



# PERCONA

## Percona Server Documentation

8.1.0-1 (2023-11-27)

*Percona Technical Documentation Team*

<a href="https://www.percona.com/about">Percona LLC</a> and/or its affiliates © 2023 – <a href="#\_\_consent">Cookie Consent</a>

# Table of contents

1. Percona Server for MySQL 8.1 – Documentation	5
1.1 For Monitoring and Management	5
1.2 Get expert help	5
2. Release notes	6
2.1 Percona Server for MySQL 8.1 release notes index	6
2.2 Percona Server for MySQL 8.1.0-1 (2023-11-27)	6
3. Features	10
3.1 Adaptive network buffers	10
3.2 Audit Log Filter	11
3.3 Limiting the disk space used by binary log files	45
3.4 Extended mysqlbinlog	46
3.5 Extended SELECT INTO OUTFILE/DUMPFILE	46
3.6 Expanded fast index creation	47
3.7 Kill idle transactions	49
3.8 The ProcFS plugin	49
3.9 Support for PROXY protocol	52
3.10 SEQUENCE_TABLE(n) function	53
3.11 Slow query log rotation and expiration	56
3.12 Thread pool	58
3.13 Trigger updates	63
3.14 Percona Toolkit UDFs	65
3.15 Utility user	66
4. Quickstart guide	72
4.1 Quickstart guide for Percona Server for MySQL	72
5. Install	74
5.1 Install Percona Server for MySQL	74
5.2 Before you start	74
5.3 Use APT	75
5.4 Use YUM	82
5.5 Use binary tarballs	87
5.6 Compile from source	89
5.7 Docker	91
5.8 Upgrade	97
6. Post-Installation	109
6.1 Binary logs and replication improvements	109

6.2	Post-installation	115
6.3	Working with AppArmor	119
6.4	Working with SELinux	127
7.	Manage	140
7.1	Extended SHOW GRANTS	140
7.2	Restrict dynamic log file locations	142
7.3	MySQL Clone plugin	142
7.4	Manage components	144
8.	Back up and restore	146
8.1	Backup and restore overview	146
8.2	Backup locks	148
8.3	Extended mysqldump	151
8.4	Start transaction with consistent snapshot	151
9.	Secure	154
9.1	LDAP features	154
9.2	Data masking	172
9.3	PAM authentication plugin	192
9.4	SSL improvements	194
9.5	FIDO authentication plugin	195
9.6	Encryption functions	197
9.7	Data-at-rest-encryption	208
10.	Replicate	229
10.1	Manage group replication flow control	229
10.2	Group replication system variables	229
10.3	Replication Security	232
11.	Monitor	234
11.1	Jemalloc memory allocation profiling	234
11.2	User statistics	238
11.3	Slow query log	247
11.4	Process list	256
11.5	Misc. INFORMATION_SCHEMA tables	257
12.	Troubleshoot	259
12.1	Use Percona Monitoring and Management (PMM) Advisors	259
12.2	Too many connections warning	259
12.3	Handle corrupted tables	260
12.4	Thread based profiling	261
12.5	Stack trace	261
12.6	Using libcoredumper	262

13. Storage engines	264
13.1 Percona XtraDB storage engine	264
13.2 Percona MyRocks	295
14. Reference	385
14.1 Percona Server for MySQL 8.1 variables	385
14.2 List of features available in Percona Server for MySQL releases	387
14.3 Percona Server for MySQL feature comparison	389
14.4 Understand version numbers	392
14.5 Development of Percona Server for MySQL	393
14.6 Telemetry on Percona Server for MySQL	394
14.7 Trademark policy	396
14.8 Index of INFORMATION_SCHEMA tables	397
14.9 Frequently asked questions	398
14.10 Copyright and licensing information	399
14.11 Glossary	399

# 1. Percona Server for MySQL 8.1 – Documentation

This documentation is for the latest release: Percona Server for MySQL 8.1.0-1 ([Release Notes](#)).

Percona Server for MySQL is a freely available, fully compatible, enhanced, and open source drop-in replacement for any MySQL database. It provides superior and optimized performance, greater scalability, and availability, enhanced backups, increased visibility and instrumentation.

Percona Server for MySQL is trusted by thousands of enterprises to provide better performance and concurrency for their most demanding workloads.

## 1.1 For Monitoring and Management

Percona Monitoring and Management (PMM ) monitors and provides actionable performance data for MySQL variants, including Percona Server for MySQL, Percona XtraDB Cluster, Oracle MySQL Community Edition, Oracle MySQL Enterprise Edition, and MariaDB. PMM captures metrics and data for the InnoDB, XtraDB, and MyRocks storage engines, and has specialized dashboards for specific engine details.

[Install PMM and connect your MySQL instances to it.](#)

## 1.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-11-23

## 2. Release notes

### 2.1 Percona Server for MySQL 8.1 release notes index

- [Percona Server for MySQL 8.1.0-1 \(2023-11-27\)](#)

#### 2.1.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

### 2.2 Percona Server for MySQL 8.1.0-1 (2023-11-27)

Get started with [Quickstart Guide for Percona Server for MySQL](#).

[Percona Server for MySQL 8.1.0-1](#) includes all the features and bug fixes available in the MySQL 8.1 Community Edition in addition to enterprise-grade features developed by Percona.

Percona Server for MySQL is a freely available, fully compatible, enhanced, and open source drop-in replacement for any MySQL database. It provides superior and optimized performance, greater scalability, and availability, enhanced backups, increased visibility and instrumentation.

Percona Server for MySQL is trusted by thousands of enterprises to provide better performance and concurrency for their most demanding workloads.

#### 2.2.1 Release highlights

Percona Server for MySQL implements telemetry that fills in the gaps in our understanding of how you use Percona Server to improve our products. Participation in the anonymous program is optional. You can opt-out if you prefer not to share this information. Find more information in the [Telemetry on Percona Server for MySQL](#) document.

The following user-defined function (UDF) shared objects (.so) are converted to components:

- The `data_masking` plugin converted into the `component_masking_functions` component
- The `binlogs_utils_udf` UDF shared object (.so) converted to the `component_binlog_utils` component
- The `percona-udf` UDF shared object (.so) converted to the `component_percona-udf` component

A user does not need to execute a separate `CREATE FUNCTION ... SONAME ...` statement for each function. Installing the components with the `INSTALL COMPONENT 'file://component_XXX'` statement performs the auto-registration operations.

The `keyring_vault` plugin converted into the `component_keyring_vault` component. This conversion aligns the `keyring_vault` with the KMIP and KMS keyrings and supports “ALTER INSTANCE RELOAD KEYRING” to update the configuration automatically.

The `audit_log_filter` plugin converted to the `component_audit_log_filter` component. The following changes are also available:

- Adds the `mysql_event_tracking_parse` audit log event
- Reworked, optimized, and reorganized the audit event data members
- Data deduplication within the audit event data members

The current version of `percona-release` does not support the `setup` subcommand with the `ps-8x-innovation` repository. Use `percona-release enable` instead. The support of the `ps-8x-innovation` repository for the `setup` subcommand will be added in the next release of `percona-release`.

The PS 8.1.0 MTR suites are reorganized. The existing percona-specific MTR test cases are regrouped and put into separate test suites:

- `component_encryption_udf`
- `percona`
- `percona_innodb`

Improvements and bug fixes introduced by Oracle for MySQL 8.1 and included in Percona Server for MySQL are the following:

- The `EXPLAIN FORMAT=JSON` can output the data to a user variable.
- New messages written to the MySQL error log during shutdown:
  - Startup and shutdown log messages, including when the server was started with `--initialize`
  - Start and end of shutdown phases for plugins and components
  - Start-of-phase and end-of-phase messages for connection closing phases
  - The number and ID of threads still alive after being forcibly disconnected and potentially causing a wait

Find the full list of bug fixes and changes in the [MySQL 8.1 Release Notes](#).

## 2.2.2 New feature

- [PS-5945](#): Add the mysql server host name in the `audit_log` details when writing to syslog.
- [PS-8042](#): The Audit Log version 2 initial design and implementation.
- [PS-8848](#): Increase the verbosity of the connection control plugin to help with debugging.

## 2.2.3 Bug fixes

- [PS-8932](#): The documentation mentioned `ALL_0_DIRECT` parameter for Percona Server for MySQL 8.x but flag was removed.
- [PS-8949](#): The MyRocks configuration failed on a Mac with M1.

