

How to Deploy Open Source Databases



Table of Contents

Introduction	4
Popular Vendors	5
Percona	5
MariaDB	5
Oracle MySQL	5
MongoDB	6
PostgreSQL	6
How to Deploy Open Source Databases	7
Percona Server for MySQL	7
Installation	7
Default Configuration	10
Optional Percona Server Configuration	12
Oracle MySQL Community Server	14
Installation	14
Default Configuration	18
Optional MySQL Community Server Configuration	19
Group Replication	21
MariaDB	24
Installation	24
Default Configuration	28
Optional MariaDB Server Configuration	31
MariaDB Cluster Configuration	32
Percona XtraDB Cluster	34
Installation	34
Default Configuration	37
Optional Percona Server Configuration	41
NDB Cluster	42
Installation	43
Default Configuration	48
Optional NDB Cluster Server Configuration	52
MongoDB	54
Installation	54
Default Configuration	55
Optional MongoDB Configuration	57
Percona Server for MongoDB	57
Installation	57
Default Configuration	59
Optional Percona MongoDB Configuration	61
PostgreSQL	61
Installation	61
Default Configuration	63
Synopsis	66



Table of Contents

How to Deploy Open Source Databases by Using ClusterControl	67
Deploy	67
Scaling	71
Load Balancing	72
Management	75
Conclusion	76
About ClusterControl	77
About Severalnines	77
Related Resources	78



Introduction

Choosing which DB engine to use between all the options we have today is not an easy task. And that is just the beginning. After deciding which engine to use, you need to learn about it and actually deploy it to play with it. We plan to help you on that second step, and show you how to install, configure and secure some of the most popular open source DB engines. In this whitepaper we are going to cover these points, with the aim of fast tracking you on the deploy task.

Popular Vendors

Let's take a look at some of the most popular open source database vendors, in no specific order.

Percona

Percona is a leading provider of unbiased open source database solutions that allow organizations to easily, securely and affordably maintain business agility, minimize risks, and stay competitive.

Percona is committed to producing open-source database software for MySQL and its variants, and develops their own version: Percona Server for MySQL, Percona XtraDB Cluster and Percona Server for MongoDB.

All Percona software is open source and free of charge.



MariaDB

MariaDB is a commercially supported fork of the MySQL relational database management system, intended to remain free and open-source software under the GNU GPL. The development of this engine is led by some of the original developers of MySQL, who forked it due to concerns over its acquisition by the Oracle Corporation.

MariaDB intends to maintain a compatibility with MySQL by using the same main storage engine. It includes the InnoDB storage engine, as well as a new storage engine, Aria, that intends to be both a transactional and non-transactional engine. There are some differences with MySQL to keep in mind like GTID, SQL syntax, encryption and tools compatibilities (e.g. xtrabackup).



Oracle MySQL

MySQL is a free and open-source RDBMS under the terms of the GNU General Public License, and is also available under a variety of proprietary licenses. MySQL was owned and developed by the Swedish company MySQL AB, which was bought by Sun Microsystems (now



Oracle Corporation). In 2010, when Oracle acquired Sun, Monty Widenius forked the open-source MySQL project to create MariaDB.

MySQL is a component of the LAMP web application software stack (and others), which is an acronym for Linux, Apache, MySQL, Perl/PHP/Python. MySQL is used by many database-driven web applications, including WordPress, Drupal, and phpBB.

MongoDB

MongoDB was founded in 2007 by Dwight Merriman, Eliot Horowitz and Kevin Ryan.

MongoDB is the leading modern, general purpose database platform, designed to unleash the power of software and data for developers and the applications they build. MongoDB has more than 6,600 customers in more than 100 countries. The MongoDB database platform has been downloaded over 40 million times and there have been more than 1 million MongoDB University registrations.



PostgreSQL

PostgreSQL is a powerful, open source object-relational database system that uses and extends the SQL language combined with many features that safely store and scale the most complicated data workloads. The origins of PostgreSQL date back to 1986 as part of the POSTGRES project at the University of California at Berkeley and has more than 30 years of active development on the core platform.

PostgreSQL is developed by the PostgreSQL Global Development Group, a diverse group of many companies and individual contributors. It is free and open-source, released under the terms of the PostgreSQL License, a permissive software license.



How to Deploy Open Source Databases

As we could see, when choosing your storage engine, there are different vendors. Let's see some of the most common open source database servers manual installations and configurations.

Note: In some cases, the installation or configuration could be similar and maybe could look a bit repetitive, but we have added it anyway to allow you to follow a specific database server without the need to read the entire document.

Percona Server for MySQL

Percona Server for MySQL is a fully compatible, enhanced and open source drop-in replacement for any MySQL database. It is trusted by thousands of enterprises to provide better performance and concurrency for their most demanding workloads, and delivers greater value to MySQL server users with optimized performance, greater performance scalability and availability, enhanced backups and increased visibility.



Installation

In [this link](#) you have the latest packages to install Percona Server for MySQL.

If you prefer, you can follow the [yum repository installation](#) or the [apt repository installation](#).

In our example, let's see the yum repository installation for a Percona Server on CentOS 7.

Percona Server IP Address: 192.168.100.117

```
1 [root@WP1 ~]# yum install http://www.percona.com/downloads/
  percona-release/redhat/0.1-6/percona-release-0.1-6.noarch.
  rpm
2 =====
  =====
3 Package                               Arch      Version
  Repository                             Size
4 =====
  =====
5 Installing:
6 percona-release                       noarch    0.1-6
```

```

7 /percona-release-0.1-6.noarch 16 k
8 Transaction Summary
9 =====
10 Install 1 Package
11
12 [root@WP1 ~]# yum install Percona-Server-server-57
13 =====
14 Package Arch Version
15 Repository Size
16 =====
17 Installing:
18 Percona-Server-server-57 x86_64 5.7.24-
26.1.el7 percona-release-x86_64 39 M
19 Installing for dependencies:
20 Percona-Server-client-57 x86_64 5.7.24-
26.1.el7 percona-release-x86_64 6.8 M
21 Percona-Server-shared-57 x86_64 5.7.24-
26.1.el7 percona-release-x86_64 748 k
22 Percona-Server-shared-compat-57 x86_64 5.7.24-
26.1.el7 percona-release-x86_64 1.2 M
23 libaio x86_64 0.3.109-13.el7
base 24 k
24 numactl-libs x86_64 2.0.9-7.el7
base 29 k
25 Transaction Summary
26 =====
27 Install 1 Package (+5 Dependent packages)

```

Now, you need to start the MySQL service.

```

1 [root@WP1 ~]# service mysql start
2 Redirecting to /bin/systemctl start mysql.service

```

A new password for the root user is created. To know it, you need to check the MySQL log and filter by temporary password. You should have something similar to this:

```

1 [root@WP1 ~]# grep "temporary password" /var/log/mysqld.log
2 2018-12-18T22:04:22.887873Z 1 [Note] A temporary password is
generated for root@localhost: W59c:%3/#sld

```

Then, you can run the `mysql_secure_installation` script, to configure a basic secure setup for your MySQL database.


```
1 [root@WP1 ~]# mysql_secure_installation
2 Securing the MySQL server deployment.
3
4 Enter password for user root:
5 The 'validate_password' plugin is installed on the server.
6 The subsequent steps will run with the existing configuration
7 of the plugin.
8 Using existing password for root.
9
10 Estimated strength of the password: 100
11 Change the password for root ? ((Press y|Y for Yes, any oth-
12 er key for No) : y
13
14 New password:
15 Re-enter new password:
16
17 Estimated strength of the password: 100
18 Do you wish to continue with the password provided?(Press
19 y|Y for Yes, any other key for No) : y
20
21 By default, a MySQL installation has an anonymous user,
22 allowing anyone to log into MySQL without having to have
23 a user account created for them. This is intended only for
24 testing, and to make the installation go a bit smoother.
25 You should remove them before moving into a production
26 environment.
27
28 Remove anonymous users? (Press y|Y for Yes, any other key
29 for No) : y
30 Success.
31
32 Normally, root should only be allowed to connect from
33 'localhost'. This ensures that someone cannot guess at
34 the root password from the network.
35
36 Disallow root login remotely? (Press y|Y for Yes, any other
37 key for No) : y
38 Success.
39
40 By default, MySQL comes with a database named 'test' that
41 anyone can access. This is also intended only for testing,
42 and should be removed before moving into a production
43 environment.
44
45 Remove test database and access to it? (Press y|Y for Yes,
46 any other key for No) : y
47 - Dropping test database...
48 Success.
```

```

45 | - Removing privileges on test database...
46 | Success.
47 |
48 | Reloading the privilege tables will ensure that all changes
49 | made so far will take effect immediately.
50 |
51 | Reload privilege tables now? (Press y|Y for Yes, any other
    | key for No) : y
52 | Success.
53 |
54 | All done!

```

Now your database is running, but it's not ready yet. It's just the beginning.

Default Configuration

By default, the Percona's my.cnf config file includes the /etc/my.cnf.d/ and /etc/percona-server.conf.d/ directories:

```

1 | [root@WP1 ~]# cat /etc/my.cnf
2 | #
3 | # The Percona Server 5.7 configuration file.
4 | #
5 | #
6 | # * IMPORTANT: Additional settings that can override those
    | from this file!
7 | #   The files must end with '.cnf', otherwise they'll be ig-
    | nored.
8 | #   Please make any edits and changes to the appropriate
    | sectional files
9 | #   included below.
10 | #
11 | !includedir /etc/my.cnf.d/
12 | !includedir /etc/percona-server.conf.d/

```

The /etc/my.cnf.d/ directory is empty by default, and in the /etc/percona-server.conf.d/ we have the following content:

```

1 | [root@WP1 ~]# ls /etc/percona-server.conf.d/
2 | mysql_d.cnf  mysql_d_safe.cnf

```

mysql_d_safe is the recommended way to start a mysqld server on Unix. It adds some safety features such as restarting the server when an error occurs and logging runtime information to an error log.

mysql_d_safe reads options from both [mysqld] and [mysqld_safe] sections in the configuration files.

The content of these configuration files are:

- mysqld.cnf

```
1 [root@WP1 ~]# cat /etc/percona-server.conf.d/mysqld.cnf
2 # Percona Server template configuration
3 [mysqld]
4 #
5 # Remove leading # and set to the amount of RAM for
6 # the most important data
7 # cache in MySQL. Start at 70% of total RAM for dedi-
8 # cated server, else 10%.
9 # innodb_buffer_pool_size = 128M
10 #
11 # Remove leading # to turn on a very important data
12 # integrity option: logging
13 # changes to the binary log between backups.
14 # log_bin
15 #
16 # Remove leading # to set options mainly useful for
17 # reporting servers.
18 # The server defaults are faster for transactions and
19 # fast SELECTs.
20 # Adjust sizes as needed, experiment to find the opti-
21 # mal values.
22 # join_buffer_size = 128M
23 # sort_buffer_size = 2M
24 # read_rnd_buffer_size = 2M
25 datadir=/var/lib/mysql
26 socket=/var/lib/mysql/mysql.sock
27 # Disabling symbolic-links is recommended to prevent
28 # assorted security risks
29 symbolic-links=0
30 log-error=/var/log/mysqld.log
31 pid-file=/var/run/mysqld/mysqld.pid
```

- mysqld_safe.cnf

```
1 [root@WP1 ~]# cat /etc/percona-server.conf.d/mysqld_
2 safe.cnf
3 #
4 # The Percona Server 5.7 configuration file.
5 #
6 # One can use all long options that the program sup-
7 # ports.
8 # Run program with --help to get a list of available
9 # options and with
10 # --print-defaults to see which it would actually un-
11 # derstand and use.
12 #
```

```

9   # For explanations see
10  # http://dev.mysql.com/doc/mysql/en/server-system-variables.html
11  [mysqld_safe]
12  pid-file = /var/run/mysqld/mysqld.pid
13  socket   = /var/run/mysqld/mysqld.sock
14  nice     = 0

```

Let's see these parameters in detail.

- `datadir`: The path to the MySQL server data directory.
- `socket`: On Unix platforms, this variable is the name of the socket file that is used for local client connections.
- `symbolic-links`: Enable or disable symbolic link support. On Unix, enabling symbolic links means that you can link a MyISAM index file or data file to another directory with the `INDEX DIRECTORY` or `DATA DIRECTORY` option of the `CREATE TABLE` statement.
- `log-error`: Write the error log and startup messages to this file.
- `pid-file`: The path name of the file in which the server should write its process ID. The server creates the file in the data directory unless an absolute path name is given to specify a different directory.
- `nice`: Use the `nice` program to set the server's scheduling priority to the given value.

Optional Percona Server Configuration

Let's see some variables that we can add in our config files. It's recommended as a basic configuration, but it ought to be further tuned based on your infrastructure.

```

1   #GENERAL
2   user=mysql
3   basedir=/usr/
4   port=3306
5   skip_name_resolve
6   ignore-db-dir=lost+found
7   #LOGGING
8   log_warnings=2
9   slow_query_log_file=/var/log/mysql/mysql-slow.log
10  long_query_time=2
11  slow_query_log=OFF
12  log_queries_not_using_indexes=OFF
13  log_slow_admin_statements=ON
14  #INNODB
15  innodb_buffer_pool_size=128M
16  innodb_flush_log_at_trx_commit=2
17  innodb_file_per_table=1
18  innodb_data_file_path = ibdata1:100M:autoextend
19  innodb_read_io_threads=4

```

```

20 innodb_write_io_threads=4
21 innodb_doublewrite=1
22 innodb_log_file_size=64M
23 innodb_log_buffer_size=16M
24 innodb_buffer_pool_instances=1
25 innodb_log_files_in_group=2
26 innodb_thread_concurrency=64
27 innodb_flush_method = O_DIRECT
28 innodb_rollback_on_timeout=ON
29 innodb_autoinc_lock_mode=2
30 innodb_stats_on_metadata=0
31 default_storage_engine=innodb
32 #REPLICATION
33 server_id=1
34 binlog_format=ROW
35 log_bin=binlog
36 log_slave_updates=1
37 gtid_mode=ON
38 enforce_gtid_consistency=1
39 relay_log=relay-bin
40 expire_logs_days=7
41 read_only=ON
42 sync_binlog=1
43 report_host=192.168.100.117
44 master_info_repository=TABLE
45 relay_log_info_repository=TABLE
46 relay_log_recovery=ON
47 #OTHER THINGS
48 tmp_table_size = 64M
49 max_heap_table_size = 64M
50 max_allowed_packet = 512M
51 sort_buffer_size = 256K
52 read_buffer_size = 256K
53 read_rnd_buffer_size = 512K
54 myisam_sort_buffer_size = 8M
55 memlock=0
56 sysdate_is_now=1
57 max_connections=500
58 thread_cache_size=512
59 query_cache_type = 0
60 query_cache_size = 0
61 table_open_cache=1024
62 lower_case_table_names=0

```

To see in detail these variables, you can follow [this link](#).

You can use the !include parameter, to split the configuration in different files, for example, the backup credentials.

Into /etc/percona-server.conf.d/mysqlld.cnf add the following line:

```
1 | !include /etc/percona-server.conf.d/secrets-backup.cnf
```

And then create the secrets-backup.cnf file:

```
1 | [root@WP1 ~]# cat /etc/percona-server.conf.d/secrets-backup.cnf
2 | # Security credentials for backup.
3 | [mysqldump]
4 | user=backupuser
5 | password=Dse0s0k0ZvXoHIItv
6 | [xtrabackup]
7 | user=backupuser
8 | password=Dse0s0k0ZvXoHIItv
```

Oracle MySQL Community Server

MySQL is the world's most popular open source database. With its proven performance, reliability, and ease-of-use, MySQL has become the leading database choice for web-based applications, used by high profile web properties including Facebook, Twitter and YouTube. Additionally, it is an extremely popular choice as embedded database, distributed by thousands of ISVs and OEMs.



Installation

To install the MySQL Community Server packages manually, you can follow [this link](#).

Another way to install it is by using [yum](#) or [apt](#) repositories.

In our example, let's see the yum repository installation of MySQL Community Server 5.7 on CentOS 7.

