Management dependencies as needed and click Continue.
Select Knox to establish a secure single entry point for accessing services.
4. Assign roles to Cloudera Edge Management by selecting the appropriate hosts where you want to install your new roles, and then click Continue to proceed.
5. Review the configuration changes and click Continue to start the service installation.
 - Provide a value for your Initial Admin Identity.
 - Fill in the database related properties as configured in *Installing and configuring a database for Cloudera Edge Management*.
 - Choose between LDAP and KNOX authentication methods according to your requirements. For more information, see the related section in *User authentication*.
 - Optionally, you can further customize your Cloudera Edge Management configuration here.
6. Click Continue and then Finish to complete the installation process.

Results

You have completed your Cloudera Edge Management installation. To access the Cloudera Edge Management UI, open a web browser and enter the following URL: `https://[***HOSTNAME***]:10090/efm/ui`

What to do next

Verify that the new service is added properly by checking its health status. If the health status is good, the service has been added correctly.