## 2.2.4 Deprecation or removal

- The `mysql_native_password` authentication plugin is deprecated and subject to removal in a future version.
- The TokuDB is removed. The following items are also removed:
  - Percona-TokuBackup submodule
  - PerconaFT submodule
  - TokuDB storage engine code
  - TokuDB MTR test suites
  - `plugin/tokudb-backup-plugin`
- The MyRocks ZenFS is removed. The following items are also removed:
  - `zenfs` submodule
  - `libzdb` submodule
  - RocksDB MTR changes are reverted
- Travis CI integration
- Supporting `readline` as a alternative to editline library is removed.
- The `audit_log` (audit version 1) plugin is removed
- The “include/ext” pre-C++17 compatibility headers are removed.
- The `keyring_vault` plugin is removed.
- The `data_masking` UDF shared object (.so) is removed.
- The `binlog_utils_udf` UDF shared object (.so) is removed.
- The `percona_udf` UDF shared object (.so) is removed.

## 2.2.5 Platform support

- Percona Server for MySQL 8.1.0-1 is not supported on Ubuntu 18.04.

## 2.2.6 Useful links

- Install [Percona Server for MySQL 8.1](#)
- The [Percona Server for MySQL GitHub repository]
- Download product binaries, packages, and tarballs at [Percona Product Downloads](#)
- [Contribute to the documentation](#)
- For [training](#), contact [Percona Training – Start learning now](#)

## 2.2.7 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 3. Features

### 3.1 Adaptive network buffers

To find the buffer size of the current connection, use the `network_buffer_length` status variable. Add `SHOW GLOBAL` to review the cumulative buffer sizes for all connections. This variable can help to estimate the maximum size of the network buffer's overhead.

Network buffers grow towards the `max_allowed_packet` size and do not shrink until the connection is terminated. For example, if the connections are selected at random from the pool, an occasional big query eventually increases the buffers of all connections. The combination of `max_allowed_packet` set to a value between 64MB to 128MB and the connection number between 256 to 1024 can create a large memory overhead.

Percona Server for MySQL implemented the `net_buffer_shrink_interval` variable to solve this issue. The default value is 0 (zero). If you set the value higher than 0, Percona Server records the network buffer's maximum use size for the number of seconds set by `net_buffer_shrink_interval`. When the next interval starts, the network buffer is set to the recorded size. This action removes spikes in the buffer size.

You can achieve similar results by disconnecting and reconnecting the TCP connections, but this solution is a heavier process. This process disconnects and reconnects connections with small buffers.

#### `net_buffer_shrink_interval`

Option	Description
Command-line:	<code>--net-buffer-shrink-interval=#</code>
Scope:	Global
Dynamic:	Yes
Data type:	integer
Default value:	0

The interval is measured in seconds. The default value is 0, which disables the functionality. The minimum value is 0, and the maximum value is 31536000.

#### 3.1.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 3.2 Audit Log Filter

### 3.2.1 Audit Log Filter overview

The Audit Log Filter component allows you to monitor, log, and block a connection or query actively executed on the selected server.

Enabling the component produces a log file that contains a record of server activity. The log file has information on connections and databases accessed by that connection.

The component uses the `mysql` system database to store filter and user account data. Set the `audit_log_filter.database` variable at server startup to select a different database.

The `AUDIT_ADMIN` privilege is required to enable users to manage the Audit Log Filter component.

#### Privileges

Define the privilege at runtime at the startup of the server. The associated Audit Log Filter privilege can be unavailable if the component is not enabled.

##### `AUDIT_ADMIN`

This privilege is defined by the server and enables the user to configure the component.

##### `AUDIT_ABORT_EXEMPT`

This privilege allows queries from a user account to always be executed. An `abort` item does not block them. This ability lets the user account regain access to a system if an audit is misconfigured. The query is logged due to the privilege. User accounts with the `SYSTEM_USER` privilege have the `AUDIT_ABORT_EXEMPT` privilege.

#### Audit Log Filter tables

The Audit Log Filter component uses `mysql` system database tables in the `InnoDB` storage engine. These tables store user account data and filter data. When you start the server, change the component's database with the `audit_log_filter.database` variable.

The `audit_log_filter` table stores the definitions of the filters and has the following column definitions:

Column name	Description
NAME	Name of the filter
FILTER	Definition of the filter linked to the no

The `audit_log_user` table stores account data and has the following column definitions:

Column name	Description
USER	The account name of the user
HOST	The account name of the host
FILTERNAME	The account filter name

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 3.2.2 Install the Audit Log Filter

The `plugin_dir` system variable defines the component library location. If needed, at server startup, set the `plugin_dir` variable.

In the `share` directory, locate the `audit_log_filter_linux_install.sql` script.

At the time you run the script, you can select the database used to store the JSON filter tables.

- If the component is loaded, the installation script takes the database name from the `audit_log_filter.database` variable
- If the component is not loaded, but passes the `-D db_name` to the `mysql` client when the installation script runs, uses the `db_name`.
- If the component is not loaded and the `-D` option is not provided, the installation script creates the required tables in the default database name `mysql`.

You can also designate a different database with the `audit_log_filter.database` system variable. The database name cannot be NULL or exceed 64 characters. If the database name is invalid, the audit log filter tables are not found.

To install the component, run the following command:

```
mysql> INSTALL COMPONENT 'file://component_audit_log_filter';
```

Find more information in the [INSTALL COMPONENT](#) document.

After the installation, you can use the `--audit_log_filter` option when restarting the server. To prevent the server from not running the plugin use `--audit_log_filter` with either the `FORCE` or the `FORCE_PLUS_PERMANENT` values.

To upgrade from `audit_log_filter` plugin in Percona Server 8.0 to `component_audit_log_filter` component in Percona Server 8.1, do the [manual upgrade](#).

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#) [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

### 3.2.3 Formats

#### Audit Log Filter file format overview

When an auditable event occurs, the component writes a record to the log file.

After the component starts, the first record lists the description of the server and the options at startup. After the first record, the auditable events are connections, disconnections, SQL statements executed, and so on. Statements within stored procedures or triggers are not logged, only the top-level statements.

If files are referenced by `LOAD_DATA`, the contents are not logged.

Set with the `audit_log_filter.format` system variable at startup. The available format types are the following;

Format Type	Command	Description
<a href="#">XML (new style)</a>	<code>audit_log_filter.format=NEW</code>	The default
<a href="#">XML (old style)</a>	<code>audit_log_filter.format=OLD</code>	The original
<a href="#">JSON</a>	<code>audit_log_filter.format=JSON</code>	Files written in JSON format

By default, the file contents in the new-style XML format are not compressed or encrypted.

Changing the `audit_log_filter.format`, you should also change the `audit_log_filter.file` name. For example, changing the `audit_log_filter.format` to JSON, change the `audit_log_filter.file` to `audit.json`. If you don't change the `audit_log_filter.file` name, then all audit log filter files have the same base name and you won't be able to easily find when the format changed.

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#) [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

#### Audit Log Filter format - XML (new style)

The filter writes the audit log filter file in XML. The XML file uses UTF-8.

The is the root element and this element contains elements. Each element contains specific information about an event that is audited.