MySQL Server IP Address: 192.168.100.118

```
1 | [root@WP2 ~]# wget https://dev.mysql.com/get/mysql80-community-release-el7-1.noarch.rpm
2 | --2019-01-03 01:41:20-- https://dev.mysql.com/get/mysql80-community-release-el7-1.noarch.rpm
3 | Resolving dev.mysql.com (dev.mysql.com)... 137.254.60.11
4 | Connecting to dev.mysql.com (dev.mysql.com)|137.254.60.11|:443... connected.
5 | HTTP request sent, awaiting response... 302 Found
6 | Location: https://repo.mysql.com//mysql80-community-release-el7-1.noarch.rpm [following]
7 | --2019-01-03 01:41:22-- https://repo.mysql.com//mysql80-community-release-el7-1.noarch.rpm
```

```

8   Resolving repo.mysql.com (repo.mysql.com)... 23.208.182.226
9   Connecting to repo.mysql.com (repo.mysql.
    com)|23.208.182.226|:443... connected.
10  HTTP request sent, awaiting response... 200 OK
11  Length: 25820 (25K) [application/x-redhat-package-manager]
12  Saving to: 'mysql80-community-release-el7-1.noarch.rpm'
13
14  100%[=====
    =====
    =====>] 25,820      --.-K/s   in 0.04s
15
16  2019-01-03 01:41:22 (698 KB/s) - 'mysql80-community-re-
    lease-el7-1.noarch.rpm' saved [25820/25820]
17
18  [root@WP2 ~]# rpm -Uvh mysql80-community-release-el7-1.
    noarch.rpm
19  warning: mysql80-community-release-el7-1.noarch.rpm: Header
    V3 DSA/SHA1 Signature, key ID 5072e1f5: NOKEY
20  Preparing...                               #####
    ##### [100%]
21  Updating / installing...
22     1:mysql80-community-release-el7-1 #####
    ##### [100%]

```

Edit the /etc/yum.repos.d/mysql-community.repo file and set the enabled parameter to 1 for MySQL 5.7 and 0 for MySQL 8.0:

```

1   # Enable to use MySQL 5.7
2   [mysql57-community]
3   name=MySQL 5.7 Community Server
4   baseurl=http://repo.mysql.com/yum/mysql-5.7-community/
    el/7/$basearch/
5   enabled=1
6   gpgcheck=1
7   gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-mysql
8   [mysql80-community]
9   name=MySQL 8.0 Community Server
10  baseurl=http://repo.mysql.com/yum/mysql-8.0-community/
    el/7/$basearch/
11  enabled=0
12  gpgcheck=1
13  gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-mysql

```

And then, install mysql-community-server:

```

1   [root@WP2 ~]# yum install mysql-community-server
2   =====
    =====

```

```

3 | Package                               Arch      Version
   | Repository                             Size
4 | =====
   | =====
5 | Installing:
6 | mysql-community-server                 x86_64    5.7.24-1.el7
   | mysql57-community                       165 M
7 | Installing for dependencies:
8 | libaio                                  x86_64    0.3.109-13.el7
   | base                                     24 k
9 | mysql-community-client                 x86_64    5.7.24-1.el7
   | mysql57-community                       24 M
10 | mysql-community-common                 x86_64    5.7.24-1.el7
   | mysql57-community                       274 k
11 | mysql-community-libs                   x86_64    5.7.24-1.el7
   | mysql57-community                       2.2 M
12 | numactl-libs                           x86_64    2.0.9-7.el7
   | base                                     29 k
13 |
14 | Transaction Summary
15 | =====
   | =====
16 | Install 1 Package (+5 Dependent packages)

```

Now, you need to start the MySQL service.

```

1 | [root@WP2 ~]# service mysqld start
2 | Redirecting to /bin/systemctl start mysqld.service

```

A new password for the root user is created. To know it, you need to check the MySQL log and filter by temporary password. You should have something similar to this:

```

1 | [root@WP2 ~]# grep "temporary password" /var/log/mysqld.log
2 | 2019-01-03T01:52:30.896979Z 1 [Note] A temporary password is
   | generated for root@localhost: Dx6!MWC5cKYg

```

Then, you can run the `mysql_secure_installation` script, to configure a basic secure setup for your MySQL database.

```

1 | [root@WP2 ~]# mysql_secure_installation
2 | Securing the MySQL server deployment.
3 |
4 | Enter password for user root:
5 | The 'validate_password' plugin is installed on the server.
6 | The subsequent steps will run with the existing configuration
7 | of the plugin.
8 | Using existing password for root.
9 |
10 | Estimated strength of the password: 100

```



```
11 Change the password for root ? ((Press y|Y for Yes, any oth-
12 er key for No) : y
13 New password:
14
15 Re-enter new password:
16
17 Estimated strength of the password: 100
18 Do you wish to continue with the password provided?(Press
19 y|Y for Yes, any other key for No) : y
20 By default, a MySQL installation has an anonymous user,
21 allowing anyone to log into MySQL without having to have
22 a user account created for them. This is intended only for
23 testing, and to make the installation go a bit smoother.
24 You should remove them before moving into a production
25 environment.
26 Remove anonymous users? (Press y|Y for Yes, any other key
27 for No) : y
28 Success.
29 Normally, root should only be allowed to connect from
30 'localhost'. This ensures that someone cannot guess at
31 the root password from the network.
32
33 Disallow root login remotely? (Press y|Y for Yes, any other
34 key for No) : y
35 Success.
36
37 By default, MySQL comes with a database named 'test' that
38 anyone can access. This is also intended only for testing,
39 and should be removed before moving into a production
40 environment.
41 Remove test database and access to it? (Press y|Y for Yes,
42 any other key for No) : y
43 - Dropping test database...
44 Success.
45 - Removing privileges on test database...
46 Success.
47
48 Reloading the privilege tables will ensure that all changes
49 made so far will take effect immediately.
50
51 Reload privilege tables now? (Press y|Y for Yes, any other
52 key for No) : y
53 Success.
54 All done!
```

Now your database is running, but it's not ready yet. We need to configure it.

Default Configuration

The MySQL installation creates the `my.cnf` config file and the `/etc/my.cnf.d/` directory in `/etc/`.

The `/etc/my.cnf.d/` directory is empty by default, and the content of `my.cnf` is:

```
1 [root@WP2 ~]# cat /etc/my.cnf
2 # For advice on how to change settings please see
3 # http://dev.mysql.com/doc/refman/5.7/en/server-configura-
4 # tion-defaults.html
5 [mysqld]
6 #
7 # Remove leading # and set to the amount of RAM for the most
8 # important data
9 # cache in MySQL. Start at 70% of total RAM for dedicated
10 # server, else 10%.
11 # innodb_buffer_pool_size = 128M
12 #
13 # Remove leading # to turn on a very important data integri-
14 # ty option: logging
15 # changes to the binary log between backups.
16 # log_bin
17 #
18 # Remove leading # to set options mainly useful for report-
19 # ing servers.
20 # The server defaults are faster for transactions and fast
21 # SELECTs.
22 # Adjust sizes as needed, experiment to find the optimal val-
23 # ues.
24 # join_buffer_size = 128M
25 # sort_buffer_size = 2M
26 # read_rnd_buffer_size = 2M
27 datadir=/var/lib/mysql
28 socket=/var/lib/mysql/mysql.sock
29 # Disabling symbolic-links is recommended to prevent assort-
30 # ed security risks
31 symbolic-links=0
32 log-error=/var/log/mysqld.log
33 pid-file=/var/run/mysqld/mysqld.pid
```

Let's see these parameters in detail.

- `datadir`: The path to the MySQL server data directory.
- `socket`: On Unix platforms, this variable is the name of the socket file that is used for local client connections.
- `symbolic-links`: Enable or disable symbolic link support. On Unix, enabling

symbolic links means that you can link a MyISAM index file or data file to another directory with the INDEX DIRECTORY or DATA DIRECTORY option of the CREATE TABLE statement.

- log-error: Write the error log and startup messages to this file.
- pid-file: The path name of the file in which the server should write its process ID. The server creates the file in the data directory unless an absolute path name is given to specify a different directory.

Optional MySQL Community Server Configuration

Let's see some variables that we can add in our config files. It's recommended as a basic configuration, but ought to be tuned based on your infrastructure.

```
1 #GENERAL
2 user=mysql
3 basedir=/usr/
4 port=3306
5 skip_name_resolve
6 ignore-db-dir=lost+found
7 #LOGGING
8 log_warnings=2
9 slow_query_log_file=/var/log/mysql/mysql-slow.log
10 long_query_time=2
11 slow_query_log=OFF
12 log_queries_not_using_indexes=OFF
13 log_slow_admin_statements=ON
14 #INNODB
15 innodb_buffer_pool_size=128M
16 innodb_flush_log_at_trx_commit=2
17 innodb_file_per_table=1
18 innodb_data_file_path = ibdata1:100M:autoextend
19 innodb_read_io_threads=4
20 innodb_write_io_threads=4
21 innodb_doublewrite=1
22 innodb_log_file_size=64M
23 innodb_log_buffer_size=16M
24 innodb_buffer_pool_instances=1
25 innodb_log_files_in_group=2
26 innodb_thread_concurrency=64
27 innodb_flush_method = O_DIRECT
28 innodb_rollback_on_timeout=ON
29 innodb_autoinc_lock_mode=2
30 innodb_stats_on_metadata=0
31 default_storage_engine=innodb
32 #REPLICATION
33 server_id=1
34 binlog_format=ROW
35 log_bin=binlog
36 log_slave_updates=1
```

```

37 gtid_mode=ON
38 enforce_gtid_consistency=1
39 relay_log=relay-bin
40 expire_logs_days=7
41 read_only=ON
42 sync_binlog=1
43 report_host=192.168.100.118
44 master_info_repository=TABLE
45 relay_log_info_repository=TABLE
46 relay_log_recovery=ON
47 #OTHER THINGS
48 tmp_table_size = 64M
49 max_heap_table_size = 64M
50 max_allowed_packet = 512M
51 sort_buffer_size = 256K
52 read_buffer_size = 256K
53 read_rnd_buffer_size = 512K
54 myisam_sort_buffer_size = 8M
55 memlock=0
56 sysdate_is_now=1
57 max_connections=500
58 thread_cache_size=512
59 query_cache_type = 0
60 query_cache_size = 0
61 table_open_cache=1024
62 lower_case_table_names=0

```

To see in detail these variables, you can follow [this link](#).

You can use the `!include` parameter, to split the configuration in different files, for example, the backup credentials.

In `/etc/my.cnf`, add the following line:

```

1 | !include /etc/my.cnf.d/secrets-backup.cnf

```

And then create the `secrets-backup.cnf` file:

```

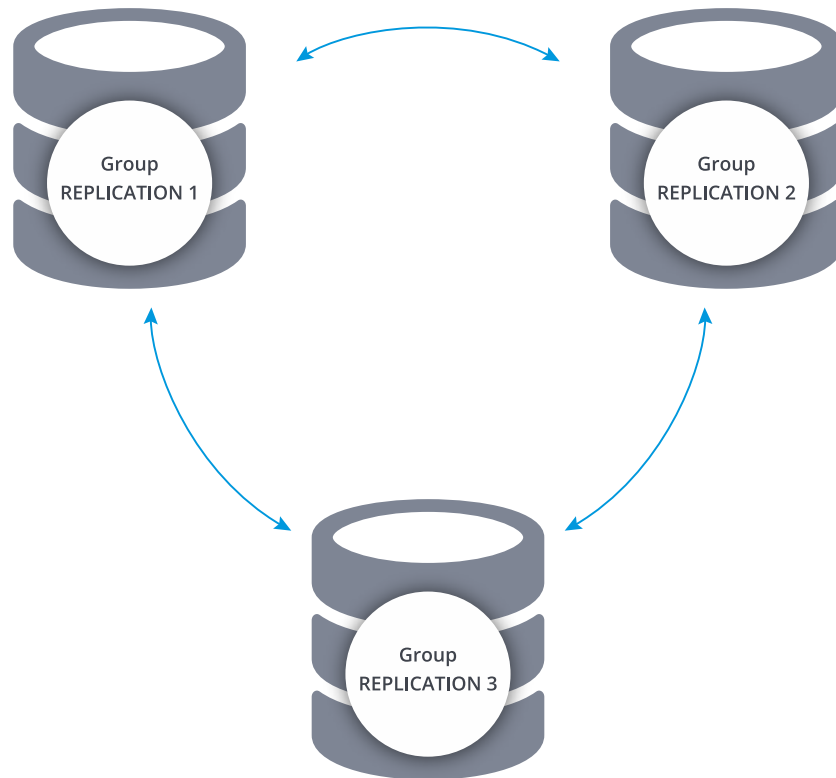
1 | [root@WP2 ~]# cat /etc/my.cnf.d/secrets-backup.cnf
2 | # Security credentials for backup.
3 | [mysqldump]
4 | user=backupuser
5 | password=Dse0s0k0ZvXoHIItv
6 | [xtrabackup]
7 | user=backupuser
8 | password=Dse0s0k0ZvXoHIItv

```

Group Replication

MySQL Group Replication is a MySQL Server plugin that enables you to create elastic, highly-available, fault-tolerant replication topologies. It can operate in a single-primary mode with automatic primary election, where only one server accepts updates at a time. Alternatively, for more advanced users, groups can be deployed in multi-primary mode, where all servers can accept updates, even if they are issued concurrently.

Let's see the basic configuration:



Node 1 IP Address: 192.168.100.186

Node 2 IP Address: 192.168.100.187

Node 3 IP Address: 192.168.100.188

In each node, add the following lines in the my.cnf file:

```
1 # GROUP REPLICATION
2 server_id=1
3 gtid_mode=ON
4 enforce_gtid_consistency=ON
5 master_info_repository=TABLE
6 relay_log_info_repository=TABLE
7 binlog_checksum=NONE
8 log_slave_updates=ON
9 log-bin=binlog
10 binlog_format=ROW
11 plugin-load=group_replication.so
12 transaction_write_set_extraction=XXHASH64
```

```

13 | loose-group_replication_group_name="4bd8d654-a40d-4792-8b88-
    | 16eff8b85f44"
14 | loose-group_replication_start_on_boot=off
15 | loose-group_replication_local_ad-
    | dress="192.168.100.186:33066"
16 | loose-group_replication_group_sees="192.168.100.186:33066,19
    | 2.168.100.187:33066,192.168.100.188:33066"
17 | loose-group_replication_bootstrap_group=off
18 | loose-group_replication_single_primary_mode=FALSE
19 | loose-group_replication_enforce_update_everywhere_
    | checks=TRUE
20 | loose-group_replication_ip_whitelist=192.168.100.0/24
21 | loose-group_replication_recovery_retry_count=3
22 | loose-group_replication_recovery_reconnect_interval=120

```

You need to change the parameters `server_id` and `loose-group_replication_local_address` according to each node, and restart the MySQL service to take the changes.

```

1 | [root@WP2 ~]# service mysqld restart
2 | Redirecting to /bin/systemctl restart mysqld.service

```

You can verify if the `group_replication` plugin is installed by running the following command in the MySQL server:

```

1 | mysql> SHOW PLUGINS;
2 | ***** 45. row
    | *****
3 |      Name: group_replication
4 |      Status: ACTIVE
5 |      Type: GROUP REPLICATION
6 |      Library: group_replication.so
7 |      License: GPL

```

If you can't see the plugin, you can install it by using the following command:

```

1 | mysql> INSTALL PLUGIN group_replication SONAME 'group_repli-
    | cation.so';

```

Then, you need to create the replication user and assign it to the Group Replication in each node:

```

1 | mysql> CREATE USER rep_user@'%' IDENTIFIED BY 'rep_pass';
2 | Query OK, 0 rows affected (0.00 sec)

```

```

1 | mysql> GRANT REPLICATION SLAVE ON *.* TO rep_user@'%' ;
2 | Query OK, 0 rows affected (0.00 sec)

```

```
1 | mysql> CHANGE MASTER TO MASTER_USER='rep_user', MASTER_PASS-
  | WORD='rep_pass' FOR CHANNEL 'group_replication_recovery';
2 | Query OK, 0 rows affected, 2 warnings (0.05 sec)
```

Now, in the first node, you need to initialize the cluster:

```
1 | mysql> SET GLOBAL group_replication_bootstrap_group=ON;
2 | Query OK, 0 rows affected (0.00 sec)
```

```
1 | mysql> START GROUP_REPLICATION;
2 | Query OK, 0 rows affected (2.11 sec)
```

```
1 | mysql> SET GLOBAL group_replication_bootstrap_group=OFF;
2 | Query OK, 0 rows affected (0.00 sec)
```

And check if the cluster is up:

```
1 | mysql> SELECT * FROM performance_schema.replication_group_
  | members\G
2 | ***** 1. row
  | *****
3 | CHANNEL_NAME: group_replication_applier
4 |     MEMBER_ID: a3de1ba4-35eb-11e9-bbde-067c0c0a7c38
5 |     MEMBER_HOST: Host28
6 |     MEMBER_PORT: 3306
7 |     MEMBER_STATE: ONLINE
8 | 1 row in set (0.01 sec)
```

Then, you must start the Group Replication in the rest of the nodes:

```
1 | mysql> START GROUP_REPLICATION;
2 | Query OK, 0 rows affected (5.85 sec)
```

And check the cluster status again:

```
1 | mysql> SELECT * FROM performance_schema.replication_group_
  | members\G
2 | ***** 1. row
  | *****
3 | CHANNEL_NAME: group_replication_applier
4 |     MEMBER_ID: a3de1ba4-35eb-11e9-bbde-067c0c0a7c38
5 |     MEMBER_HOST: Host28
6 |     MEMBER_PORT: 3306
7 |     MEMBER_STATE: ONLINE
8 | ***** 2. row
  | *****
```

```

9 | CHANNEL_NAME: group_replication_applier
10 |     MEMBER_ID: a5ae763f-35eb-11e9-bcf8-2ab9ec7193a3
11 |     MEMBER_HOST: Host29
12 |     MEMBER_PORT: 3306
13 |     MEMBER_STATE: ONLINE
14 | ***** 3. row
15 | *****
16 | CHANNEL_NAME: group_replication_applier
17 |     MEMBER_ID: a6bae589-35eb-11e9-bc18-aeb046c0331d
18 |     MEMBER_HOST: Host30
19 |     MEMBER_PORT: 3306
20 |     MEMBER_STATE: ONLINE
20 | 3 rows in set (0.00 sec)

```

Now, you have your Group Replication up and running.