For each new file, the Audit Log Filter component writes the XML declaration and the root element tag. The component writes the closing root element when closing the file. If the file is open, this closing element is not available.

```
<?xml version="1.0" encoding="utf-8"?>
<AUDIT>
  <AUDIT_RECORD>
    <NAME>Audit</NAME>
    <RECORD_ID>0_2023-03-29T11:11:43</RECORD_ID>
    <TIMESTAMP>2023-03-29T11:11:43</TIMESTAMP>
    <SERVER_ID>1</SERVER_ID>
  </AUDIT_RECORD>
  <AUDIT_RECORD>
    <NAME>Command Start</NAME>
    <RECORD_ID>1_2023-03-29T11:11:45</RECORD_ID>
    <TIMESTAMP>2023-03-29T11:11:45</TIMESTAMP>
    <STATUS>0</STATUS>
    <CONNECTION_ID>1</CONNECTION_ID>
    <COMMAND_CLASS>query</COMMAND_CLASS>
  </AUDIT_RECORD>
  <AUDIT_RECORD>
    <NAME>Query</NAME>
    <RECORD_ID>2_2023-03-29T11:11:45</RECORD_ID>
    <TIMESTAMP>2023-03-29T11:11:45</TIMESTAMP>
    <COMMAND_CLASS>create_table</COMMAND_CLASS>
    <CONNECTION_ID>11</CONNECTION_ID>
    <HOST>localhost</HOST>
    <IP></IP>
    <USER>root[root] @ localhost []</USER>
    <OS_LOGIN></OS_LOGIN>
    <SQLTEXT>CREATE TABLE t1 (c1 INT)</SQLTEXT>
    <STATUS>0</STATUS>
  </AUDIT_RECORD>
  <AUDIT_RECORD>
    <NAME>Query Start</NAME>
    <RECORD_ID>3_2023-03-29T11:11:45</RECORD_ID>
    <TIMESTAMP>2023-03-29T11:11:45</TIMESTAMP>
    <STATUS>0</STATUS>
    <CONNECTION_ID>11</CONNECTION_ID>
    <COMMAND_CLASS>create_table</COMMAND_CLASS>
    <SQLTEXT>CREATE TABLE t1 (c1 INT)</SQLTEXT>
  </AUDIT_RECORD>
  <AUDIT_RECORD>
    <NAME>Query</NAME>
    <RECORD_ID>4_2023-03-29T11:11:45</RECORD_ID>
    <TIMESTAMP>2023-03-29T11:11:45</TIMESTAMP>
    <COMMAND_CLASS>create_table</COMMAND_CLASS>
    <CONNECTION_ID>11</CONNECTION_ID>
    <HOST>localhost</HOST>
    <IP></IP>
    <USER>root[root] @ localhost []</USER>
    <OS_LOGIN></OS_LOGIN>
    <SQLTEXT>CREATE TABLE t1 (c1 INT)</SQLTEXT>
    <STATUS>0</STATUS>
  </AUDIT_RECORD>
</AUDIT>
```

```

<AUDIT_RECORD>
  <NAME>Command End</NAME>
  <RECORD_ID>5_2023-03-29T11:11:45</RECORD_ID>
  <TIMESTAMP>2023-03-29T11:11:45</TIMESTAMP>
  <STATUS>0</STATUS>
  <CONNECTION_ID>1</CONNECTION_ID>
  <COMMAND_CLASS>query</COMMAND_CLASS>
</AUDIT_RECORD>
</AUDIT>

```

The order of the attributes within an can vary. Certain attributes are in every element. Other attributes are optional and depend on the type of audit record.

The attributes in every element are the following:

Attribute Name	Description
<NAME>	The action that generated the audit record.
<RECORD_ID>	The <RECORD_ID> consists of a sequence number and a timestamp. The sequence number is initialized with the value 1 in the filter file.
<TIMESTAMP>	Displays the date and time when the audit record was generated.

The optional attributes are the following:

Attribute Name	Description
<COMMAND_CLASS>	Contains the type of performed action.
<CONNECTION_ID>	Contains the client connection identifier.
<CONNECTION_ATTRIBUTES>	Contains the client connection attributes in the form of a <NAME> <VALUE> pair.
<CONNECTION_TYPE>	Contains the type of connection security.
<DB>	Contains the database name.
<HOST>	Contains the client's hostname.
<IP>	Contains the client's IP address.
<MYSQL_VERSION>	Contains the MySQL server version.
<OS_LOGIN>	Contains the user name used during authentication. If the user is authenticated through a proxy, this component does not set a value or the user name. If authentication fails, this value is empty.
<OS_VERSION>	Contains the server's operating system version.
<PRIV_USER>	Contains the user name used by the proxy. The user name may be different than <USER>.
<PROXY_USER>	Contains the proxy user. If a proxy is used, this value is not empty.
<SERVER_ID>	Contains the server ID.
<SQLTEXT>	Contains the text of the SQL statement.
<STARTUP_OPTIONS>	

Attribute Name	Description
	Contains the startup options. These can be specified on the command line or files.
<STATUS>	Contains the status of a command. A non-zero value indicates an error.
<STATUS_CODE>	Contains the status of a command, with the error code that occurred (i).
<TABLE>	Contains the table name.
<USER>	Contains the user name from the client. If the user is not known, <PRIV_USER>.
<VERSION>	Contains the audit log filter format.

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

### Audit Log Filter format - XML (old style)

The old style XML format uses `<AUDIT>` tag as the root element and adds the `</AUDIT>` tag when the file closes. Each audited event is contained in an element.

The order of the attributes within an can vary. Certain attributes are in every element. Other attributes are optional and depend on the type of audit record.

```
<?xml version="1.0" encoding="utf-8"?>
<AUDIT>
  <AUDIT_RECORD
    NAME="Audit"
    RECORD_ID="0_2023-03-29T11:15:52"
    TIMESTAMP="2023-03-29T11:15:52"
    SERVER_ID="1"/>
  <AUDIT_RECORD
    NAME="Command Start"
    RECORD_ID="1_2023-03-29T11:15:53"
    TIMESTAMP="2023-03-29T11:15:53"
    STATUS="0"
    CONNECTION_ID="1"
    COMMAND_CLASS="query"/>
  <AUDIT_RECORD
    NAME="Query"
    RECORD_ID="2_2023-03-29T11:15:53"
    TIMESTAMP="2023-03-29T11:15:53"
    COMMAND_CLASS="create_table"
    CONNECTION_ID="11"
```

```

HOST="localhost"
IP=""
USER="root[root] @ localhost []"
OS_LOGIN=""
SQLTEXT="CREATE TABLE t1 (c1 INT)"
STATUS="0"/>
<AUDIT_RECORD
  NAME="Query Start"
  RECORD_ID="3_2023-03-29T11:15:53"
  TIMESTAMP="2023-03-29T11:15:53"
  STATUS="0"
  CONNECTION_ID="11"
  COMMAND_CLASS="create_table"
  SQLTEXT="CREATE TABLE t1 (c1 INT)"/>
<AUDIT_RECORD
  NAME="Query Status End"
  RECORD_ID="4_2023-03-29T11:15:53"
  TIMESTAMP="2023-03-29T11:15:53"
  STATUS="0"
  CONNECTION_ID="11"
  COMMAND_CLASS="create_table"
  SQLTEXT="CREATE TABLE t1 (c1 INT)"/>
<AUDIT_RECORD
  NAME="Query"
  RECORD_ID="5_2023-03-29T11:15:53"
  TIMESTAMP="2023-03-29T11:15:53"
  COMMAND_CLASS="create_table"
  CONNECTION_ID="11"
  HOST="localhost"
  IP=""
  USER="root[root] @ localhost []"
  OS_LOGIN=""
  SQLTEXT="CREATE TABLE t1 (c1 INT)"
  STATUS="0"/>
<AUDIT_RECORD
  NAME="Command End"
  RECORD_ID="6_2023-03-29T11:15:53"
  TIMESTAMP="2023-03-29T11:15:53"
  STATUS="0"
  CONNECTION_ID="1"
  COMMAND_CLASS="query"/>
</AUDIT>

```

The required attributes are the following:

Attribute Name	Description
NAME	The action that generated the audit
RECORD_ID	The RECORD_ID consists of a sequence number. The sequence number is initialized when the audit file is created.
TIMESTAMP	Displays the date and time when the audit record was generated.

The optional attributes are the following:

Attribute Name	Description
COMMAND_CLASS	Type of action performed
CONNECTION_ID	Client connection identifier
CONNECTION_TYPE	Connection security type
DB	Database name
HOST	Client's hostname
IP	Client's IP address
MYSQL_VERSION	Server version
OS_LOGIN	The user name used during an external authentication. If the user is authenticated through an LDAP component, this value is empty.
OS_VERSION	Server's operating system
PRIV_USER	The user name used by the server when the proxy user is different than USER.
PROXY_USER	The proxy user. If a proxy is not used, this value is empty.
SERVER_ID	Server Identifier
SQLTEXT	SQL statement text
STARTUP_OPTIONS	Server startup options, either command line or configuration file.
STATUS	Command's status - a 0 (zero) is a success, a non-zero is a failure.
STATUS_CODE	A 0 (zero) is a success, a non-zero is a failure.
TABLE	Table name
USER	Client's user name - this name may be different than the user name used by the server.
VERSION	Format of audit log filter

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

### Audit Log Filter format - JSON

The JSON format has one top-level JSON array, which contain JSON objects with key-value pairs. These objects represent an event in the audit. Some pairs are listed in every audit record. The audit record type determines if other key-value pairs are listed. The order of the pairs within an audit record is not guaranteed. The value description may be truncated.