MariaDB

MariaDB Server is one of the most popular database servers. It's made by the original developers of MySQL and guaranteed to stay open source.

MariaDB turns data into structured information in a wide array of applications, ranging from banking to websites. It is an enhanced, drop-in replacement for MySQL. MariaDB is used because it is fast, scalable and robust, with a rich ecosystem of storage engines, plugins and many other tools which make it very versatile for a wide variety of use cases.



MariaDB Cluster is a synchronous multi-master cluster based on Galera replication. It is available on Linux only, and only supports the XtraDB/InnoDB storage engines

These installation steps are for both MariaDB Server and MariaDB Cluster. MariaDB uses the same binaries for both databases. The difference between them is the configuration parameters. We're going to see how to configure both in some minutes.

Installation

In [this link](#) you can download the latest packages to install MariaDB, or if you prefer, you can follow the [repository installation link](#).

In our example, let's see the installation of MariaDB on CentOS 7 from the repository.

IP Address: 192.168.100.143

To add the MariaDB repository, you can run:

```

1 | cat > /etc/yum.repos.d/MariaDB.repo <<- EOF
2 | # MariaDB 10.3 CentOS repository

```



```

3 [mariadb]
4 name = MariaDB
5 baseurl = http://yum.mariadb.org/10.3/centos7-amd64
6 gpgkey=https://yum.mariadb.org/RPM-GPG-KEY-MariaDB
7 gpgcheck=1
8 EOF

```

And then, install MariaDB-server and MariaDB-client packages:

```

1 [root@WP3 ~]# yum install MariaDB-server MariaDB-client
2 =====
3 Package Arch Version
4 Repository Size
5 =====
5 Installing:
6 MariaDB-client x86_64 10.3.11-1.el7.
7 centos mariadb 53 M
8 MariaDB-server x86_64 10.3.11-1.el7.
9 centos mariadb 123 M
10 Installing for dependencies:
11 MariaDB-common x86_64 10.3.11-1.el7.
12 centos mariadb 157 k
13 MariaDB-compat x86_64 10.3.11-1.el7.
14 centos mariadb 2.8 M
15 boost-program-options x86_64 1.53.0-27.el7
16 base 156 k
17 galera x86_64 25.3.24-1.
18 rhel7.el7.centos mariadb 8.1 M
19 libaio x86_64 0.3.109-13.el7
20 base 24 k
21 lsof x86_64 4.87-6.el7
22 base 331 k
23 make x86_64 1:3.82-23.el7
24 base 420 k
25 openssl x86_64 1:1.0.2k-16.
26 el7 base 493 k
27 perl-Compress-Raw-Bzip2 x86_64 2.061-3.el7
28 base 32 k
29 perl-Compress-Raw-Zlib x86_64 1:2.061-4.el7
30 base 57 k
31 perl-DBI x86_64 1.627-4.el7
32 base 802 k
33 perl-Data-Dumper x86_64 2.145-3.el7
34 base 47 k
35 perl-IO-Compress noarch 2.061-2.el7
36 base 260 k
37 perl-Net-Daemon noarch 0.48-5.el7

```

```

23 | base                    51 k
    | perl-PlRPC              noarch    0.2020-14.e17
    | base                    36 k
24 | Updating for dependencies:
25 | openssl-libs           x86_64    1:1.0.2k-16.
    | e17                    base        1.2 M
26 |
27 | Transaction Summary
28 | =====
    | =====
29 | Install  2 Packages (+15 Dependent packages)
30 | Upgrade   ( 1 Dependent package)

```

After this, you need to start the MySQL service.

```

1 | [root@WP3 ~]# service mysql start
2 | Starting mysql (via systemctl): [ OK ]

```

By default, MariaDB is installed without root password, so you only need to run the mysql command to access the database with root privileges.

```

1 | [root@WP3 ~]# mysql
2 | Welcome to the MariaDB monitor.  Commands end with ; or \g.
3 | Your MariaDB connection id is 8
4 | Server version: 10.3.11-MariaDB MariaDB Server
5 |
6 | Copyright (c) 2000, 2018, Oracle, MariaDB Corporation Ab and
   | others.
7 |
8 | Type 'help;' or '\h' for help. Type '\c' to clear the cur-
   | rent input statement.
9 |
10 | MariaDB [(none)]>

```

It's recommended to use the mysql_secure_installation to improve your database security.

```

1 | [root@WP3 ~]# mysql_secure_installation
2 | NOTE: RUNNING ALL PARTS OF THIS SCRIPT IS RECOMMENDED FOR
   | ALL MariaDB
3 |     SERVERS IN PRODUCTION USE!  PLEASE READ EACH STEP
   | CAREFULLY!
4 | In order to log into MariaDB to secure it, we'll need the
   | current
5 | password for the root user.  If you've just installed Mari-
   | aDB, and
6 | you haven't set the root password yet, the password will be
   | blank,
7 | so you should just press enter here.

```

```
8
9 Enter current password for root (enter for none):
10 OK, successfully used password, moving on...
11
12 Setting the root password ensures that nobody can log into
13 the MariaDB
14 root user without the proper authorisation.
15
16 Set root password? [Y/n] y
17 New password:
18 Re-enter new password:
19 Password updated successfully!
20 Reloading privilege tables..
21 ... Success!
22
23 By default, a MariaDB installation has an anonymous user,
24 allowing anyone
25 to log into MariaDB without having to have a user account
26 created for
27 them. This is intended only for testing, and to make the
28 installation
29 go a bit smoother. You should remove them before moving
30 into a
31 production environment.
32
33 Remove anonymous users? [Y/n]
34 ... Success!
35
36 Normally, root should only be allowed to connect from 'lo-
37 calhost'. This
38 ensures that someone cannot guess at the root password from
39 the network.
40
41 Disallow root login remotely? [Y/n]
42 ... Success!
43
44 By default, MariaDB comes with a database named 'test' that
45 anyone can
46 access. This is also intended only for testing, and should
47 be removed
48 before moving into a production environment.
49
50 Remove test database and access to it? [Y/n]
51 - Dropping test database...
52 ... Success!
53 - Removing privileges on test database...
54 ... Success!
55
56 Reloading the privilege tables will ensure that all changes
```

```

made so far
48 will take effect immediately.
49
50 Reload privilege tables now? [Y/n]
51 ... Success!
52
53 Cleaning up...
54
55 All done! If you've completed all of the above steps, your
MariaDB
56 installation should now be secure.
57
58 Thanks for using MariaDB!

```

Now your database is running, but it's not ready yet. We need more work here.

Default Configuration

For MariaDB, the my.cnf config file includes the /etc/my.cnf.d directory:

```

1 [root@WP3 ~]# cat /etc/my.cnf
2 #
3 # This group is read both both by the client and the server
4 # use it for options that affect everything
5 #
6 [client-server]
7 #
8 # include all files from the config directory
9 #
10 !includedir /etc/my.cnf.d

```

And there, you have the following files by default:

```

1 [root@WP3 ~]# ls /etc/my.cnf.d/
2 enable_encryption.preset mysql-clients.cnf server.cnf

```

For encryption configuration, you have the enable_encryption file:

```

1 [root@WP3 ~]# cat /etc/my.cnf.d/enable_encryption.preset
2 #
3 # !include this file into your my.cnf (or any of *.cnf files
4 # in /etc/my.cnf.d)
5 # and it will enable data at rest encryption. This is a sim-
6 # ple way to
7 # ensure that everything that can be encrypted will be and
8 # your
9 # data will not leak unencrypted.
10 #

```

```

8   # DO NOT EDIT THIS FILE! On MariaDB upgrades it might be re-
9   # placed with a
10  # newer version and your edits will be lost. Instead, add
11  # your edits
12  # to the .cnf file after the !include directive.
13  #
14  # NOTE that you also need to install an encryption plugin
15  # for the encryption
16  # to work. See https://mariadb.com/kb/en/mariadb/data-at-rest-encryption/#encryption-key-management
17  #
18  [mariadb]
19  aria-encrypt-tables
20  encrypt-binlog
21  encrypt-tmp-disk-tables
22  encrypt-tmp-files
23  loose-innodb-encrypt-log
24  loose-innodb-encrypt-tables

```

For clients configuration, you can use the `mysql-clients.cnf` file where you can use a different configuration for each client:

```

1   [root@WP3 ~]# cat /etc/my.cnf.d/mysql-clients.cnf
2   #
3   # These groups are read by MariaDB command-line tools
4   # Use it for options that affect only one utility
5   #
6   [mysql]
7   [mysql_upgrade]
8   [mysqladmin]
9   [mysqlbinlog]
10  [mysqlcheck]
11  [mysqldump]
12  [mysqlimport]
13  [mysqlshow]
14  [mysqlslap]

```

And finally, you have the `server.cnf` configuration file:

```

1   [root@WP3 ~]# cat /etc/my.cnf.d/server.cnf
2   #
3   # These groups are read by MariaDB server.
4   # Use it for options that only the server (but not clients)
5   # should see
6   #
7   # See the examples of server my.cnf files in /usr/share/
8   # mysql/
9   #
10  # this is read by the standalone daemon and embedded servers
11  [server]

```

```

10 # this is only for the mysqld standalone daemon
11 [mysqld]
12 #
13 # * Galera-related settings
14 #
15 [galera]
16 # Mandatory settings
17 #wsrep_on=ON
18 #wsrep_provider=
19 #wsrep_cluster_address=
20 #binlog_format=row
21 #default_storage_engine=InnoDB
22 #innodb_autoinc_lock_mode=2
23 #
24 # Allow server to accept connections on all interfaces.
25 #
26 #bind-address=0.0.0.0
27 #
28 # Optional setting
29 #wsrep_slave_threads=1
30 #innodb_flush_log_at_trx_commit=0
31 # this is only for embedded server
32 [embedded]
33 # This group is only read by MariaDB servers, not by MySQL.
34 # If you use the same .cnf file for MySQL and MariaDB,
35 # you can put MariaDB-only options here
36 [mariadb]
37 # This group is only read by MariaDB-10.3 servers.
38 # If you use the same .cnf file for MariaDB of different ver-
39 # sions,
40 # use this group for options that older servers don't under-
41 # stand
42 [mariadb-10.3]

```

Let's see these parameters in detail.

- `aria-encrypt-tables`: Enables automatic encryption of all user-created Aria tables that have the `ROW_FORMAT` table option set to `PAGE`.
- `encrypt-binlog`: Encrypt binary logs (including relay logs).
- `encrypt-tmp-disk-tables`: Enables automatic encryption of all internal on-disk temporary tables that are created during query execution if `aria_used_for_temp_tables=ON` is set.
- `encrypt-tmp-files`: Enables automatic encryption of temporary files, such as those created for filesort operations, binary log file caches, etc.
- `loose-innodb-encrypt-log`: Enables encryption of the InnoDB redo log. This also enables encryption of some temporary files created internally by InnoDB, such as those used for merge sorts and row logs.
- `loose-innodb-encrypt-tables`: Enables automatic encryption of all InnoDB tablespaces.

You can also add your own configuration file in `/etc/my.cnf.d/` or just add a different path in the `my.cnf` file using the parameter `!include`.

Optional MariaDB Server Configuration

Let's see some variables that we can add in our config files. It's recommended as a basic configuration, but ought to be tuned based on your infrastructure.

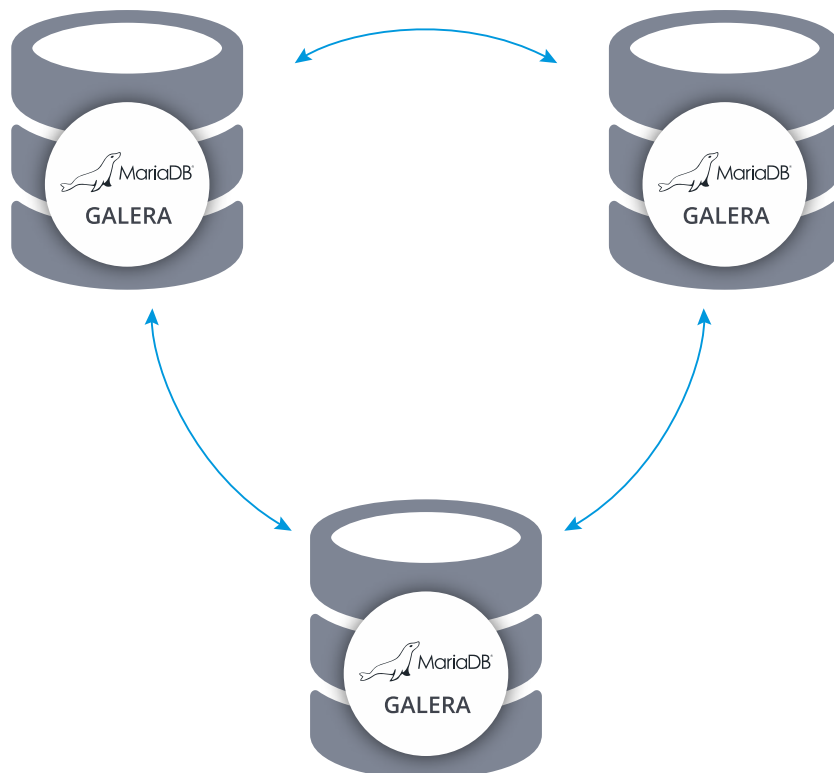
```
1 #GENERAL
2 user=mysql
3 basedir=/usr/
4 datadir=/var/lib/mysql
5 socket=/var/lib/mysql/mysql.sock
6 pid_file=/var/lib/mysql/mysql.pid
7 port=3306
8 ignore-db-dir=lost+found
9 #LOGGING
10 log_error=/var/log/mysql/mysqld.log
11 log_warnings=2
12 slow_query_log_file=/var/log/mysql/mysql-slow.log
13 long_query_time=2
14 slow_query_log=OFF
15 log_queries_not_using_indexes=OFF
16 #INNODB
17 innodb_buffer_pool_size=128M
18 innodb_flush_log_at_trx_commit=2
19 innodb_file_per_table=1
20 innodb_data_file_path = ibdata1:100M:autoextend
21 innodb_read_io_threads=4
22 innodb_write_io_threads=4
23 innodb_doublewrite=1
24 innodb_log_file_size=64M
25 innodb_log_buffer_size=16M
26 innodb_buffer_pool_instances=1
27 innodb_log_files_in_group=2
28 innodb_thread_concurrency=64
29 innodb_flush_method = O_DIRECT
30 innodb_rollback_on_timeout=ON
31 innodb_autoinc_lock_mode=2
32 innodb_stats_on_metadata=0
33 default_storage_engine=innodb
34 #REPLICATION
35 server_id=1
36 binlog_format=ROW
37 log_bin=binlog
38 log_slave_updates=1
39 relay_log=relay-bin
40 expire_logs_days=7
41 read_only=ON
42 report_host=192.168.100.143
```

```
43 # OTHER THINGS
44 key_buffer_size = 24M
45 tmp_table_size = 64M
46 max_heap_table_size = 64M
47 max_allowed_packet = 512M
48 skip_name_resolve
49 memlock=0
50 sysdate_is_now=1
51 max_connections=500
52 thread_cache_size=512
53 query_cache_type = 0
54 query_cache_size = 0
55 table_open_cache=1024
56 lower_case_table_names=0
```

To see these variables in detail, you can follow [this link](#).

MariaDB Cluster Configuration

As we mentioned previously, MariaDB uses the same binaries for both MariaDB Server and MariaDB Galera Cluster (they used to be separate binaries). There are some variables that we must configure to have MariaDB Cluster enabled.