Certain statistics, such as query time and size, are only available in the JSON format and help detect activity outliers when analyzed.

```
[
  {
    "timestamp": "2023-03-29 11:17:03",
    "id": 0,
    "class": "audit",
    "server_id": 1
  },
  {
    "timestamp": "2023-03-29 11:17:05",
    "id": 1,
    "class": "command",
    "event": "command_start",
    "connection_id": 1,
    "command_data": {
      "name": "command_start",
      "status": 0,
      "command": "query"
    }
  },
  {
    "timestamp": "2023-03-29 11:17:05",
    "id": 2,
    "class": "general",
    "event": "log",
    "connection_id": 11,
    "account": { "user": "root[root] @ localhost []", "host": "localhost" },
    "login": { "user": "root[root] @ localhost []", "os": "", "ip": "", "proxy": "" },
    "general_data": {
      "command": "Query",
      "sql_command": "create_table",
      "query": "CREATE TABLE t1 (c1 INT)",
      "status": 0
    }
  },
  {
    "timestamp": "2023-03-29 11:17:05",
    "id": 3,
    "class": "query",
    "event": "query_start",
    "connection_id": 11,
    "query_data": {
      "query": "CREATE TABLE t1 (c1 INT)",
      "status": 0,
      "sql_command": "create_table"
    }
  },
  {
    "timestamp": "2023-03-29 11:17:05",
    "id": 4,
    "class": "query",
    "event": "query_status_end",
    "connection_id": 11,
    "query_data": {
      "query": "CREATE TABLE t1 (c1 INT)",
      "status": 0,
      "sql_command": "create_table"
    }
  },
]
```

```

{
  "timestamp": "2023-03-29 11:17:05",
  "id": 5,
  "class": "general",
  "event": "status",
  "connection_id": 11,
  "account": { "user": "root[root] @ localhost []", "host": "localhost" },
  "login": { "user": "root[root] @ localhost []", "os": "", "ip": "", "proxy":
"" },
  "general_data": {
    "command": "Query",
    "sql_command": "create_table",
    "query": "CREATE TABLE t1 (c1 INT)",
    "status": 0}
},
{
  "timestamp": "2023-03-29 11:17:05",
  "id": 6,
  "class": "command",
  "event": "command_end",
  "connection_id": 1,
  "command_data": {
    "name": "command_end",
    "status": 0,
    "command": "query"}
}
]

```

The order of the attributes within the JSON object can vary. Certain attributes are in every element. Other attributes are optional and depend on the type of event and the filter settings or component settings.

The following fields are contained in each object:

- `timestamp`
- `id`
- `class`
- `event`

The possible attributes in a JSON object are the following:

Name	Description
<code>class</code>	Defines the type of event
<code>account</code>	Defines the MySQL account associated with the event.
<code>connection_data</code>	Defines the client connection.
<code>connection_id</code>	Defines the client connection identifier.
<code>event</code>	Defines a subclass of the <code>event</code> class.
<code>general_data</code>	Defines the executed statement or command. Only applicable for the <code>general</code> class value of <code>general</code> .
<code>id</code>	Defines the event ID
<code>login</code>	Defines how the client connected to the server.

Name	Description
<code>query_statistics</code>	Defines optional query statistics and
<code>shutdown_data</code>	Defines the audit log filter termination
<code>startup_data</code>	Defines the initialization of the audit l
<code>table_access_data</code>	Defines access to a table
<code>time</code>	Defines an integer that represents a
<code>timestamp</code>	Defines a UTC value in the YYYY-MM-DD

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

### 3.2.4 Audit Log Filter security

The Audit Log Filter component generates audit log filter files. The directory that contains these files should be accessible only to the following:

- Users who must be able to view the log
- Server must be able to write to the directory

The files are not encrypted by default and may contain sensitive information.

The default name for the file in the data directory is `audit_filter.log`. If needed, use the `audit_log_filter.file` system variable at server startup to change the location. Due to the log rotation, multiple audit log files may exist.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 3.2.5 Audit Log Filter compression and encryption

### Compression

You can enable compression for any `format` by setting the `audit_log_filter.compression` system variable when the server starts.

The `audit_log_filter.compression` variable can be either of the following:

- NONE (no compression) - the default value
- GZIP - uses the GNU Zip compression

If compression and encryption are enabled, the component applies compression before encryption. If you must manually recover a file with both settings, first decrypt the file and then uncompress the file.

### Encryption

You can encrypt any audit log filter file in any `format`. The audit log filter component generates the initial password, but you can use user-defined passwords after that. The component stores the passwords in the keyring, so that feature must be enabled.

Set the `audit_log_filter.encryption` system variable with the server starts. The allowed values are the following:

- NONE - no encryption, the default value
- AES - AES-256-CBC (Cipher Block Chaining) encryption

The AES uses the 256-bit key size.

The following audit log filter functions are used with encryption:

Function name	Description
<code>audit_log_encryption_password_set()</code>	Stores the password in the keyring. If rotates the log file by renaming the o encrypted with the password.
<code>audit_log_encryption_password_get()</code>	Invoking this function without an arg password. An argument that specifie or current password returns that pas

The `audit_log_filter.password_history_keep_days` variable is used with encryption. If the variable is not zero (0), invoking `audit_log_encryption_password_set()` causes the expiration of archived audit log passwords.

When the component starts with encryption enabled, the component checks if the keyring has an audit log filter encryption password. If no password is found, the component generates a random password and stores this password in the keyring. Use `audit_log_encryption_password_get()` to review this password.

If compression and encryption are enabled, the component applies compression before encryption. If you must manually recover a file with both settings, first decrypt the file and then uncompress the file.

## Manually uncompressing and decrypting audit log filter files

To decrypt an encrypted log file, use the openssl command. For example:

```
openssl enc -d -aes-256-cbc -pass pass:password
  -iter iterations -md sha256
  -in audit.timestamp.log.pwd_id.enc
  -out audit.timestamp.log
```

To execute that command, you must obtain a password and iterations. To do this, use `audit_log_encryption_password_get()`.

This function gets the encryption password, and the iterations count and returns this data as a JSON-encoded string. For example, if the audit log file name is `audit.20190415T151322.log.20190414T223342-2.enc`, the password ID is `20190414T223342-2` and the keyring ID is `audit-log-20190414T223342-2`.

Get the keyring password:

```
mysql> SELECT audit_log_encryption_password_get('audit-log-20190414T223342-2');
```

The return value of this function may look like the following:

### Expected output

```
{"password":"6i1Nx35CKMIaQaFzBrDzLpCq6gC","iterations":568977}
```

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 3.2.6 Reading Audit Log Filter files

The Audit Log Filter functions can provide a SQL interface to read JSON-format audit log files. The functions cannot read log files in other formats. Configuring the component for JSON logging lets the functions use the directory that contains the current audit log filter file and search in that location for readable files. The value of the `audit_log_filter.file` system variable provides the file location, base name, and the suffix and then searches for names that match the pattern.

If the file is renamed and no longer fits the pattern, the file is ignored.

## Functions used for reading the files

The following functions read the files in the JSON-format:

- `audit_log_read` - reads audit log filter events
- `[audit_log_read_bookmark()](audit-log-filter-variables.md#audit_log_read_bookmark)` - for the most recently read event, returns a bookmark. The bookmark can be passed to `audit_log_read()`.

Initialize a read sequence by using a bookmark or an argument that specifies the start position:

```
mysql> SELECT audit_log_read(audit_log_read_bookmark());
```