Node 1 IP Address: 192.168.100.131

Node 2 IP Address: 192.168.100.132

Node 3 IP Address: 192.168.100.133

- `wsrep_provider=/usr/lib64/galera/libgalera_smm.so`: Path to the Galera library.
- `wsrep_cluster_address=gcomm://192.168.100.131,192.168.100.132,192.168.100.133`: `gcomm` is the option to use for a working implementation.
- `binlog_format=ROW`: There are three formats for binary logging: statement-based, row-based and mixed.
- `default_storage_engine=InnoDB`
- `innodb_autoinc_lock_mode=2`: Locking mode used for generating auto-increment values. 0 is the traditional lock mode, 1 - the consecutive, and 2 - the interleaved.
- `innodb_doublewrite=1`: This is the default value. To improve fault tolerance InnoDB first stores data to a doublewrite buffer before writing it to data file.
- `query_cache_size=0`: Only mandatory for MariaDB versions prior to MariaDB Galera Cluster 5.5.40, MariaDB Galera Cluster 10.0.14, and MariaDB 10.1.2.
- `wsrep_on=ON`: Enable wsrep replication (starting 10.1.1)

Also, there are some optional variables to add to configure your MariaDB Cluster.

```

1  wsrep_node_address=192.168.100.131
2  wsrep_provider_options="base_port=4567; gcache.size=1024M;
   gmcast.segment=0 "
3  wsrep_cluster_name="MariaDB1"
4  wsrep_cluster_address=g-
   comm://192.168.100.131,192.168.100.132,192.168.100.133
5  wsrep_node_name=192.168.100.131
6  wsrep_slave_threads=4
7  wsrep_certify_nonPK=1
8  wsrep_max_ws_rows=131072
9  wsrep_max_ws_size=1073741824
10 wsrep_debug=0
11 wsrep_convert_LOCK_to_trx=0
12 wsrep_retry_autocommit=1
13 wsrep_auto_increment_control=1
14 wsrep_replicate_myisam=1
15 wsrep_drupal_282555_workaround=0
16 wsrep_causal_reads=0
17 wsrep_sst_method=mariabackup
18 wsrep_log_conflicts=1
19 wsrep_gtid_domain_id=9999
20 wsrep_gtid_mode=1

```

Keep in mind that some of these values depend on your infrastructure.

Percona XtraDB Cluster

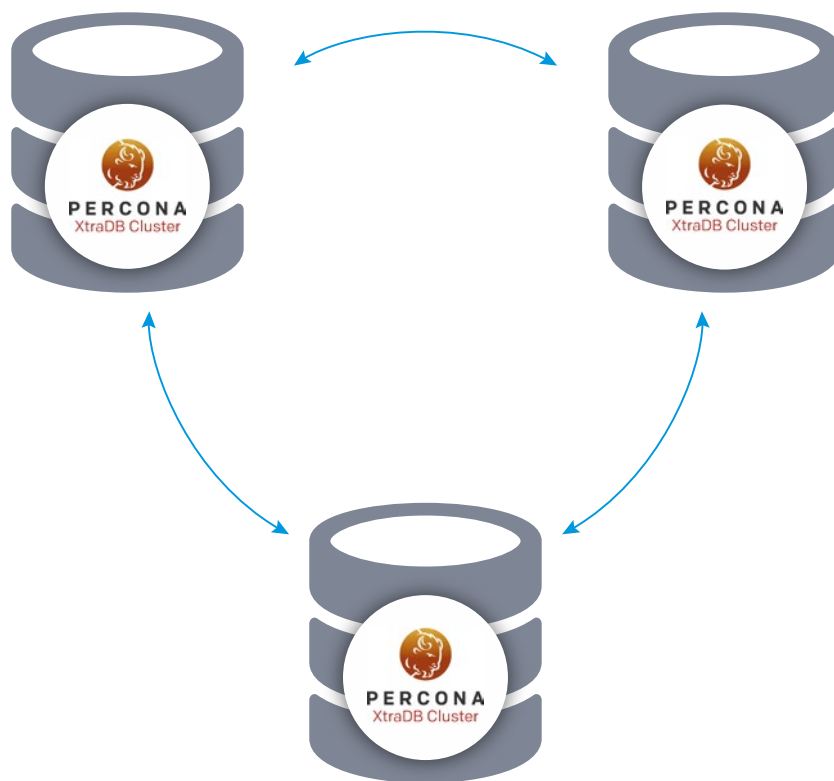
Percona XtraDB Cluster is an open source, cost-effective, and robust MySQL clustering solution for businesses. Organizations use Percona XtraDB Cluster to power highly available applications in the most demanding public, private and hybrid cloud environments. It is an excellent solution for your MySQL cluster and database needs.

Installation

In [this link](#) you have the latest packages to install Percona XtraDB Cluster.

If you prefer, you can use the [yum repository](#) or [apt repository](#) installation.

In our example, let's see the yum repository installation of Percona XtraDB Cluster on CentOS 7.



Node 1 IP Address: 192.168.100.154

Node 2 IP Address: 192.168.100.155

Node 3 IP Address: 192.168.100.156

```
1 [root@WP4 ~]# yum install http://www.percona.com/downloads/
2 percona-release/redhat/0.1-6/percona-release-0.1-6.noarch.
3 rpm
4 =====
5 =====
6 Package                               Arch    Ver-
7 sion                                 Repo-   Size
8 =====
```

```

=====
5 Installing:
6   percona-release                                noarch      0.1-6
   /percona-release-0.1-6.noarch                  16 k
7
8 Transaction Summary
9 =====
10 Install 1 Package
11
12 [root@WP4 ~]# yum install Percona-XtraDB-Cluster-57
13 =====
14 Package                                          Arch      Ver-
   sion              Repository                               Size
15 =====
16 Installing:
17   Percona-XtraDB-Cluster-57                      x86_64
   5.7.23-31.31.2.el7  percona-release-x86_64                27 k
18 Installing for dependencies:
19   Percona-XtraDB-Cluster-client-57                x86_64
   5.7.23-31.31.2.el7  percona-release-x86_64                7.2 M
20   Percona-XtraDB-Cluster-server-57                x86_64
   5.7.23-31.31.2.el7  percona-release-x86_64                51 M
21   Percona-XtraDB-Cluster-shared-57                x86_64
   5.7.23-31.31.2.el7  percona-release-x86_64                737 k
22   Percona-XtraDB-Cluster-shared-compat-57         x86_64
   5.7.23-31.31.2.el7  percona-release-x86_64                1.1 M
23   libaio                                           x86_64
   0.3.109-13.el7      base                                    24 k
24   libev                                           x86_64
   7.el7                extras                                  44 k
25   lsof                                           x86_64
   6.el7                base                                    331 k
26   numactl-libs                                   x86_64
   7.el7                base                                    29 k
27   percona-xtrabackup-24                          x86_64
   2.4.12-1.el7         percona-release-x86_64                7.5 M
28   perl-Compress-Raw-Bzip2                        x86_64
   3.el7                base                                    32 k
29   perl-Compress-Raw-Zlib                         x86_64
   1:2.061-4.el7        base                                    57 k
30   perl-DBD-MySQL                                 x86_64
   6.el7                base                                    140 k
31   perl-DBI                                       x86_64
   4.el7                base                                    802 k
32   perl-Data-Dumper                               x86_64
   3.el7                base                                    47 k

```

```

33 | perl-Digest                noarch  1.17-
    | 245.e17                    base    23 k
34 | perl-Digest-MD5           x86_64  2.52-
    | 3.e17                      base    30 k
35 | perl-IO-Compress          noarch  2.061-
    | 2.e17                      base    260 k
36 | perl-Net-Daemon           noarch  0.48-
    | 5.e17                      base    51 k
37 | perl-PIRPC                noarch
    | 0.2020-14.e17              base    36 k
38 | qpress                    x86_64  11-1.
    | e17                        percona-release-x86_64 31 k
39 | socat                     x86_64
    | 1.7.3.2-2.e17              base    290 k
40 |
41 | Transaction Summary
42 | =====
    | =====
43 | Install 1 Package (+21 Dependent packages)

```

You need to repeat these installation steps for each node. It's recommended to have at least three nodes in a cluster to improve the HA environment.

Before we start to use our cluster, we need to configure it. For this, please edit the `/etc/percona-xtradb-cluster.conf.d/wsrep.cnf` configuration file:

```

1 | wsrep_cluster_address=g-
  | comm://192.168.100.154,192.168.100.155,192.168.100.156
2 | #Replace the IP Address for the IP of each node.
3 | wsrep_node_address=192.168.100.154
4 | #Replace the IP Address for the IP of the current node
5 | wsrep_cluster_name=cluster1
6 | #Replace for your Cluster Name
7 | wsrep_node_name=node1
8 | #Replace for your Node Name
9 | wsrep_sst_auth="sstuser:Password!"
10 | #Replace for the SST credential that you want to use in your
   | new Cluster
11 | wsrep_on=ON

```

Now, you need to start the MySQL service. If it's the first node, you need to bootstrap the cluster:

```
1 | [root@WP4 ~]# systemctl start mysql@bootstrap.service
```

A new password for the root user is created. To know it, you need to check the MySQL log and filter by temporary password. You should have something similar to this:

```
1 | [root@WP4 ~]# grep "temporary password" /var/log/mysqld.log
```

```
2 | 2019-01-03T22:58:39.518754Z 1 [Note] A temporary password is
generated for root@localhost: eFrhjwjh+1tH
```

Connect to the database using the temporary password and change the current root password (it's required by the server before creating a new user)

```
1 | [root@WP4 ~]# mysql -p
2 | mysql> SET PASSWORD='*****';
3 | Query OK, 0 rows affected (0.00 sec)
```

Create a SST user for localhost (the user that we have in the `wsrep_sst_auth` variable):

```
1 | mysql> GRANT RELOAD, LOCK TABLES, PROCESS, REPLICATION CLI-
  | ENT ON *.* TO 'sstuser'@'192.168.100.%' IDENTIFIED BY 'Pass-
  | word!';
2 | Query OK, 0 rows affected, 1 warning (0.01 sec)
```

```
1 | mysql> FLUSH PRIVILEGES;
2 | Query OK, 0 rows affected (0.01 sec)
```

For the rest of the nodes, you only need to start the MySQL service as usual:

```
1 | [root@WP4 ~]# systemctl start mysql
```

This process will perform a SST and start the service in the node.

After starting all the nodes, your database cluster is running, but it's not ready yet. Let's see the default and optional configuration.

Default Configuration

As we could see previously on Percona Server for MySQL, by default, the Percona's `my.cnf` config file includes the `/etc/my.cnf.d/` and `/etc/percona-server.conf.d/` directories:

```
1 | [root@WP4 ~]# cat /etc/my.cnf
2 | #
3 | # The Percona XtraDB Cluster 5.7 configuration file.
4 | #
5 | #
6 | # * IMPORTANT: Additional settings that can override those
  | from this file!
7 | #   The files must end with '.cnf', otherwise they'll be ig-
  | nored.
8 | #   Please make any edits and changes to the appropriate
  | sectional files
9 | #   included below.
10 | #
11 | !includedir /etc/my.cnf.d/
12 | !includedir /etc/percona-xtradb-cluster.conf.d/
```

The `/etc/my.cnf.d/` directory is empty by default, and in `/etc/percona-server.conf.d/` we have the following content:

```
1 [root@WP4 ~]# ls /etc/percona-xtradb-cluster.conf.d/
2 mysql.cnf  mysqld_safe.cnf  wsrep.cnf
```

`mysqld_safe` is the recommended way to start a `mysqld` server on Unix. It adds some safety features such as restarting the server when an error occurs and logging runtime information to an error log.

`mysqld_safe` reads options from both `[mysqld]` and `[mysqld_safe]` sections in the configuration files.

The content of these configuration files are:

- `mysql.cnf`

```
1 [root@WP4 ~]# cat /etc/percona-xtradb-cluster.conf.d/
  mysql.cnf
2 # Template my.cnf for PXC
3 # Edit to your requirements.
4 [client]
5 socket=/var/lib/mysql/mysql.sock
6 [mysqld]
7 server-id=1
8 datadir=/var/lib/mysql
9 socket=/var/lib/mysql/mysql.sock
10 log-error=/var/log/mysqld.log
11 pid-file=/var/run/mysqld/mysqld.pid
12 log-bin
13 log_slave_updates
14 expire_logs_days=7
15 # Disabling symbolic-links is recommended to prevent
  assorted security risks
16 symbolic-links=0
```

- `mysqld_safe.cnf`

```
1 [root@WP4 ~]# cat /etc/percona-xtradb-cluster.conf.d/
  mysqld_safe.cnf
2 #
3 # The Percona Server 5.7 configuration file.
4 #
5 # One can use all long options that the program sup-
  ports.
6 # Run program with --help to get a list of available
  options and with
7 # --print-defaults to see which it would actually un-
  derstand and use.
8 #
9 # For explanations see
```

```

10 # http://dev.mysql.com/doc/mysql/en/server-sys-
    tem-variables.html
11 [mysqld_safe]
12 pid-file = /var/run/mysqld/mysqld.pid
13 socket   = /var/lib/mysql/mysql.sock
14 nice     = 0

```

- wsrep.cnf

```

1 [root@WP4 ~]# cat /etc/percona-xtradb-cluster.conf.d/
  wsrep.cnf
2 [mysqld]
3 # Path to Galera library
4 wsrep_provider=/usr/lib64/galera3/libgalera_smm.so
5 # Cluster connection URL contains IPs of nodes
6 #If no IP is found, this implies that a new cluster
  needs to be created,
7 #in order to do that you need to bootstrap this node
8 wsrep_cluster_address=gcomm://
9 # In order for Galera to work correctly binlog format
  should be ROW
10 binlog_format=ROW
11 # MyISAM storage engine has only experimental support
12 default_storage_engine=InnoDB
13 # Slave thread to use
14 wsrep_slave_threads= 8
15 wsrep_log_conflicts
16 # This changes how InnoDB autoincrement locks are man-
  aged and is a requirement for Galera
17 innodb_autoinc_lock_mode=2
18 # Node IP address
19 #wsrep_node_address=192.168.70.63
20 # Cluster name
21 wsrep_cluster_name=pxc-cluster
22 #If wsrep_node_name is not specified, then system
  hostname will be used
23 wsrep_node_name=pxc-cluster-node-1
24 #pxc_strict_mode allowed values: DISABLED,PERMIS-
  SIVE,ENFORCING,MASTER
25 pxc_strict_mode=ENFORCING
26 # SST method
27 wsrep_sst_method=xtrabackup-v2
28 #Authentication for SST method
29 #wsrep_sst_auth="sstuser:s3cretPass"

```

Let's see these parameters in detail.

- server-id: Specifies the server ID. The server_id system variable is set to 0 by default.
- datadir: The path to the MySQL server data directory.

- `socket`: On Unix platforms, this variable is the name of the socket file that is used for local client connections.
- `log-error`: Write the error log and startup messages to this file.
- `pid-file`: The path name of the file in which the server should write its process ID. The server creates the file in the data directory unless an absolute path name is given to specify a different directory.
- `log-bin`: Enables binary logging. With binary logging enabled, the server logs all statements that change data to the binary log, which is used for backup and replication. The binary log is a sequence of files with a base name and a numeric extension.
- `log_slave_updates`: Whether updates received by a slave server from a master server should be logged to the slave's own binary log.
- `expire_logs_days`: The number of days for automatic binary log file removal. The default is 0, which means "no automatic removal."
- `symbolic-links`: Enable or disable symbolic link support. On Unix, enabling symbolic links means that you can link a MyISAM index file or data file to another directory with the `INDEX DIRECTORY` or `DATA DIRECTORY` option of the `CREATE TABLE` statement.
- `nice`: Use the `nice` program to set the server's scheduling priority to the given value.
- `wsrep_provider`: Specifies the path to the Galera library. This is usually `/usr/lib64/libgalera_smm.so` on CentOS/RHEL and `/usr/lib/libgalera_smm.so` on Debian/Ubuntu.
- `wsrep_cluster_address`: Defines the back-end schema, IP addresses, ports, and options that the node uses when connecting to the cluster. This variable needs to specify at least one other node's address, which is alive and a member of the cluster.
- `binlog_format`: This variable sets the binary logging format, and can be any one of `STATEMENT`, `ROW`, or `MIXED`.
- `default_storage_engine`: Galera fully supports only the InnoDB storage engine. It will not work correctly with MyISAM or any other non-transactional storage engines.
- `wsrep_slave_threads`: Specifies the number of threads that can apply replication transactions in parallel. Galera supports true parallel replication that applies transactions in parallel only when it is safe to do so.
- `wsrep_log_conflicts`: Defines whether the node should log additional information about conflicts.
- `innodb_autoinc_lock_mode`: Galera supports only interleaved (2) lock mode for InnoDB.
- `wsrep_cluster_name`: Specify the logical name for your cluster. It must be the same for all nodes in your cluster.
- `wsrep_node_name`: Specify the logical name for each individual node. If this variable is not specified, the host name will be used.
- `pxc_strict_mode`: PXC Strict Mode is enabled by default and set to `ENFORCING`, which blocks the use of experimental and unsupported features in Percona XtraDB Cluster.
- `wsrep_sst_method`: By default, Percona XtraDB Cluster uses Percona XtraBackup for State Snapshot Transfer (SST). Setting `wsrep_sst_method=xtrabackup-v2` is

highly recommended. This method requires a user for SST to be set up on the initial node. Provide SST user credentials with the `wsrep_sst_auth` variable.

Optional Percona Server Configuration

Let's see some variables that we can add in our config files. It's recommended as a basic configuration, but it depends on your infrastructure.

```
1  #GENERAL
2  user=mysql
3  basedir=/usr/
4  port=3306
5  skip_name_resolve
6  ignore-db-dir=lost+found
7  #LOGGING
8  log_warnings=2
9  slow_query_log_file=/var/log/mysql/mysql-slow.log
10 long_query_time=2
11 slow_query_log=OFF
12 log_queries_not_using_indexes=OFF
13 log_slow_admin_statements=ON
14 log_throttle_queries_not_using_indexes=1
15 #INNODB
16 innodb_buffer_pool_size=128M
17 innodb_flush_log_at_trx_commit=2
18 innodb_file_per_table=1
19 innodb_data_file_path = ibdata1:100M:autoextend
20 innodb_read_io_threads=4
21 innodb_write_io_threads=4
22 innodb_doublewrite=1
23 innodb_log_file_size=64M
24 innodb_log_buffer_size=16M
25 innodb_buffer_pool_instances=1
26 innodb_log_files_in_group=2
27 innodb_thread_concurrency=64
28 innodb_flush_method = O_DIRECT
29 innodb_autoinc_lock_mode=2
30 innodb_stats_on_metadata=0
31 default_storage_engine=innodb
32 #REPLICATION
33 server_id=1
34 binlog_format=ROW
35 #OTHER THINGS
36 tmp_table_size = 64M
37 max_heap_table_size = 64M
38 max_allowed_packet = 512M
39 memlock=0
40 sysdate_is_now=1
41 max_connections=500
42 thread_cache_size=512
```

```

43 query_cache_type = 0
44 query_cache_size = 0
45 table_open_cache=1024
46 lower_case_table_names=0
47 # WSREP
48 wsrep_provider_options="base_port=4567; gcache.size=1024M;
   gmcast.segment=0 "
49 wsrep_certify_nonPK=1
50 wsrep_max_ws_rows=131072
51 wsrep_max_ws_size=1073741824
52 wsrep_debug=0
53 wsrep_convert_LOCK_to_trx=0
54 wsrep_retry_autocommit=1
55 wsrep_auto_increment_control=1
56 wsrep_replicate_myisam=0
57 wsrep_drupal_282555_workaround=0
58 wsrep_causal_reads=0

```

To see these variables in detail, you can follow [this link](#) or [this one](#).

You can use the !include parameter, to split the configuration in different files, for example, the backup credentials.

Into /etc/percona-server.conf.d/mysql.d.cnf add the following line:

```

1 | !include /etc/percona-server.conf.d/secrets-backup.cnf

```

And then create the secrets-backup.cnf file:

```

1 | [root@WP4 ~]# cat /etc/percona-server.conf.d/secrets-backup.
   cnf
2 | # Security credentials for backup.
3 | [mysqldump]
4 | user=backupuser
5 | password=Dse0s0k0ZvXoHItv
6 | [xtrabackup]
7 | user=backupuser
8 | password=Dse0s0k0ZvXoHItv

```

Keep in mind that some of these values depends on your infrastructure.

NDB Cluster

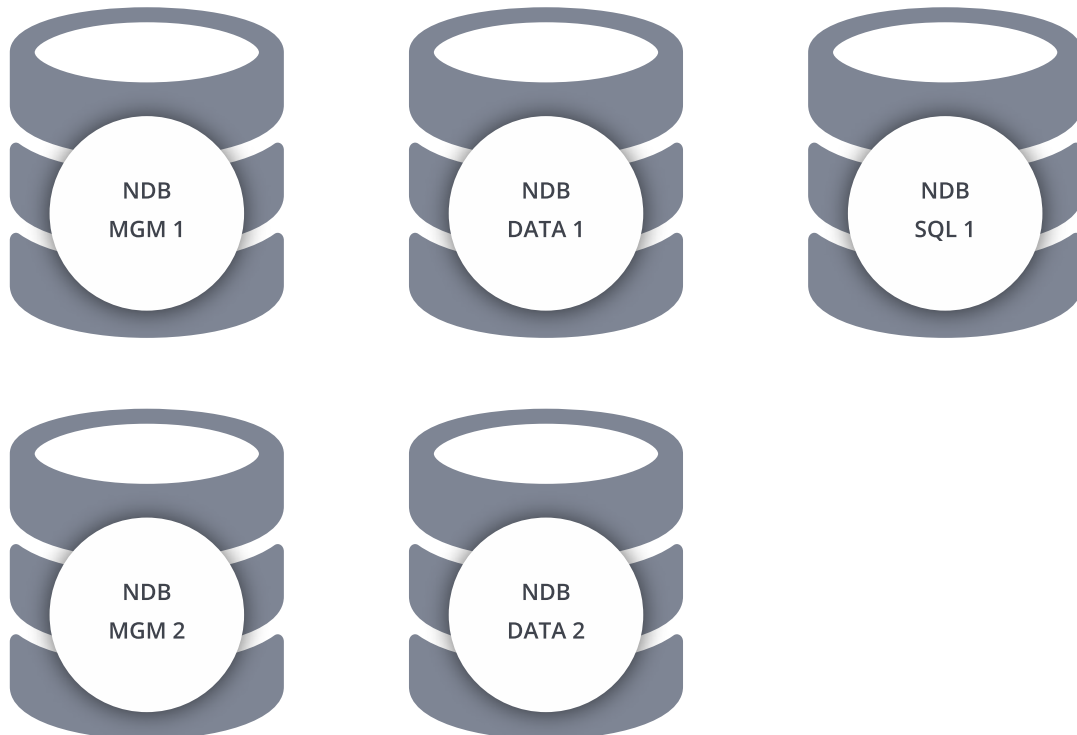
MySQL NDB Cluster is a high-availability storage engine for MySQL adapted for distributed computing environments. It consists of several elements: there are management servers, data nodes and SQL nodes.

Installation

To install the MySQL NDB Cluster packages manually, you can follow [this link](#).

Another way to install it is using [yum](#) or [apt](#) repositories.

In our example, let's see the yum repository installation of MySQL NDB Cluster 7.5 on CentOS 7.



We need to install 5 nodes:

- 2 Management Nodes: Management nodes are intended to control the cluster. IP Address: 192.168.100.175, 192.168.100.176.
- 2 Data Nodes: Data nodes stores the data using NDB engine. IP Address: 192.168.100.177, 192.168.100.178.
- 1 SQL Node: SQL nodes are used as the entry points to the cluster. They parse SQL, ask for data from the data nodes and aggregate result sets when needed. IP Address: 192.168.100.179.

For all nodes:

```
1 [root@WP5 ~]# wget https://dev.mysql.com/get/mysql80-community-release-el7-2.noarch.rpm
2 --2019-02-15 19:43:41-- https://dev.mysql.com/get/mysql80-community-release-el7-2.noarch.rpm
3 Resolving dev.mysql.com (dev.mysql.com)... 137.254.60.11
4 Connecting to dev.mysql.com (dev.mysql.com)|137.254.60.11|:443... connected.
5 HTTP request sent, awaiting response... 302 Found
6 Location: https://repo.mysql.com//mysql80-community-release-el7-2.noarch.rpm [following]
```

```

7  --2019-02-15 19:43:42-- https://repo.mysql.com//
mysql80-community-release-el7-2.noarch.rpm
8  Resolving repo.mysql.com (repo.mysql.com)... 23.208.182.226
9  Connecting to repo.mysql.com (repo.mysql.
com)|23.208.182.226|:443... connected.
10 HTTP request sent, awaiting response... 200 OK
11 Length: 25892 (25K) [application/x-redhat-package-manager]
12 Saving to: 'mysql80-community-release-el7-2.noarch.rpm'
13
14 100%[=====
=====
=====
=====>]
25,892      --.-K/s   in 0.01s
15
16 2019-02-15 19:43:42 (2.56 MB/s) - 'mysql80-community-re-
lease-el7-2.noarch.rpm' saved [25892/25892]
17
18 [root@WP5 ~]# rpm -Uvh mysql80-community-release-el7-2.
noarch.rpm
19 warning: mysql80-community-release-el7-2.noarch.rpm: Header
V3 DSA/SHA1 Signature, key ID 5072e1f5: NOKEY
20 Preparing...                               #####
##### [100%]
21 Updating / installing...
22  1:mysql80-community-release-el7-2  #####
##### [100%]

```

Edit the `/etc/yum.repos.d/mysql-community.repo` file and set the `enable` parameter in 1 for MySQL Cluster 7.5 and 0 for MySQL 8.0:

```

1  # Enable to use MySQL Cluster 7.5
2  [mysql-cluster-7.5-community]
3  name=MySQL Cluster 7.5 Community
4  baseurl=http://repo.mysql.com/yum/mysql-cluster-7.5-communi-
ty/el/7/$basearch/
5  enabled=1
6  gpgcheck=1
7  gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-mysql
8  [mysql80-community]
9  name=MySQL 8.0 Community Server
10 baseurl=http://repo.mysql.com/yum/mysql-8.0-community/
el/7/$basearch/
11 enabled=0
12 gpgcheck=1
13 gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-mysql

```

And then, install:

- For SQL Nodes (1):

```

1 [root@WP9 ~]# yum install mysql-cluster-communi-
  ty-server
2 =====
3 Package Arch Version
  Repository Size
4 =====
5 Installing:
6 mysql-cluster-community-server x86_64 7.5.13-
  1.e17 mysql-cluster-7.5-community 175 M
7 Installing for dependencies:
8 libaio x86_64 0.3.109-
  13.e17 base 24 k
9 mysql-cluster-community-client x86_64 7.5.13-
  1.e17 mysql-cluster-7.5-community 85 M
10 mysql-cluster-community-common x86_64 7.5.13-
  1.e17 mysql-cluster-7.5-community 274 k
11 mysql-cluster-community-libs x86_64 7.5.13-
  1.e17 mysql-cluster-7.5-community 2.2 M
12 numactl-libs x86_64 2.0.9-7.
  e17 base 29 k
13 perl-Class-MethodMaker x86_64 2.20-1.
  e17 epel 334 k
14 perl-Compress-Raw-Bzip2 x86_64 2.061-3.
  e17 base 32 k
15 perl-Compress-Raw-Zlib x86_64 1:2.061-
  4.e17 base 57 k
16 perl-DBI x86_64 1.627-4.
  e17 base 802 k
17 perl-Data-Dumper x86_64 2.145-3.
  e17 base 47 k
18 perl-IO-Compress noarch 2.061-2.
  e17 base 260 k
19 perl-Net-Daemon noarch 0.48-5.
  e17 base 51 k
20 perl-PlRPC noarch 0.2020-
  14.e17 base 36 k
21
22 Transaction Summary
23 =====
24 Install 1 Package (+13 Dependent packages)

```

- For Management Nodes (2):

```

1 [root@WP5 ~]# yum install mysql-cluster-community-man-

```

```

2  |  agement-server
3  |  =====
4  |  Package                      Arch
   |  Version                      Repository                Size
5  |  =====
6  |  Installing:
7  |  mysql-cluster-community-management-server  x86_64
   |  7.5.13-1.el7      mysql-cluster-7.5-community  4.8 M
8  |  Transaction Summary
9  |  =====
10 |  Install 1 Package

```

- For Data Nodes (2):

```

1  |  [root@WP7 ~]# yum install mysql-cluster-community-da-
   |  ta-node
2  |  =====
3  |  Package                      Arch      Version
   |  Repository                    Size
4  |  =====
5  |  Installing:
6  |  mysql-cluster-community-data-node  x86_64  7.5.13-
   |  1.el7      mysql-cluster-7.5-community  20 M
7  |  Transaction Summary
8  |  =====
9  |  =====
10 |  Install 1 Package

```

After this, you need to initialize the MySQL installation on the SQL Node.

```

1  |  [root@WP9 ~]# service mysqld start
2  |  Redirecting to /bin/systemctl start mysqld.service

```

A new password for the root user is created. To know it, you need to check the MySQL log and filter by temporary password. You should have something similar to this:

```

1  |  [root@WP9 ~]# grep "temporary password" /var/log/mysqld.log
2  |  2019-02-15T20:06:43.188506Z 1 [Note] A temporary password is
   |  generated for root@localhost: ESbw298Q1?z0

```

Then, you can run the `mysql_secure_installation` script, to configure a basic secure setup for your MySQL database.

```
1 [root@WP9 ~]# mysql_secure_installation
2 Securing the MySQL server deployment.
3
4 Enter password for user root:
5
6 The existing password for the user account root has expired.
  Please set a new password.
7
8 New password:
9
10 Re-enter new password:
11 The 'validate_password' plugin is installed on the server.
12 The subsequent steps will run with the existing configuration
13 of the plugin.
14 Using existing password for root.
15
16 Estimated strength of the password: 100
17 Change the password for root ? ((Press y|Y for Yes, any oth-
  er key for No) : y
18
19 New password:
20
21 Re-enter new password:
22
23 Estimated strength of the password: 100
24 Do you wish to continue with the password provided?(Press
  y|Y for Yes, any other key for No) : y
25 By default, a MySQL installation has an anonymous user,
26 allowing anyone to log into MySQL without having to have
27 a user account created for them. This is intended only for
28 testing, and to make the installation go a bit smoother.
29 You should remove them before moving into a production
30 environment.
31
32 Remove anonymous users? (Press y|Y for Yes, any other key
  for No) : y
33 Success.
34
35 Normally, root should only be allowed to connect from
36 'localhost'. This ensures that someone cannot guess at
37 the root password from the network.
38
39 Disallow root login remotely? (Press y|Y for Yes, any other
  key for No) : y
40 Success.
41
42 By default, MySQL comes with a database named 'test' that
```

```

43 | anyone can access. This is also intended only for testing,
44 | and should be removed before moving into a production
45 | environment.
46 |
47 | Remove test database and access to it? (Press y|Y for Yes,
48 | any other key for No) : y
49 |   - Dropping test database...
50 | Success.
51 |   - Removing privileges on test database...
52 | Success.
53 |
54 | Reloading the privilege tables will ensure that all changes
55 | made so far will take effect immediately.
56 |
57 | Reload privilege tables now? (Press y|Y for Yes, any other
58 | key for No) : y
59 | Success.
60 | All done!

```

Now, let's see the configuration.

Default Configuration

- SQL Node

The MySQL installation creates the my.cnf config file and the /etc/my.cnf.d/ directory in /etc/.

The /etc/my.cnf.d/ directory is empty by default, and in the content of my.cnf is:

```

1 | [root@WP9 ~]# cat /etc/my.cnf
2 | # For advice on how to change settings please see
3 | # http://dev.mysql.com/doc/refman/5.7/en/server-config-
4 | # uration-defaults.html
5 | [mysqld]
6 | #
7 | # Remove leading # and set to the amount of RAM for
8 | # the most important data
9 | # cache in MySQL. Start at 70% of total RAM for dedi-
10 | # cated server, else 10%.
11 | # innodb_buffer_pool_size = 128M
12 | #
13 | # Remove leading # to turn on a very important data
14 | # integrity option: logging
15 | # changes to the binary log between backups.
16 | # log_bin
17 | #

```



```

14 # Remove leading # to set options mainly useful for
    reporting servers.
15 # The server defaults are faster for transactions and
    fast SELECTs.
16 # Adjust sizes as needed, experiment to find the opti-
    mal values.
17 # join_buffer_size = 128M
18 # sort_buffer_size = 2M
19 # read_rnd_buffer_size = 2M
20 datadir=/var/lib/mysql
21 socket=/var/lib/mysql/mysql.sock
22 # Disabling symbolic-links is recommended to prevent
    assorted security risks
23 symbolic-links=0
24 log-error=/var/log/mysqld.log
25 pid-file=/var/run/mysqld/mysqld.pid

```

In this file, we need to add the following lines.

In the [mysqld] section:

```

1 | ndbcluster

```

And, a new section [mysql_cluster]:

```

1 | [mysql_cluster]
2 | ndb-connectstring=192.168.100.175,192.168.100.176 # IP
    address of Management Node

```

Let's see these parameters in detail.

- datadir: The path to the MySQL server data directory.
- socket: On Unix platforms, this variable is the name of the socket file that is used for local client connections.
- symbolic-links: Enable or disable symbolic link support. On Unix, enabling symbolic links means that you can link a MyISAM index file or data file to another directory with the INDEX DIRECTORY or DATA DIRECTORY option of the CREATE TABLE statement.
- log-error: Write the error log and startup messages to this file.
- pid-file: The path name of the file in which the server should write its process ID. The server creates the file in the data directory unless an absolute path name is given to specify a different directory.
- ndbcluster: The ndbcluster storage engine is necessary for using NDB Cluster.
- ndb-connectstring: When using the ndbcluster storage engine, this option specifies the management server that distributes cluster configuration data.

- Data Nodes

Create a new configuration file /etc/my.cnf in each data node:

```
1 [root@WP7 ~]# vi /etc/my.cnf
2 [mysqld]
3 ndbcluster
4 [mysql_cluster]
5 ndb-connectstring=192.168.100.175,192.168.100.176
6 #IP address of Management Nodes
```

And create the datadir:

```
1 [root@WP7 ~]# mkdir -p /usr/local/mysql/data
```

Let's see these parameters in detail.