The following example continues reading from the current position:

```
```{.bash data-prompt="mysql>"}
mysql> SELECT audit_log_read();
```

Reading a file is closed when the session ends or calling `audit_log_read()` with another argument.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 3.2.7 Manage the Audit Log Filter files

The Audit Log Filter files have the following potential results:

- Consume a large amount of disk space
- Grow large

You can manage the space by using log file rotation. This operation renames and then rotates the current log file and then uses the original name on a new current log file. You can rotate the file either manually or automatically.

If automatic rotation is enabled, you can prune the log file. This pruning operation can be based on either the log file age or combined log file size.

### Manual log rotation

The default setting for `audit_log_filter.rotate_on_size` is 1GB. If this option is set to `0`, the audit log filter component does not do an automatic rotation of the log file. You must do the rotation manually with this setting.

The `SELECT audit_log_rotate()` command renames the file and creates a new audit log filter file with the original name. You must have the `AUDIT_ADMIN` privilege.

The files are pruned if either `audit_log_filter.max_size` or `audit_log_filter.prune_seconds` have a value greater than 0 (zero) and `audit_log_filter.rotate_on_size` > 0.

After the files have been renamed, you must manually remove any archived audit log filter files. The renamed audit log filter files can be read by `audit_log_read()`. The `audit_log_read()` does not find the logs if the name pattern differs from the current pattern.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#) [:percona–logo: Get a Percona Expert](#)

Last update: 2023-11-27

### 3.2.8 Filter the Audit Log Filter logs

The audit filter log filtering is based on rules. The filter rule definition has the ability to include or exclude events based on the following attributes:

- User account
- Audit event class
- Audit event subclass
- Audit event fields (for example, `COMMAND_CLASS` or `STATUS`)

You can define multiple filters and assign any filter to multiple accounts. You can also create a default filter for specific user accounts. The filters are defined using function calls. After the filter is defined, the filter is stored in `mysql` system tables.

#### Audit Log Filter functions

The Audit Log filter functions require `AUDIT_ADMIN` or `SUPER` privilege.

The following functions are used for rule-based filtering:

Function	Description	Example
<code>audit_log_filter_flush()</code>	Manually flush the filter tables	<code>SELECT</code>
<code>audit_log_filter_set_filter()</code>	Defines a filter	<code>SELECT</code> <code>audit_log</code> <code>{}}')</code>
<code>audit_log_filter_remove_filter()</code>	Removes a filter	
<code>audit_log_filter_set_user()</code>	Assigns a filter to a specific user account	
<code>audit_log_filter_remove_user()</code>	Removes the filters from a specific user account	

Using a SQL interface, you can define, display, or modify audit log filters. The filters are stored in the `mysql` system database.

The `audit_log_session_filter_id()` function returns the internal ID of the audit log filter in the current session.

Filter definitions are `JSON` values.

The function, `audit_log_filter_flush()`, forces reloading all filters and should only be invoked when modifying the audit tables. This function affects all users. Users in current sessions must either execute `change-user` or `disconnect` and `reconnect`.

### Constraints

The `component_audit_log_filter` component must be enabled and the audit tables must exist to use the audit log filter functions. The user account must have the required privileges.

### Using the audit log filter functions

With a new connection, the audit log filter component finds the user account name in the filter assignments. If a filter has been assigned, the component uses that filter. If no filter has been assigned, but there is a default account filter, the component uses that filter. If there is no filter assigned, and there is no default account filter, then the component does not process any event.

The default account is represented by `%` as the account name.

You can assign filters to a specific user account or disassociate a user account from a filter. To disassociate a user account, either unassign a filter or assign a different filter. If you remove a filter, that filter is unassigned from all users, including current users in current sessions.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 3.2.9 Audit Log Filter restrictions

### General restrictions

The Audit Log Filter has the following general restrictions:

- Logs only SQL statements. Statements made by NoSQL APIs, such as the Memcached API, are not logged.
- Logs only the top-level statement. Statements within a stored procedure or a trigger are not logged. Does not log the file contents for statements like `LOAD_DATA`.
- If used with a cluster, the component must be installed on each server used to execute SQL on the cluster.
- If used with a cluster, the application or user is responsible for aggregating all the data of each server used in the cluster.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 3.2.10 Audit Log Filter file naming conventions

### Name qualities

The audit log filter file name has the following qualities:

- Optional directory name
- Base name
- Optional suffix

Using either [compression](#) or [encryption](#) adds the following suffixes:

- Compression adds the `.gz` suffix
- Encryption adds the `pwd_id.enc` suffix

The `pwd_id` represents the password used for encrypting the log files. The audit log filter component stores passwords in the keyring.

You can combine compression and encryption, which adds both suffixes to the `audit_filter.log` name.

The following table displays the possible ways a file can be named:

Default name	Enabled feature
audit.log	No compression or encryption
audit.log.gz	Compression
audit.log.pwd_id.enc	Encryption
audit.log.gz.pwd_id.enc	Compression, encryption

#### ENCRYPTION ID FORMAT

The format for `pwd_id` is the following:

- A UTC value in `YYYYMMDDThhmmss` format that represents when the password was created
- A sequence number that starts at `1` and increases if passwords have the same timestamp value

The following are examples of `pwd_id` values:

```
20230417T082215-1
20230301T061400-1
20230301T061400-2
```

The following example is a list of the audit log filter files with the `pwd_id`:

```
audit_filter.log.20230417T082215-1.enc
audit_filter.log.20230301T061400-1.enc
audit_filter.log.20230301T061400-2.enc
```

The current password has the largest sequence number.

#### Renaming operations

During initialization, the component checks if a file with that name exists. If it does, the component renames the file. The component writes to an empty file.

During termination, the component renames the file.

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

### 3.2.11 Disable Audit Log Filter logging

The `audit_log_filter.disable` system variable lets you disable or enable logging for all connections.

You can set the variable in the following ways:

- Option file
- Command-line startup string
- SET statement during runtime

```
mysql> SET GLOBAL audit_log_filter.disable = true;
```

Setting `audit_log_filter.disable` has the following effect:

Value	Actions
true	Generates a warning. Audit log functions generate session warnings. Disables the component.
false	Re-enables the component and generates a warning.

### Privileges required

Setting the value of `audit_log_filter.disable` at runtime requires the following:

- `AUDIT_ADMIN` privilege
- `SYSTEM_VARIABLES_ADMIN` privilege

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 3.2.12 Audit log filter functions, options and variables

The following sections describe the [functions](#), [options](#), and [variables](#) available in the audit log filter component.

### Audit log filter functions

The following audit log filter functions are available.

Function name
<code>audit_log_encryption_password_get(keyring_id)</code>
<code>audit_log_encryption_password_set(new_password)</code>
<code>audit_log_filter_flush()</code>

Function name
<code>audit_log_read()</code>
<code>audit_log_read_bookmark()</code>
<code>audit_log_session_filter_id()</code>
<code>audit_log_filter_remove_filter(filter_name)</code>
<code>audit_log_filter_remove_user(user_name)</code>
<code>audit_log_rotate()</code>
<code>audit_log_filter_set_filter(filter_name, definition)</code>
<code>audit_log_filter_set_user(user_name, filter_name)</code>

#### AUDIT\_LOG\_ENCRYPTION\_PASSWORD\_GET(KEYRING\_ID)

This function returns the encryption password. Any keyring component or keyring component can be used, but the component or component must be enabled. If the component or component is not enabled, an error occurs.

#### Parameters

`keyring_id` - If the function does not contain a `keyring_id`, the function returns the current encryption password. You can also request a specific encryption password with the keyring ID of either the current password or an archived password.

#### Returns

This function returns a JSON object containing the password, iterations count used by the password.

#### Example

```
mysql> SELECT audit_log_encryption_password_get();
```

#### Expected output

```
+-----+
| audit_log_encryption_password_get() |
+-----+
| {"password": "passw0rd", "iterations": 5689} |
+-----+
```

#### AUDIT\_LOG\_ENCRYPTION\_PASSWORD\_SET(NEW\_PASSWORD)

This function sets the encryption password and stores the new password in the keyring.

#### Parameters

`password` - the password as a string. The maximum length is 766 bytes.

#### Returns

This function returns a string. An `OK` indicates a success. `ERROR` indicates a failure.

#### Example

```
mysql> SELECT audit_log_encryption_password_set(password);
```

#### Expected output

```
+-----+
| audit_log_encryption_password_set(password) |
+-----+
| OK  |
+-----+
```

#### AUDIT\_LOG\_FILTER\_FLUSH()

This function updates the audit log filter tables and makes any changes operational.

Modifying the audit log filter tables directly with `INSERT`, `UPDATE`, or `DELETE` does not implement the modifications immediately. The tables must be flushed to have those changes take effect.

This function forces reloading all filters and should only be used if someone has modified the tables directly.

#### Important

Avoid using this function. This function performs an operation that is similar to uninstalling and reinstalling the component. Filters are detached from all current sessions. To restart logging, the current sessions must either disconnect and reconnect or do a change-user operation.

#### Parameters

None.

#### Returns

This function returns either an `OK` for success or an error message for failure.

#### Example

```
mysql> SELECT audit_log_filter_flush();
```

#### Expected output

```
+-----+
| audit_log_filter_flush() |
+-----+
| OK                          |
+-----+
```

#### AUDIT\_LOG\_READ()

If the audit log filter format is JSON, this function reads the audit log and returns an array of the audit events as a JSON string. Generates an error if the format is not JSON.

**Parameters**

None. If the start position is not provided, the read continues from the current position.

Optional: You can specify a starting position for the read with `start` or a `timestamp` and an `id`, both items are considered a bookmark and can be used to identify an event. You must include both (`timestamp` and `id`) or an error is generated. If the `timestamp` does not include a `time` section, the function assumes the time is `00:00`.

You can also provide a `max_array_length` to limit the number of log events.

Call `audit_log_read_bookmark()` to return the most recently written event.

**Returns**

This function returns a string of a JSON array of the audit events, or a JSON NULL value. Returns `NULL` and generates an error if the call fails.

**Example**

```
mysql> SELECT audit_log_read(audit_log_read_bookmark());
```

**Expected output**

```
+-----+
| audit_log_read(audit_log_read_bookmark()) |
+-----+
| [{"timestamp" : "2023-06-02 09:43:25", "id": 10, "class": "connection",}] |
+-----+
```

**AUDIT\_LOG\_READ\_BOOKMARK()**

This function provides a bookmark for the most recently written audit log event as a JSON string. Generates an error if the format is not JSON.

If this function is used with `[audit_log_read()](#audit_log_read)`, the `audit_log_read()` function starts reading at that position.

**Parameters**

None.

**Returns**

This function returns a `JSON` string containing a bookmark for success or `NULL` and an error for failure.

**Example**

```
mysql> SELECT audit_log_read_bookmark();
```

**Expected output**

```
+-----+
| audit_log_read_bookmark() |
+-----+
| {"timestamp" : "2023-06-02 09:43:25", "id": 10 } |
+-----+
```

**AUDIT\_LOG\_SESSION\_FILTER\_ID()**

This function returns the internal ID of the audit log filter in the current session.

Returns 0 (zero) if the session has no assigned filter.

**AUDIT\_LOG\_FILTER\_REMOVE\_FILTER(FILTER\_NAME)**

This function removes the selected filter from the current set of filters.

If user accounts are assigned the selected filter, the user accounts are no longer filtered. The user accounts are removed from `audit_log_user`. If the user accounts are in a current session, they are detached from the selected filter and no longer logged.

**Parameters**

`filter_name` - a selected filter name as a string.

**Returns**

This function returns either an `OK` for success or an error message for failure.

If the filter name does not exist, no error is generated.

**Example**

```
mysql> SELECT audit_log_filter_remove_filter('filter-name');
```

**Expected output**

```
+-----+
| audit_log_filter_remove_filter('filter-name') |
+-----+
| OK |
+-----+
```

**AUDIT\_LOG\_FILTER\_REMOVE\_USER(USER\_NAME)**

This function removes the assignment of a filter from the selected user account.

If the user account is in a current session, they are not affected. New sessions for this user account use the default account filter or are not logged.

If the user-name is `%`, the default account filter is removed.

**Parameters**

`user_name` - a selected user name in either the `user_name @ host_name` format or `%`.

**Returns**

This function returns either an `OK` for success or an error message for failure.

If the `user_name` has no filter assigned, no error is generated.

**Example**

```
mysql> SELECT audit_log_filter_remove_user('user-name@localhost');
```

**Expected output**

```
+-----+
| audit_log_filter_remove_user('user-name@localhost') |
+-----+
| OK |
+-----+
```

**AUDIT\_LOG\_ROTATE()****Parameters**

None.

**Returns**

This function returns renamed file name.

**Example**

```
mysql> SELECT audit_log_filter_rotate();
```

**AUDIT\_LOG\_FILTER\_SET\_FILTER(FILTER\_NAME, DEFINITION)**

This function, when provided with a filter name and definition, adds the filter.

The new filter has a different filter ID. Generates an error if the filter name exists.

**Parameters**

- `filter_name` - a selected filter name as a string.
- `definition` - Defines the definition as a JSON value.

**Returns**

This function returns either an `OK` for success or an error message for failure.

**Example**

```
mysql> SET @filter = '{ "filter_name": { "log": true } }'
mysql> SET audit_log_filter_set_filter('filter-name', @filter);
```

**Expected output**

```
+-----+
| audit_log_filter_set_filter('filter-name', @filter) |
+-----+
| OK |
+-----+
```

**AUDIT\_LOG\_FILTER\_SET\_USER(USER\_NAME, FILTER\_NAME)**

This function assigns the filter to the selected user account.

A user account can only have one filter. If the user account already has a filter, this function replaces the current filter. If the user account is in a current session, nothing happens. When the user account connects again the new filter is used.

The user name, %, is the default account. The filter assigned to % is used by any user account without a defined filter.

**Parameters**

- `user_name` - a selected user name in either the `user_name @ host_name` format or %.
- `filter_name` - a selected filter name as a string.

**Returns**

This function returns either an `OK` for success or an error message for failure.

**Example**

```
mysql> SELECT audit_log_filter_set_user('user-name@localhost', 'filter-name');
```

**Expected output**