- ndbcluster: The ndbcluster storage engine is necessary for using NDB Cluster.
 - ndb-connectstring: When using the ndbcluster storage engine, this option specifies the management server that distributes cluster configuration data.
- Management Nodes

```
1 [root@WP5 ~]# mkdir /var/lib/mysql-cluster
2 [root@WP5 ~]# cd /var/lib/mysql-cluster
3 [root@WP5 ~]# vi config.ini
4 [ndbd default]
5 # Default Options
6 NoOfReplicas=2
7 DataMemory=80M
8 [ndb_mgmd]
9 # Management Node 1
10 HostName=192.168.100.175
11 NodeId=1
12 DataDir=/var/lib/mysql-cluster
13 [ndb_mgmd]
14 # Management Node 2
15 HostName=192.168.100.176
16 NodeId=2
17 DataDir=/var/lib/mysql-cluster
18 [ndbd]
19 # Data Node 1
20 HostName=192.168.100.177
21 NodeId=3
22 DataDir=/usr/local/mysql/data
23 [ndbd]
24 # Data Node 2
25 HostName=192.168.100.178
```

```

26 | NodeId=4
27 | DataDir=/usr/local/mysql/data
28 | [mysqld]
29 | # SQL node
30 | NodeId=5
31 | HostName=192.168.100.179

```

Let's see these parameters in detail.

- **NoOfReplicas:** This global parameter can be set only in the [ndbd default] section, and defines the number of replicas for each table stored in the cluster.
- **DataMemory:** This parameter defines the amount of space available for storing database records. The amount specified by this value is allocated in memory, so it is important that the machine has sufficient physical memory to accommodate it.
- **HostName:** This parameter defines the hostname of the computer.
- **NodeId:** A unique node ID is used as the node's address for all cluster internal messages.
- **DataDir:** This parameter specifies the directory where the data files will be stored.

After the installation and configuration are completed, you should follow the initialization order to start the cluster. First, you should start the management nodes, then the data nodes and after that the SQL Nodes.

- Starting the Management Nodes

```

1 | [root@WP5 ~]# ndb_mgmd -f /var/lib/mysql-cluster/con-
  | fig.ini
2 | MySQL Cluster Management Server mysql-5.7.25 ndb-
  | 7.5.13
3 | 2019-02-19 00:25:17 [MgmtSrvr] INFO      -- The default
  | config directory '/usr/mysql-cluster' does not exist.
  | Trying to create it...
4 | 2019-02-19 00:25:17 [MgmtSrvr] INFO      -- Successfully
  | created config directory

```

- Starting the Data Nodes

```

1 | [root@WP7 ~]# ndbd
2 | 2019-02-19 00:29:29 [ndbd] INFO      -- Angel connected
  | to '192.168.100.175:1186'
3 | 2019-02-19 00:29:29 [ndbd] INFO      -- Angel allocated
  | nodeid: 3

```

- Starting the SQL Node

If we add the configuration for NDB Cluster after the initialization, we need to restart the MySQL service to apply the changes.

```
1 | [root@WP9 ~]# service mysqld restart
2 | Redirecting to /bin/systemctl restart mysqld.service
```

Then, we can check the cluster status using the `ndb_mgm` command from the SQL Node:

```
1 | [root@WP9 ~]# ndb_mgm
2 | -- NDB Cluster -- Management Client --
3 | ndb_mgm> SHOW
4 | Connected to Management Server at:
5 | 192.168.100.175:1186
6 | Cluster Configuration
7 | -----
8 | [ndbd(NDB)]      2 node(s)
9 | id=3    @192.168.100.177 (mysql-5.7.25 ndb-7.5.13,
10 | Nodegroup: 0, *)
11 | id=4    @192.168.100.178 (mysql-5.7.25 ndb-7.5.13,
12 | Nodegroup: 0)
13 | [ndb_mgmd(MGM)]  2 node(s)
14 | id=1    @192.168.100.175 (mysql-5.7.25 ndb-7.5.13)
15 | id=2    @192.168.100.176 (mysql-5.7.25 ndb-7.5.13)
16 | [mysqld(API)]   1 node(s)
17 | id=5    @192.168.100.179 (mysql-5.7.25 ndb-7.5.13)
```

Optional NDB Cluster Server Configuration

Let's see some variables that we can add in our config files. It's recommended as a basic configuration but it depends on your infrastructure.

```
1 | [TCP DEFAULT]
2 | SendBufferMemory=1M
3 | ReceiveBufferMemory=1M
4 | [NDB_MGMD DEFAULT]
5 | PortNumber=1186
6 | [NDB_MGMD]
7 | LogDestination=FILE:filename=ndb_1_cluster.log,max-
8 | size=10000000,maxfiles=6
9 | ArbitrationRank=1
10 | [NDBD DEFAULT]
11 | ServerPort=2200
12 | FileSystemPathDD=/var/lib/mysql-cluster
13 | BackupDataDir=/var/lib/mysql-cluster/backup
14 | FileSystemPathUndoFiles=/var/lib/mysql-cluster
15 | FileSystemPathDataFiles=/var/lib/mysql-cluster
16 | DataMemory=128M
17 | IndexMemory=21M
18 | LockPagesInMainMemory=1
19 | MaxNoOfConcurrentOperations=32768
```

```
19 MaxNoOfConcurrentTransactions=8192
20 StringMemory=25
21 MaxNoOfTables=2048
22 MaxNoOfOrderedIndexes=1024
23 MaxNoOfUniqueHashIndexes=256
24 MaxNoOfAttributes=12288
25 MaxNoOfTriggers=7168
26 MaxNoOfExecutionThreads=2
27 NoOfFragmentLogParts=4
28 FragmentLogFileSize=512M
29 InitFragmentLogFiles=SPARSE
30 NoOfFragmentLogFiles=3
31 RedoBuffer=8M
32 TransactionBufferMemory=8M
33 TimeBetweenGlobalCheckpoints=1000
34 TimeBetweenEpochs=100
35 TimeBetweenEpochsTimeout=32000
36 MinDiskWriteSpeed=20M
37 MaxDiskWriteSpeed=80M
38 MaxDiskWriteSpeedOtherNodeRestart=50M
39 MaxDiskWriteSpeedOwnRestart=200M
40 TimeBetweenLocalCheckpoints=20
41 HeartbeatIntervalDbDb=1500
42 HeartbeatIntervalDbApi=1500
43 MemReportFrequency=30
44 BackupReportFrequency=10
45 LogLevelStartup=15
46 LogLevelShutdown=15
47 LogLevelCheckpoint=8
48 LogLevelNodeRestart=15
49 BackupMaxWriteSize=1M
50 BackupDataBufferSize=24M
51 BackupLogBufferSize=16M
52 TimeBetweenWatchdogCheckInitial=60000
53 TransactionInactiveTimeout=60000
54 RedoOverCommitCounter=3
55 RedoOverCommitLimit=20
56 SharedGlobalMemory=20M
57 DiskPageBufferMemory=8M
58 BatchSizePerLocalScan=512
59 [MYSQLD DEFAULT]
60 DefaultOperationRedoProblemAction=ABORT
61 BatchSize=512
```

To see these variables in detail, you can follow [this link](#).

MongoDB

MongoDB is an open-source document database that provides high performance, high availability, and automatic scaling. Classified as a NoSQL database, MongoDB uses JSON-like documents with schema. MongoDB is developed by MongoDB Inc.



Installation

There are many ways to install MongoDB, you can follow [this link](#) to choose one.

You can install MongoDB by using [yum](#) or [apt](#) repositories.

In our example, let's see the yum repository installation of MongoDB on CentOS 7.

MongoDB Server IP Address: 192.168.100.181

To add the MongoDB repository, you can run:

```
1 cat > /etc/yum.repos.d/mongodb-org-4.0.repo <<- EOF
2 [mongodb-org-4.0]
3 name=MongoDB Repository
4 baseurl=https://repo.mongodb.org/yum/redhat/7/mongod
  b-org/4.0/x86_64/
5 gpgcheck=1
6 enabled=1
7 gpgkey=https://www.mongodb.org/static/pgp/server-4.0.asc
8 EOF
```

And then, install the mongodb-org packages:

```
1 [root@WP10 ~]# yum install -y mongodb-org
2 =====
3 Package Arch Version
4 Repository Size
5 =====
5 Installing:
6 mongodb-org x86_64 4.0.6-1.el7
6 mongodb-org-4.0 5.8 k
7 Installing for dependencies:
8 make x86_64 1:3.82-23.el7
8 base 420 k
9 mongodb-org-mongos x86_64 4.0.6-1.el7
9 mongodb-org-4.0 12 M
10 mongodb-org-server x86_64 4.0.6-1.el7
10 mongodb-org-4.0 21 M
11 mongodb-org-shell x86_64 4.0.6-1.el7
```

```

12  |  mongodb-org-4.0                13 M
    |  mongodb-org-tools             x86_64  4.0.6-1.el7
    |  mongodb-org-4.0                32 M
13  |  openssl                         x86_64  1:1.0.2k-16.
    |  el7 base                        493 k
14  |  Updating for dependencies:
15  |  openssl-libs                   x86_64  1:1.0.2k-16.
    |  el7 base                          1.2 M
16  |
17  |  Transaction Summary
18  |  =====
    |  =====
19  |  Install 1 Package (+6 Dependent packages)
20  |  Upgrade ( 1 Dependent package)

```

Now, you need to start the MongoDB service.

```

1  | [root@WP10 ~]# service mongod start
2  | Redirecting to /bin/systemctl start mongod.service

```

You can verify the status by filtering the “waiting” word in the log file:

```

1  | [root@WP10 ~]# grep "waiting" /var/log/mongodb/mongod.log
2  | 2019-02-20T00:40:37.449+0000 I NETWORK [initandlisten]
    | waiting for connections on port 27017

```

Default Configuration

The MongoDB installation creates the `/etc/mongod.conf` config file.

```

1  | [root@WP10 ~]# cat /etc/mongod.conf
2  | # mongod.conf
3  | # for documentation of all options, see:
4  | # http://docs.mongodb.org/manual/reference/configuration-options/
5  | # where to write logging data.
6  | systemLog:
7  |   destination: file
8  |   logAppend: true
9  |   path: /var/log/mongodb/mongod.log
10 | # Where and how to store data.
11 | storage:
12 |   dbPath: /var/lib/mongo
13 |   journal:
14 |     enabled: true
15 | # engine:
16 | # mmapv1:

```

```

17 # wiredTiger:
18 # how the process runs
19 processManagement:
20   fork: true # fork and run in background
21   pidFilePath: /var/run/mongodb/mongod.pid # location of
    pidfile
22   timeZoneInfo: /usr/share/zoneinfo
23 # network interfaces
24 net:
25   port: 27017
26   bindIp: 127.0.0.1 # Enter 0.0.0.0,:: to bind to all IPv4
    and IPv6 addresses or, alternatively, use the net.bindIpAll
    setting.
27 #security:
28 #operationProfiling:
29 #replication:
30 #sharding:
31 ## Enterprise-Only Options
32 #auditLog:
33 #snmp:

```

Let's see these parameters in detail.

- systemLog:
 - destination: The destination to which MongoDB sends all log output.
 - logAppend: When true, MongoDB appends new entries to the end of the existing log file when the instance restarts.
 - path: Specify a log file path.
- storage:
 - dbPath: Specify a data directory path.
 - journal: Enable or disable the durability journal to ensure data files remain valid and recoverable.
- processManagement:
 - fork: Enable a daemon mode that runs the MongoDB process in the background.
 - pidFilePath: Specifies a file location to hold the process ID of the MongoDB process.
 - timeZoneInfo: The full path from which to load the time zone database.
- net:
 - port: The TCP port on which the MongoDB instance listens for client connections.
 - bindIp: The hostnames and/or IP addresses and/or full Unix domain socket paths on which MongoDB should listen for client connections.

Optional MongoDB Configuration

Let's see some variables that we can add in our config files. It's recommended as a basic configuration but it depends on your infrastructure.

```
1 storage:
2   engine: wiredTiger
3   mmapv1:
4     smallFiles: false
5 net:
6   bindIp: 0.0.0.0
7 setParameter:
8   enableLocalhostAuthBypass: true
9 replication:
10  replSetName: my_mongodb_0
11 sharding:
12  clusterRole: shardsvr
13 security.keyFile: /etc/mongo-cluster.key
```

To see in detail these variables, you can follow [this link](#).

Percona Server for MongoDB

Percona Server for MongoDB is a free and open-source drop-in replacement for MongoDB Community Edition. It offers all the features and benefits of MongoDB Community Edition, plus additional enterprise-grade functionality.



Installation

To install the Percona Server for MongoDB packages manually, you can follow [this link](#).

Another way to install it is using [yum](#) or [apt](#) repositories.

In our example, let's see the yum repository installation of Percona Server for MongoDB on CentOS 7.

Percona MongoDB Server IP Address: 192.168.100.182

Install the repository:

```
1 [root@WP11 ~]# yum install https://repo.percona.com/yum/per-
2   =====
3   Package                               Arch      Version
   Repository                             Size
```

```

4  =====
5  Installing:
6  percona-release                noarch    1.0-7
   /percona-release-latest.noarch  18 k
7
8  Transaction Summary
9  =====
10 Install 1 Package

```

Enable the percona repository:

```

1  [root@WP11 ~]# percona-release enable psmdb-40 release
2  * Enabling the Percona Original repository
3  <*> All done!

```

And then, install percona-server-mongodb:

```

1  [root@WP11 ~]# yum install percona-server-mongodb
2  =====
3  Package                Arch      Version
4  Repository              Size
5  =====
6  Installing:
7  percona-server-mongodb  x86_64   4.0.5-2.el7
8  psmdb-40-release-x86_64 4.8 k
9  Installing for dependencies:
10 libpcap                 x86_64   14:1.5.3-11.
    el7 base                138 k
11 numactl                 x86_64   2.0.9-7.el7
    base                    66 k
12 numactl-libs            x86_64   2.0.9-7.el7
    base                    29 k
13 percona-server-mongodb-mongos x86_64   4.0.5-2.el7
    psmdb-40-release-x86_64 8.7 M
14 percona-server-mongodb-server x86_64   4.0.5-2.el7
    psmdb-40-release-x86_64 18 M
15 percona-server-mongodb-shell x86_64   4.0.5-2.el7
    psmdb-40-release-x86_64 9.7 M
16 percona-server-mongodb-tools x86_64   4.0.5-2.el7
    psmdb-40-release-x86_64 26 M
17
18 Transaction Summary
19 =====
20 Install 1 Package (+7 Dependent packages)

```

Now, you need to start the Percona MongoDB service.

```
1 [root@WP11 ~]# service mongod start
2 Redirecting to /bin/systemctl start mongod.service
```

You can verify the status by filtering the “waiting” word in the log file:

```
1 [root@WP11 ~]# grep “waiting” /var/log/mongo/mongod.log
2 2019-02-20T02:11:29.790+0000 I NETWORK [initandlisten]
   waiting for connections on port 27017
```

Default Configuration

The Percona MongoDB installation creates the `/etc/mongod.conf` config file.

```
1 [root@WP11 ~]# cat /etc/mongod.conf
2 # mongod.conf, Percona Server for MongoDB
3 # for documentation of all options, see:
4 #   http://docs.mongodb.org/manual/reference/configuration-options/
5 # Where and how to store data.
6 storage:
7   dbPath: /var/lib/mongo
8   journal:
9     enabled: true
10  # engine: mmapv1
11  # engine: wiredTiger
12  # engine: inMemory
13  # Storage engine various options
14  # More info for mmapv1: https://docs.mongodb.com/v4.0/reference/configuration-options/#storage-mmapv1-options
15  # mmapv1:
16  #   preallocDataFiles: true
17  #   nsSize: 16
18  #   quota:
19  #     enforced: false
20  #     maxFilesPerDB: 8
21  #     smallFiles: false
22  # More info for wiredTiger: https://docs.mongodb.com/v4.0/reference/configuration-options/#storage-wiredtiger-options
23  # wiredTiger:
24  #   engineConfig:
25  #     cacheSizeGB: 1
26  #     checkpointSizeMB: 1000
27  #     statisticsLogDelaySecs: 0
28  #     journalCompressor: snappy
29  #     directoryForIndexes: false
30  #   collectionConfig:
```

```

31 #     blockCompressor: snappy
32 #     indexConfig:
33 #     prefixCompression: true
34 # More info for inMemory: https://www.percona.com/doc/percona-server-for-mongodb/4.0/inmemory.html#configuring-percona-memory-engine
35 # inMemory:
36 #     engineConfig:
37 #     inMemorySizeGB: 1
38 #     statisticsLogDelaySecs: 0
39 # Two options below can be used for wiredTiger and inMemory storage engines
40 #setParameter:
41 #     wiredTigerConcurrentReadTransactions: 128
42 #     wiredTigerConcurrentWriteTransactions: 128
43 # where to write logging data.
44 systemLog:
45     destination: file
46     logAppend: true
47     path: /var/log/mongo/mongod.log
48 processManagement:
49     fork: true
50     pidFilePath: /var/run/mongod.pid
51 # network interfaces
52 net:
53     port: 27017
54     bindIp: 127.0.0.1
55 #security:
56 #operationProfiling:
57 #replication:
58 #sharding:
59 ## Enterprise-Only Options:
60 #auditLog:
61 #snmp:

```

Let's see these parameters in detail.

- systemLog:
 - destination: The destination to which Percona MongoDB sends all log output.
 - logAppend: When true, Percona MongoDB appends new entries to the end of the existing log file when the instance restarts.
 - path: Specify a log file path.
- storage:
 - dbPath: Specify a data directory path.
 - journal: Enable or disable the durability journal to ensure data files remain valid and recoverable.

- processManagement:
 - fork: Enable a daemon mode that runs the Percona MongoDB process in the background.
 - pidFilePath: Specifies a file location to hold the process ID of the MongoDB process.
 - timeZoneInfo: The full path from which to load the time zone database.
- net:
 - port: The TCP port on which the Percona MongoDB instance listens for client connections.
 - bindIp: The hostnames and/or IP addresses and/or full Unix domain socket paths on which Percona MongoDB should listen for client connections.

Optional Percona MongoDB Configuration

Let's see some variables that we can add in our config files. It's recommended as a basic configuration but it depends on your infrastructure.

```

1 storage:
2     engine: wiredTiger
3     mmapv1:
4         smallFiles: false
5 net:
6     bindIp: 0.0.0.0
7 setParameter:
8     enableLocalhostAuthBypass: true
9 replication:
10    replSetName: my_mongodb_0
11 sharding:
12    clusterRole: shardsvr
13 security.keyFile: /etc/mongo-cluster.key

```

To see in detail these variables, you can follow [this link](#).

PostgreSQL

PostgreSQL has earned a strong reputation for its proven architecture, reliability, data integrity, robust feature set, extensibility, and the dedication of the open source community behind the software to consistently deliver performant and innovative solutions.



Installation

To install the PostgreSQL packages manually, you can follow [this link](#).

Another way to install it is using [yum](#) or [apt](#) repositories.

In our example, let's see the yum repository installation of PostgreSQL 11 on CentOS 7.

PostgreSQL Server IP Address: 192.168.100.185

Install the repository:

```
1 [root@WP12 ~]# yum install https://download.postgresql.org/
2 pub/repos/yum/11/redhat/rhel-7-x86_64/pgdg-centos11-11-2.
3 noarch.rpm
4 =====
5 Package Arch Version
6 Repository Size
7 =====
8 Installing:
9 pgdg-centos11 noarch 11-2
10 /pgdg-centos11-11-2.noarch 2.7 k
11
12 Transaction Summary
13 =====
14
15 Install 1 Package
```

Installing PostgreSQL client:

```
1 [root@WP12 ~]# yum install postgresql11
2 =====
3 Package Arch Version
4 Repository Size
5 =====
6 Installing:
7 postgresql11 x86_64 11.2-1PGDG.
8 rhel7 pgdg11 1.6 M
9 Installing for dependencies:
10 libicu x86_64 50.1.2-17.el7
11 base 6.9 M
12 postgresql11-libs x86_64 11.2-1PGDG.
13 rhel7 pgdg11 360 k
14
15 Transaction Summary
16 =====
17
18 Install 1 Package (+2 Dependent packages)
```

Installing PostgreSQL server:

```
1 [root@WP12 ~]# yum install postgresql11-server
2 =====
3 Package Arch Version
4 Repository Size
5 =====
5 Installing:
6 postgresql11-server x86_64 11.2-1PGDG.
7 rhel7 pgdg11 4.7 M
8
8 Transaction Summary
9 =====
10 Install 1 Package
```

You can initialize your PostgreSQL database:

```
1 [root@WP12 ~]# /usr/pgsql-11/bin/postgresql-11-setup initdb
2 Initializing database ... OK
```

Enable the PostgreSQL service:

```
1 [root@WP12 ~]# systemctl enable postgresql-11
2 Created symlink from /etc/systemd/system/multi-user.target.wants/postgresql-11.service to /usr/lib/systemd/system/postgresql-11.service.
```

Now, you can start the PostgreSQL service.

```
1 [root@WP12 ~]# systemctl start postgresql-11
```

Default Configuration

In the PostgreSQL datadir, by default `/var/lib/pgsql/11/data/`, you have different configuration files:

- `pg_hba.conf`: Client authentication is controlled by this file.
- `pg_ident.conf`: User name maps are defined in this ident map file.
- `postgresql.conf`: It's the main server configuration file.
- `postmaster.opts`: A file recording the command-line options the server was last started with.

Let's see these files one by one.

- pg_hba.conf

```

1  # TYPE  DATABASE          USER              ADDRESS
   METHOD
2  # "local" is for Unix domain socket connections only
3  local   all              all
   peer
4  # IPv4 local connections:
5  host    all              all              127.0.0.1/32
   ident
6  # IPv6 local connections:
7  host    all              all              ::1/128
   ident
8  # Allow replication connections from localhost, by a
   user with the
9  # replication privilege.
10 local   replication   all
   peer
11 host    replication   all              127.0.0.1/32
   ident
12 host    replication   all              ::1/128
   ident

```

- pg_ident.conf

```

1  # MAPNAME          SYSTEM-USERNAME   PG-USERNAME

```

- postgresql.conf (We will only see the uncommented lines for space reasons)

```

1  max_wal_size = 1GB
2  min_wal_size = 80MB
3  log_timezone = 'UTC'
4  datestyle = 'iso, mdy'
5  timezone = 'UTC'
6  default_text_search_config = 'pg_catalog.english'

```

- postmaster.opts

```

1  /usr/pgsql-11/bin/postgres "-D" "/var/lib/pgsql/11/
   data/"

```

Let's see the parameters in detail.

- max_wal_size: Maximum size the WAL is allowed to grow between the control points.
- min_wal_size: When the WAL file is kept below this value, it is recycled for future use at a checkpoint, instead of being deleted.
- log_timezone: Sets the time zone used for timestamps written in the server log.
- datestyle: Sets the display formats for date and time values, as well as the rules

for interpreting ambiguous date input values.

- `timezone`: Sets the time zone for displaying and interpreting time stamps.
- `default_text_search_config`: Selects the text search configuration that is used by those variants of the text search functions that do not have an explicit argument specifying the configuration.

To see these variables in detail, you can follow the [official documentation](#).



Synopsis

We have showed you some examples of how to install, configure and secure some of the most popular DB engines. To be able to perform get to the above procedures, you need to research, test, and analyse your available resources in order to get a well tuned deploy.

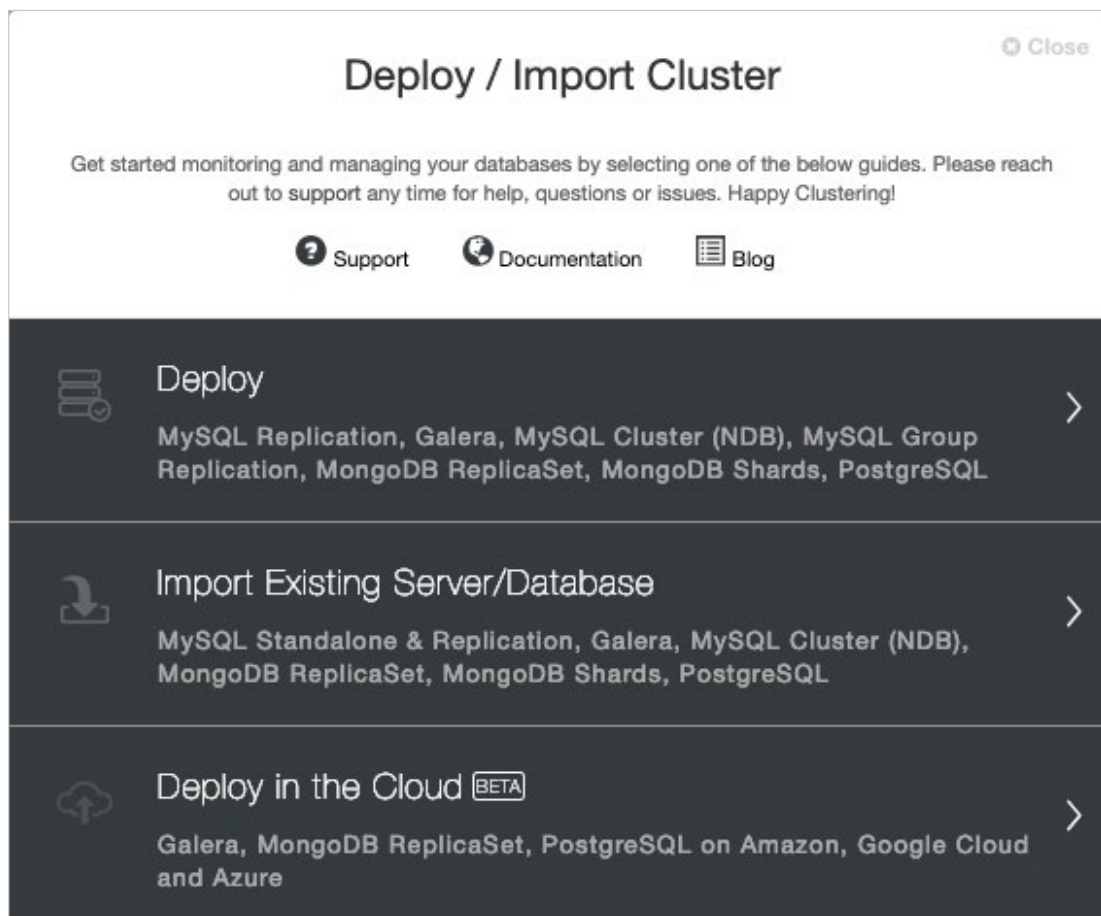
We will now look into ClusterControl, how it handles these tasks for you, delivering a secure well tuned environment.

How to Deploy Open Source Databases by Using ClusterControl

After seeing how we can deploy some of the most common open source databases manually, let's see how ClusterControl can make our lives easier.

Deploy

To perform a new installation from ClusterControl, simply select the option "Deploy" and follow the instructions that appear. Note that if you already have a database instance running, then you need to select the 'Import Existing Server/Database' instead.



There are some differences depending on which technology we want to deploy.

Deploy Database Cluster
Close

MySQL Replication
MySQL Galera
MySQL Cluster (NDB)
NEW
TimescaleDB
PostgreSQL
MongoDB ReplicaSet
MongoDB Shards

1 General & SSH Settings
2 Define PostgreSQL Servers
3 Define Topology
4 Deployment Summary

SSH User ?

SSH Key Path ?

Sudo Password ?

SSH Port ?

Cluster Name ?

x
✓
Install Software

x
✓
Disable Firewall?

x
✓
Disable AppArmor/SELinux?

i Use clean and minimal VMs. Existing package dependencies might be removed if required. New packages will be installed and existing packages can be uninstalled when provisioning the node with required software.

Back
Continue

In the deploy section, we need first to select the database technology, then, we must specify User, Key or Password and port to connect by SSH to our new database host. We also need a name for our new cluster and if we want ClusterControl to install the corresponding software and configurations for us.

Proper passwordless SSH setup from ClusterControl node to all nodes (including ClusterControl node) is mandatory. Before performing any operation on the managed node, the node must be accessible via SSH without using password but using key-based authentication instead.

Deploy Database Cluster
Close

MySQL Replication
MySQL Galera
MySQL Cluster (NDB)
NEW
TimescaleDB
PostgreSQL
MongoDB ReplicaSet
MongoDB Shards

✓ 1 General & SSH Settings
2 Define MySQL Servers
3 Define Topology

Vendor ?

Version ?

Server Data Directory ?

Server Port ?

Configuration Template ?

Admin/Root Password ?

Repository

i Setup and use the vendor's repositories - learn more.

Back
Continue

After setting up the SSH access information, we must define the database information, like vendor, version and user. We can also specify which repository to use.

The asked information will be different depending on which database technology we selected.

The screenshot shows the 'Deploy Database Cluster' wizard at step 2, 'Define PostgreSQL Servers'. The top navigation bar includes options for MySQL Replication, MySQL Galera, MySQL Cluster (NDB), TimescaleDB (marked as NEW), PostgreSQL (selected), MongoDB ReplicaSet, and MongoDB Shards. The progress bar shows four steps: 1. General & SSH Settings (completed), 2. Define PostgreSQL Servers (current), 3. Define Topology, and 4. Deployment Summary. The main form contains the following fields:

- Server Port:** Input field with value 5432.
- User:** Input field with value admin.
- Password:** Password input field with masked characters and a visibility toggle.
- Version:** Selection buttons for 9.6, 10, and 11.
- Datadir:** Input field with placeholder text 'Enter directory (optional)'. A note below it says 'Setup and use the vendor's repositories - learn more.'
- Repository:** A dropdown menu currently set to 'Use Vendor Repositories'.

At the bottom, there are 'Back' and 'Continue' buttons.

In the next step, we need to specify will be our cluster topology.

The screenshot shows the 'Deploy Database Cluster' wizard at step 3, 'Define Topology'. The top navigation bar is the same as in the previous step, but the 'PostgreSQL' option is now highlighted. The progress bar shows: 1. General & SSH Settings (completed), 2. Define PostgreSQL Servers (completed), 3. Define Topology (current), and 4. Deployment Summary. The main form is titled 'Master A - IP/Hostname' and displays a list of three server entries:

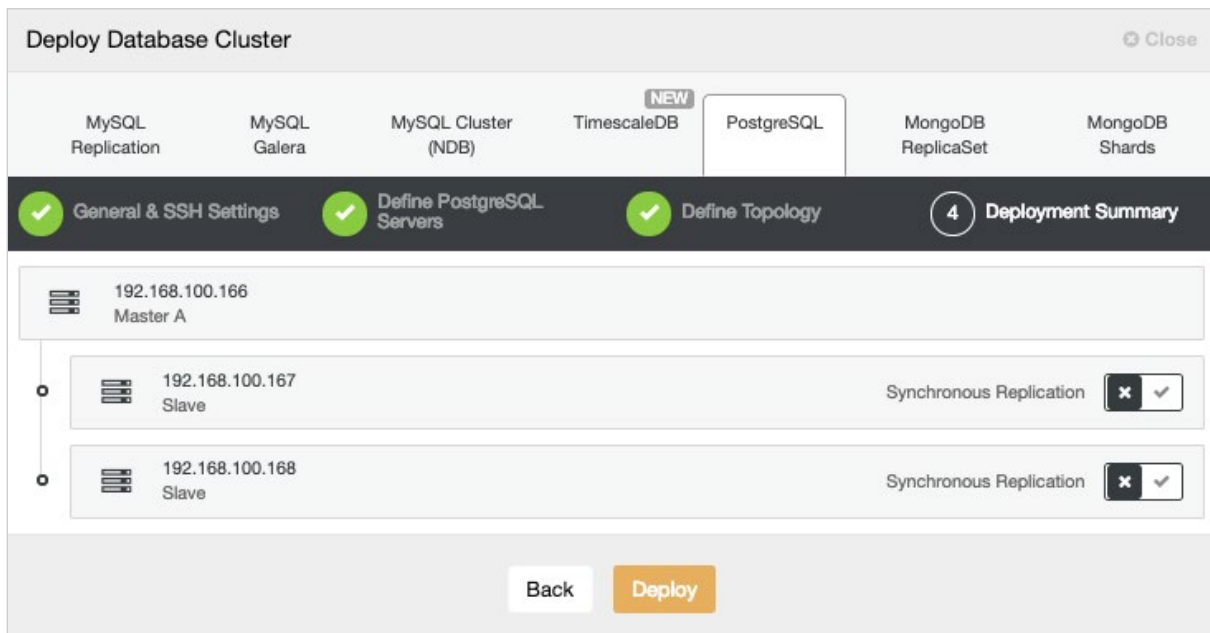
- 192.168.100.166 (checked)
- 192.168.100.167 (checked)
- 192.168.100.168 (checked)

Each entry has a green checkmark on the left and a close button (X) on the right. A vertical line connects the first entry to the second and third, indicating a topology. Below this list is a dashed box titled 'Add slaves to master A' containing an 'Add Slave' button.

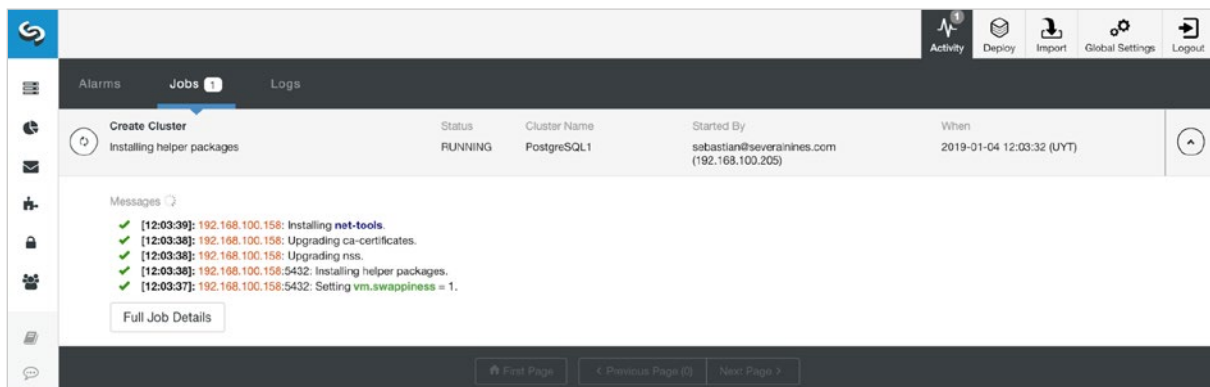
At the bottom, there are 'Back' and 'Continue' buttons.

When adding our servers, we can enter IP or hostname.

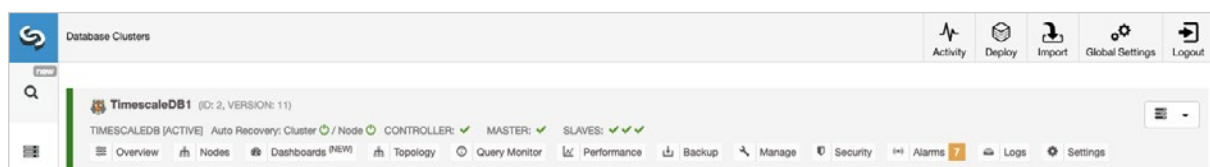
In the last step, we can choose if our replication will be Synchronous or Asynchronous.



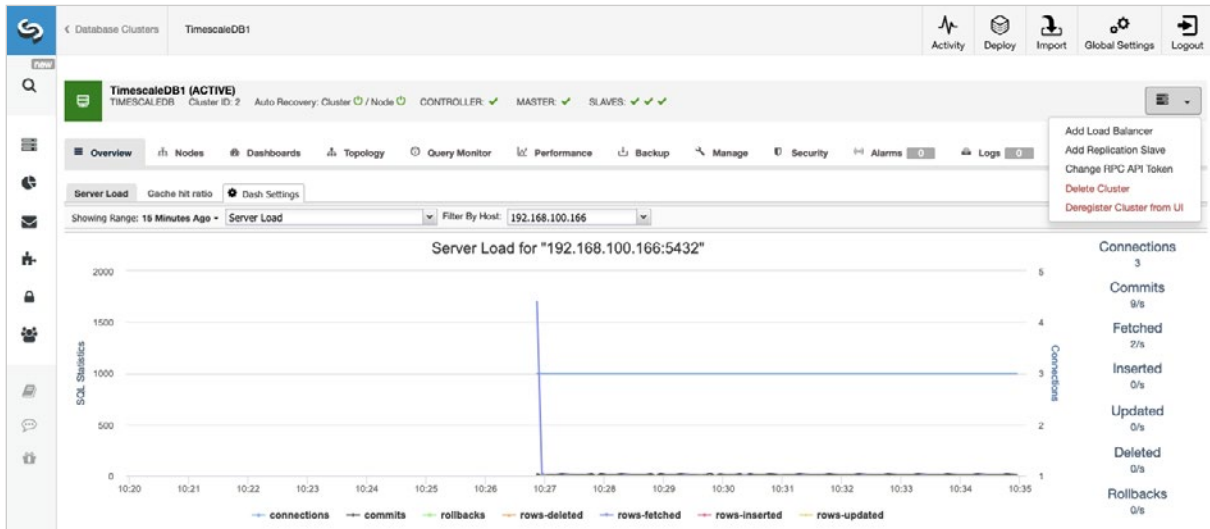
We can monitor the status of the creation of our new cluster from the ClusterControl activity monitor.



Once the task is finished, we can see our new cluster in the main ClusterControl screen.



Once we have our cluster created, we can perform several tasks on it, like adding a load balancer (HAProxy) or a new replica.



Scaling

If we go to cluster actions and select "Add Replication Slave", we can either create a new replica from scratch, or add an existing database as a replica.

Add Replication Slave ✕ Close

1
2

Add new Replication slave

Import existing Replication Slave

Cancel
Continue

Let's see how adding a new replication slave can be a really easy task.

Add Replication Slave ✕ Close

✓ 2

Master server:

Slave hostname:

Slave port:

Use Package Default for Datadir

Install PostgreSQL software:

Synchronous Replication:

The slave will be staged with data from the master. The master's configuration will be altered to allow the slave to join the master.

The slave server must be reachable by SSH (key-based auth.) from the controller.

As you can see in the image, we only need to choose our Master server, enter the IP address for our new slave server and the database port. Then, we can choose if we want ClusterControl to install the software for us, and if the replication slave should be Synchronous or Asynchronous.

In this way, we can add as many replicas as we want and spread read traffic between them using a load balancer, which we can also implement with ClusterControl.

Load Balancing

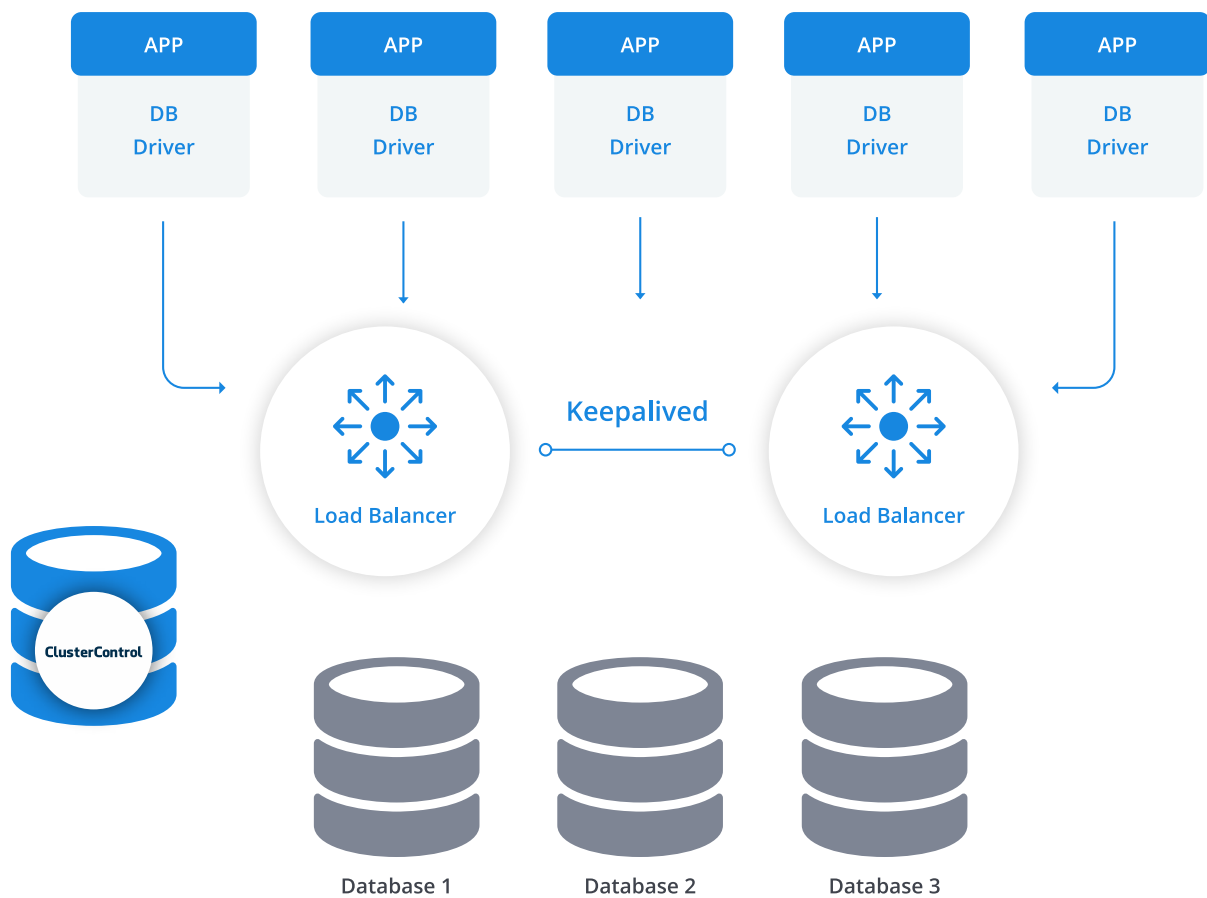
Load balancers can be used to manage the traffic from your application, to get the most out of your database architecture.

Not only is it useful for balancing the load of our databases, it also helps applications get redirected to the available/healthy nodes and even specify ports with different roles.

With ClusterControl, you can deploy ProxySQL, HAProxy or MaxScale as Load Balancer.

By using a load balancer, you can distribute the traffic from one origin to one or more destinations and can define specific rules and/or protocols for this task. If any of the destinations stops responding, it is marked as offline, and the traffic is sent to the rest of the available destinations.

Keepalived is a service that allows you to configure a virtual IP within an active/passive group of servers. This virtual IP is assigned to an active server. If this server fails, the IP is automatically migrated to the "Secondary" passive server, allowing it to continue working with the same IP in a transparent way for the systems.



To perform a load balancer deployment, select the option "Add Load Balancer" in the cluster actions and fill the asked information.

Load Balancer
Deploy & Import Load Balancers

ProxySQL HAProxy **Keepalived** Garbd MaxScale

Deploy HAProxy Import HAProxy

Choose where to install

Server Address: Policy:

Listen Port (Read/Write):

Install for read/write splitting (master-slave replication)

Installation Settings

Build from Source (latest available source package will be used)

Overwrite Existing /usr/local/sbin/mysqldchk on targets

Show advanced settings

To perform a keepalived deployment, select the cluster, go to “Manage” menu and “Load Balancer” section, and then select “Keepalived” option.

Load Balancer
Deploy & Import Load Balancers

ProxySQL HAProxy **Keepalived** Garbd MaxScale

Deploy Keepalived Import Keepalived

Load balancer type:

Keepalived 1:

Type a virtual host Address

Virtual IP:

Network Interface:

For your HA environment, you need to select the load balancer servers and the virtual IP address.

Keepalived uses a virtual IP and migrates it from one load balancer to another in case of failure, so your setup can continue to function normally.

Management

From ClusterControl, you can also perform different management tasks like scheduling backups and verifying them for integrity, automatic failover, encryption of traffic, topology changes, and so on. The options depend on the database engine that you are using.

The screenshot displays the ClusterControl management interface for a Galera1 cluster. The top navigation bar shows the cluster status as 'Galera1 (ACTIVE)' with various components like CONTROLLER, GALERA NODES, MASTER, KEEPALIVED, and PROXYSQL all in a 'checked' state. The main interface is divided into several sections:

- Controller:** Shows the Controller Node at 192.168.100.110.
- MySQL:** Lists Galera Nodes, with the Master Node at 192.168.100.120 highlighted as 'Synced (Primary) | Master'. Other nodes at 192.168.100.121 and 192.168.100.122 are also 'Synced (Primary)'.
- Keepalived:** Shows nodes at 192.168.100.111.
- Proxysql:** Shows nodes at 192.168.100.111 and 192.168.100.112.
- Prometheus:** Shows Prometheus Nodes.

The central focus is the **Master Node** (192.168.100.120) overview, which includes:

- A status indicator: 'Node is OK'.
- Navigation tabs: Overview, Top, DB Performance, DB Status, and DB Variables.
- A 'Managed Process' section showing 'mysqld <pid: 1433>' is running.
- A 'Node Stat' section with a 'Show Range' dropdown set to '1 Hour Ago'.
- Two performance graphs: 'CPU Usage' (All cores) and 'Network Usage' (Interface eth0).

On the right side, a 'Node Actions' dropdown menu is open, listing the following management tasks:

- SSH Console
- Schedule Maintenance Mode
- Reboot Host
- Restart Node
- Resync Node
- Bootstrap Cluster
- Disable Binary Logging
- Stop Node
- Unregister Node



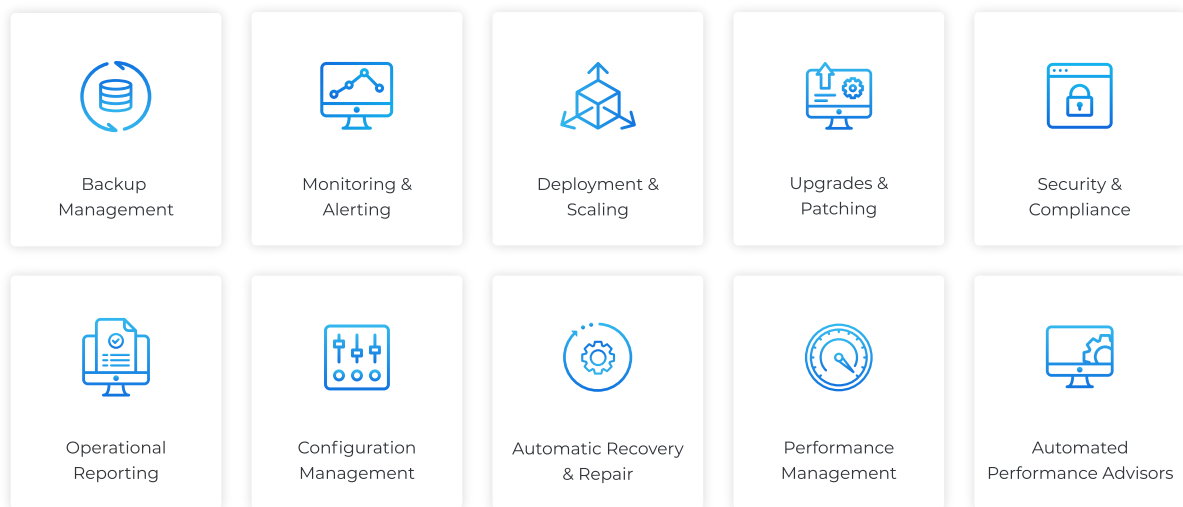
Conclusion

In this Whitepaper, we have listed some of the most popular DB engines. After that we went through some examples on how to install, tune and secure each of them. Getting to those step by step procedures involves research, experimentation and testing. You need to understand each step of what you're doing and the available options, and research them to pick the correct ones for your needs.

We finally introduced ClusterControl, a system that can help you close that gap and get a well secured, well tuned deploy of your chosen engine.

About ClusterControl

ClusterControl is the all-inclusive open source database management system for users with mixed environments that removes the need for multiple management tools. ClusterControl provides advanced deployment, management, monitoring, and scaling functionality to get your MySQL, MongoDB, and PostgreSQL databases up-and-running using proven methodologies that you can depend on to work. At the core of ClusterControl is its automation functionality that lets you automate many of the database tasks you have to perform regularly like deploying new databases, adding and scaling new nodes, running backups and upgrades, and more.



About Severalnines

Severalnines provides automation and management software for database clusters. We help companies deploy their databases in any environment, and manage all operational aspects to achieve high-scale availability.

Severalnines' products are used by developers and administrators of all skills levels to provide the full 'deploy, manage, monitor, scale' database cycle, thus freeing them from the complexity and learning curves that are typically associated with highly available database clusters. Severalnines is often called the "anti-startup" as it is entirely self-funded by its founders. The company has enabled over 12,000 deployments to date via its popular product ClusterControl. Currently counting BT, Orange, Cisco, CNRS, Technicolor, AVG, Ping Identity and Paytrail as customers. Severalnines is a private company headquartered in Stockholm, Sweden with offices in Singapore, Japan and the United States. To see who is using Severalnines today visit:

<https://www.severalnines.com/company>

Related Resources



Deployment & Scaling with ClusterControl

ClusterControl provides a suite of database deployment tools, allowing cluster deployment, database importing, load balancing, hybrid deployments and more!



How to Deploy a Production-Ready MySQL or MariaDB Galera Cluster using ClusterControl

ClusterControl can be used to deploy open-source database clusters that are configured in complex topologies. The full high availability stack includes both database and proxy layers.

In this blog, we are going to show you how to deploy a production-grade Galera Cluster, complete with load balancers, for a high availability setup.



How to Deploy PostgreSQL for High Availability

In this blog, we will review some important concepts of High Availability, possible database HA architectures and useful components when implementing PostgreSQL HA. Then we will see how to use ClusterControl to deploy an entire high availability stack for PostgreSQL.



How to Deploy MongoDB for High Availability

MongoDB provides ReplicaSets to help you address high availability database requirements. Although support for replication and failover is built-in, it is not enough for the database to be considered production-ready. That requires a set of policies and procedures, like getting alerted in case of performance slowdowns, anomalies or failures in a live environment. Backups are essential. Being able to automatically recover from different types of failures can drastically reduce downtime.

Choosing which DB engine to use between all the options we have today is not an easy task. And that is just the beginning. After deciding which engine to use, you need to learn about it and actually deploy it to play with it. We plan to help you on that second step, and show you how to install, configure and secure some of the most popular open source DB engines. In this whitepaper we are going to cover these points, with the aim of fast tracking you on the deploy task.