```
+-----+
| audit_log_filter_set_user('user-name@localhost', 'filter-name') |
+-----+
| OK |
+-----+
```

**Audit log filter options and variables**

Name
<code>audit-log-filter</code>
<code>audit_log_filter.buffer_size</code>
<code>audit_log_filter.compression</code>
<code>audit_log_filter.database</code>
<code>audit_log_filter.disable</code>
<code>audit_log_filter.encryption</code>

Name
<code>audit_log_filter.file</code>
<code>audit_log_filter.format</code>
<code>audit_log_filter.format_unix_timestamp</code>
<code>audit_log_filter.handler</code>
<code>audit_log_filter.key_derivation_iterations_count_mean</code>
<code>audit_log_filter.max_size</code>
<code>audit_log_filter.keep_password_history_keep_days</code>
<code>audit_log_filter.prune_seconds</code>
<code>audit_log_filter.read_buffer_size</code>
<code>audit_log_filter.rotate_on_size</code>
<code>audit_log_filter.strategy</code>
<code>audit_log_filter.syslog_tag</code>
<code>audit_log_filter.syslog_priority</code>

**AUDIT-LOG-FILTER**

Option	Description
Command-line	<code>-audit-log-filter[=value]</code>
Dynamic	No
Scope	
Data type	Enumeration
Default	ON

This option determines how, at startup, the server loads the `audit_log_filter` component. The component must be registered.

The valid values are the following:

- ON
- OFF
- FORCE
- FORCE\_PLUS\_PERMANENT

**AUDIT\_LOG\_FILTER.BUFFER\_SIZE**

Option name	Description
Command-line	<code>-audit-log-filter.buffer-size</code>
Dynamic	No
Scope	Global
Data type	Integer

Option name	Description
Default	1048576
Minimum value	4096
Maximum value	18446744073709547520
Units	bytes
Block size	4096

This variable defines the buffer size in multiples of 4096 when logging is asynchronous. The contents for events are stored in a buffer. The contents are stored until the contents are written.

The component initializes a single buffer and removes the buffer when the component terminates.

#### AUDIT\_LOG\_FILTER.COMPRESSION

Option name	Description
Command-line	-audit-log-filter.compression
Dynamic	Yes
Scope	Global
Data type	Enumeration
Default	NONE
Valid values	NONE or GZIP

This variable defines the compression type for the audit log filter file. The values can be either `NONE`, the default value and file has no compression, or `GZIP`.

#### AUDIT\_LOG\_FILTER.DATABASE

Option name	Description
Command-line	-audit-log-filter.database
Dynamic	No
Scope	Global
Data type	String
Default	mysql

This variable defines the `audit_log_filter` database. This read-only variable stores the necessary tables. Set this option at system startup. The database name cannot exceed 64 characters or be `NULL`.

An invalid database name prevents the use of audit log filter component.

#### AUDIT\_LOG\_FILTER.DISABLE

Option name	Description
Command-line	-audit-log-filter.disable

Option name	Description
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

This variable disables the component logging for all connections and any sessions.

This variable requires the user account to have `SYSTEM_VARIABLES_ADMIN` and `AUDIT_ADMIN` privileges.

#### AUDIT\_LOG\_FILTER. ENCRYPTION

Option name	Description
Command-line	<code>-audit-log-filter.encryption</code>
Dynamic	No
Scope	Global
Data type	Enumeration
Default	NONE
Valid values	NONE or AES

This variable defines the encryption type for the audit log filter file. The values can be either of the following:

- NONE - the default value, no encryption
- AES

#### AUDIT\_LOG\_FILTER.FILE

Option name	Description
Command-line	<code>-audit-log-filter.file</code>
Dynamic	No
Scope	Global
Data type	String
Default	<code>audit_filter.log</code>

This variable defines the name and suffix of the audit log filter file. The component writes events to this file.

The file name and suffix can be either of the following:

- a relative path name - the component looks for this file in the data directory
- a full path name - the component uses the given value

If you use a full path name, ensure the directory is accessible only to users who need to view the log and the server.

For more information, see [Naming conventions](#)

#### AUDIT\_LOG\_FILTER.FORMAT

Option name	Description
Command-line	<code>-audit-log-filter.format</code>
Dynamic	No
Scope	Global
Data type	Enumeration
Default	NEW
Available values	OLD, NEW, JSON

This variable defines the audit log filter file format.

The available values are the following:

- [OLD \(old-style XML\)](#)
- [NEW \(new-style XML\)](#) and
- [JSON](#).

#### AUDIT\_LOG\_FILTER.FORMAT\_UNIX\_TIMESTAMP

Option name	Description
Command-line	<code>-audit-log-filter.format-unix-timestamp</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

This option is only supported for JSON-format files.

Enabling this option adds a `time` field to JSON-format files. The integer represents the UNIX timestamp value and indicates the date and time when the audit event was generated. Changing the value causes a file rotation because all records must either have or do not have the `time` field. This option requires the `AUDIT_ADMIN` and `SYSTEM_VARIABLES_ADMIN` privileges.

This option does nothing when used with other format types.

#### AUDIT\_LOG\_FILTER.HANDLER

Option name	Description
Command-line	<code>-audit-log-filter.handler</code>
Dynamic	No
Scope	Global
Data type	String
Default	FILE

Defines where the component writes the audit log filter file. The following values are available:

- FILE - component writes the log to a location specified in `audit_log_filter.file`
- SYSLOG - component writes to the syslog

#### AUDIT\_LOG\_FILTER.KEY\_DERIVATION\_ITERATIONS\_COUNT\_MEAN

Option name	Description
Command-line	<code>-audit-log-filter.key-derivation-itera</code>
Dynamic	Yes
Scope	Global
Data type	Integer
Default	60000
Minimum value	1000
Maximum value	1000000

Defines the mean value of iterations used by the password-based derivation routine while calculating the encryption key and iv values. A random number represents the actual iteration count and deviates no more than 10% from this value.

#### AUDIT\_LOG\_FILTER.MAX\_SIZE

Option name	Description
Command-line	<code>-audit-log-filter.max-size</code>
Dynamic	Yes
Scope	Global
Data type	Integer
Default	1GB
Minimum value	0
Maximum value	18446744073709551615
Unit	bytes
Block size	4096

Defines pruning based on the combined size of the files:

The default value is 1GB.

A value of 0 (zero) disables pruning based on size.

A value greater than 0 (zero) enables pruning based on size and defines the combined size limit. When the files exceed this limit, they can be pruned.

The value is based on 4096 (block size). A value is truncated to the nearest multiple of the block size. If the value is less than 4096, the value is treated as 0 (zero).

If the values for `audit_log_filter.rotate_on_size` and `audit_log_filter.max_size` are greater than 0, we recommend that `audit_log_filter.max_size` value should be at least seven times the `audit_log_filter.rotate_on_size` value.

Pruning requires the following options:

- `audit_log_filter.max_size`
- `audit_log_filter.rotate_on_size`
- `audit_log_filter.prune_seconds`

#### AUDIT\_LOG\_FILTER.PASSWORD\_HISTORY\_KEEP\_DAYS

Option name	Description
Command-line	<code>-audit-log-filter.password-history-k</code>
Dynamic	Yes
Scope	Global
Data type	Integer
Default	0

Defines when passwords may be removed and measured in days.

Encrypted log files have passwords stored in the keyring. The component also stores a password history. A password does not expire, despite being past the value, in case the password is used for rotated audit logs. The operation of creating a password also archives the previous password.

The default value is 0 (zero). This value disables the expiration of passwords. Passwords are retained forever.

If the component starts and encryption is enabled, the component checks for an audit log filter encryption password. If a password is not found, the component generates a random password.

Call `audit_log_filter_encryption_set()` to set a specific password.

#### AUDIT\_LOG\_FILTER.PRUNE\_SECONDS

Option name	Description
Command-line	<code>-audit-log-filter.prune-seconds</code>
Dynamic	Yes
Scope	Global
Data type	Integer
Default	0
Minimum value	0
Maximum value	1844674073709551615
Unit	seconds

Defines when the audit log filter file is pruned. This pruning is based on the age of the file. The value is measured in seconds.

A value of 0 (zero) is the default and disables pruning. The maximum value is 18446744073709551615.

A value greater than 0 enables pruning. An audit log filter file can be pruned after this value.

To enable log pruning, you must set one of the following:

- Enable log rotation by setting `audit_log_filter.rotate_on_size`
- Add a value greater than 0 (zero) for either `audit_log_filter.max_size` or `audit_log_filter.prune_seconds`

#### AUDIT\_LOG\_FILTER.READ\_BUFFER\_SIZE

Option name	Description
Command-line	<code>-audit-log-filter.read-buffer-size</code>
Dynamic	Yes
Scope	Global
Data type	Integer
Unit	Bytes
Default	32768

This option is only supported for JSON-format files.

The size of the buffer for reading from the audit log filter file. The `audit_log_filter_read()` reads only from this buffer size.

#### AUDIT\_LOG\_FILTER.ROTATE\_ON\_SIZE

Option name	Description
Command-line	<code>-audit-log-filter.rotate-on-size</code>
Dynamic	Yes
Scope	Global
Data type	Integer
Default	1GB

Performs an automatic log file rotation based on the size. The default value is 1GB. If the value is greater than 0, when the log file size exceeds the value, the component renames the current file and opens a new log file using the original name.

If you set the value to less than 4096, the component does not automatically rotate the log files. You can rotate the log files manually using `audit_log_rotate()`. If the value is not a multiple of 4096, the component truncates the value to the nearest multiple.

#### AUDIT\_LOG\_FILTER.STRATEGY

Option name	Description
Command-line	<code>-audit-log-filter.strategy</code>
Dynamic	No

Option name	Description
Scope	Global
Data type	Enumeration
Default	ASYNCHRONOUS

Defines the Audit Log filter component's logging method. The valid values are the following:

Values	Description
ASYNCHRONOUS	Waits until there is outer buffer space
PERFORMANCE	If the outer buffer does not have enough space
SEMISYNCHRONOUS	Operating system permits caching
SYNCHRONOUS	Each request calls <code>sync()</code>

#### AUDIT\_LOG\_FILTER.SYSLOG\_TAG

Option	Description
Command-line	<code>-audit-log-filter.syslog-tag=</code>
Dynamic	No
Scope	Global
Data type	String
Default	<code>audit-filter</code>

#### AUDIT\_LOG\_FILTER.SYSLOG\_FACILITY

Option name	Description
Command-line	<code>-audit-log-filter.syslog-facility</code>
Dynamic	No
Scope	Global
Data type	String
Default	<code>LOG_USER</code>

Specifies the syslog `facility` value. The option has the same meaning as the appropriate parameter described in the [syslog\(3\) manual](#).

#### AUDIT\_LOG\_FILTER.SYSLOG\_PRIORITY

Option name	Description
Command-line	<code>-audit-log-filter.syslog-priority</code>
Dynamic	No
Scope	Global
Data type	String

Option name	Description
Default	LOG_INFO

Defines the `priority` value for the syslog. The option has the same meaning as the appropriate parameter described in the [syslog\(3\) manual](#).

### Audit log filter status variables

The audit log filter component exposes status variables. These variables provide information on the operations.

Name	Description
<code>audit_log_filter_current_size</code>	The current size of the audit log filter to 0.
<code>audit_log_filter_direct_writes</code>	Identifies when the <code>log_strategy_type</code> bypass the write buffer and are written.
<code>audit_log_filter_max_drop_size</code>	In the performance logging mode, the
<code>audit_log_filter_events</code>	The number of audit log filter events
<code>audit_log_filter_events_filtered</code>	The number of filtered audit log filter
<code>audit_log_filter_events_lost</code>	If the event is larger than the available space, the event is lost
<code>audit_log_filter_events_written</code>	The number of audit log filter events
<code>audit_log_filter_total_size</code>	The total size of the events written to the log. This value increases even when a log is rotated.
<code>audit_log_filter_write_waits</code>	In the asynchronous logging mode, the number of times the space in the audit log filter buffer

#### Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 3.2.13 Uninstall Audit Log Filter

To remove the component, run the following:

```
mysql> UNINSTALL COMPONENT 'file://component_audit_log_filter';
```

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 3.3 Limiting the disk space used by binary log files

It is a challenge to control how much disk space is used by the binary logs. The size of a binary log can vary because a single transaction must be written to a single binary log and cannot be split between multiple binary log files.

### 3.3.1 binlog\_space\_limit

Attribute	Description
Uses the command line	Yes
Uses the configuration file	Yes
Scope	Global
Dynamic	No
Variable type	ULONG_MAX
Default value	0 (unlimited)
Maximum value - 64-bit platform	18446744073709547520

This variable places an upper limit on the total size in bytes of all binary logs. When the limit is reached, the oldest binary logs are purged until the total size is under the limit or only the active log remains.

The default value of `0` disables the feature. No limit is set on the log space. The binary logs accumulate indefinitely until the disk space is full.

### 3.3.2 Example

Set the `binlog_space_limit` to 50 GB in the `my.cnf` file:

```
[mysqld]
...
binlog_space_limit = 50G
...
```

### 3.3.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-07-27

## 3.4 Extended mysqlbinlog

### Note

The `--compress` option was marked as deprecated and may be removed in a future version of Percona Server for MySQL.

Percona Server for MySQL has implemented compression support for `mysqlbinlog`. This is similar to support that both `mysql` and `mysqldump` programs include (the `-C`, `--compress` options “Use compression in server/client protocol”). Using the compressed protocol helps reduce the bandwidth use and speed up transfers.

Percona Server for MySQL has also implemented support for `SSL`. `mysqlbinlog` now accepts the `SSL` connection options as all the other client programs. This feature can be useful with `--read-from-remote-server` option.

### 3.4.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-31

## 3.5 Extended SELECT INTO OUTFILE/DUMPFILE

Percona Server for MySQL extends the `SELECT INTO ... OUTFILE` and `SELECT INTO DUMPFILE` commands to add support for UNIX sockets and named pipes. Before this was implemented the database would return an error for such files.

This feature allows using `LOAD DATA LOCAL INFILE` in combination with `SELECT INTO OUTFILE` to quickly load multiple partitions across the network or in other setups, without having to use an intermediate file that wastes space and I/O.

## 3.5.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 3.6 Expanded fast index creation

Percona has implemented several changes related to *MySQL*'s fast index creation feature. Fast index creation was implemented in *MySQL* as a way to speed up the process of adding or dropping indexes on tables with many rows.

This feature implements a session variable that enables extended fast index creation. Besides optimizing DDL directly, [expand\\_fast\\_index\\_creation](#) may also optimize index access for subsequent DML statements because using it results in much less fragmented indexes.

### 3.6.1 The mysqldump command

A new option, `--innodb-optimize-keys`, was implemented in **mysqldump**. It changes the way *InnoDB* tables are dumped, so that secondary and foreign keys are created after loading the data, thus taking advantage of fast index creation. More specifically:

- `KEY`, `UNIQUE KEY`, and `CONSTRAINT` clauses are omitted from `CREATE TABLE` statements corresponding to *InnoDB* tables.
- An additional `ALTER TABLE` is issued after dumping the data, in order to create the previously omitted keys.

### 3.6.2 ALTER TABLE

When `ALTER TABLE` requires a table copy, secondary keys are now dropped and recreated later, after copying the data. The following restrictions apply:

- Only non-unique keys can be involved in this optimization.
- If the table contains foreign keys, or a foreign key is being added as a part of the current `ALTER TABLE` statement, the optimization is disabled for all keys.

### 3.6.3 OPTIMIZE TABLE

Internally, `OPTIMIZE TABLE` is mapped to `ALTER TABLE ... ENGINE=innodb` for *InnoDB* tables. As a consequence, it now also benefits from fast index creation, with the same restrictions as for `ALTER TABLE`.

### 3.6.4 Caveats

*InnoDB* fast index creation uses temporary files in `tmpdir` for all indexes being created. So make sure you have enough `tmpdir` space when using `expand_fast_index_creation`. It is a session variable, so you can temporarily switch it off if you are short on `tmpdir` space and/or don't want this optimization to be used for a specific table.

There's also a number of cases when this optimization is not applicable:

- `UNIQUE` indexes in `ALTER TABLE` are ignored to enforce uniqueness where necessary when copying the data to a temporary table;
- `ALTER TABLE` and `OPTIMIZE TABLE` always process tables containing foreign keys as if `expand_fast_index_creation` is `OFF` to avoid dropping keys that are part of a `FOREIGN KEY` constraint;
- `mysqldump --innodb-optimize-keys` ignores foreign keys because *InnoDB* requires a full table rebuild on foreign key changes. So adding them back with a separate `ALTER TABLE` after restoring the data from a dump would actually make the restore slower;
- `mysqldump --innodb-optimize-keys` ignores indexes on `AUTO_INCREMENT` columns, because they must be indexed, so it is impossible to temporarily drop the corresponding index;
- `mysqldump --innodb-optimize-keys` ignores the first `UNIQUE` index on non-nullable columns when the table has no `PRIMARY KEY` defined, because in this case *InnoDB* picks such an index as the clustered one.

### 3.6.5 System variables

#### `expand_fast_index_creation`

Option	Description
Command Line:	Yes
Config file	No
Scope:	Local/Global
Dynamic:	Yes
Data type	Boolean
Default value	ON/OFF

#### See also

[Improved InnoDB fast index creation](#)

[Thinking about running OPTIMIZE on your InnoDB Table? Stop!](#)

### 3.6.6 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-31

## 3.7 Kill idle transactions

This feature limits the age of idle transactions, for all transactional storage engines. If a transaction is idle for more seconds than the threshold specified, it will be killed. This prevents users from blocking *InnoDB* purge by mistake.

### 3.7.1 System variables

#### `kill_idle_transaction`

Option	Description
Config file	Yes
Scope:	Global
Dynamic:	Yes
Data type	Integer
Default value	0 (disabled)
Units	Seconds

If non-zero, any idle transaction will be killed after being idle for this many seconds.

### 3.7.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 3.8 The ProcFS plugin

The ProcFS plugin provides access to the Linux performance counters by running SQL queries against a Percona Server for MySQL 8.1.

You may be unable to capture operating system metrics in certain environments, such as Cloud installations or MySQL-as-a-Service installations. These metrics are essential for complete system performance monitoring.

The plugin does the following:

- Reads selected files from the `/proc` file system and the `/sys` file system.
- Populates the file names and their content as rows in the `INFORMATION_SCHEMA.PROCFS` view.

The system variable `procfs_files_spec` provides access to the `/proc` and the `/sys` files and directories. This variable cannot be changed at run time, preventing a compromised account from giving itself greater access to those file systems.

### 3.8.1 Install the PLUGIN manually

We recommend installing the plugin as part of the package. If needed, you can install this plugin manually. Copy the `procfs.so` file to the mysql plugin installation directory and execute the following command:

```
INSTALL PLUGIN procfs SONAME 'procfs.so';
```

### 3.8.2 Access privileges required

Only users with the `ACCESS_PROCFS` dynamic privilege can access the `INFORMATION_SCHEMA.PROCFS` view. During the plugin startup, this dynamic privilege is registered with the server.

After the plugin installation, grant a user access to the `INFORMATION_SCHEMA.PROCFS` view by executing the following command:

```
GRANT ACCESS_PROCFS ON *.* TO 'user'@'host';
```

#### Important

An SELinux policy or an AppArmor profile may prevent access to file locations needed by the ProcFS plugin, such as the `/proc/sys/fs/file-nr` directory or any sub-directories or files under `/proc/iq/`. Either edit the policy or profile to ensure that the plugin has the necessary access. If the policy and profile do not allow access, the plugin may have unexpected behavior.

For more information, see [Working with SELinux](#) and [Working with AppArmor](#).

### 3.8.3 Using the ProcFS plugin

Authorized users can obtain information from individual files by specifying the exact file name within a `WHERE` clause. Files that are not included are ignored and considered not to exist.

All files that match the `procfs_files_spec` are opened, read, stored in memory, and, finally, returned to the client. It is critical to add a `WHERE` clause to return only specific files to limit the impact of the plugin on the server's performance. A failure to use a `WHERE` clause can lead to lengthy query response times, high load, and high memory usage on the server. The `WHERE` clause can contain either an equality operator, the `LIKE` operator, or the `IN` operator. The `LIKE` operator limits file globbing. You can write file access patterns in the [glob\(7\) style](#), such as `/sys/block/sd[a-z]/stat;/proc/version\*`

The following example returns the `proc/version`:

```
SELECT * FROM INFORMATION_SCHEMA.PROCFS WHERE FILE = '/proc/version';
```

## 3.8.4 Tables

### PROCFS

The schema definition of the INFORMATION\_SCHEMA.PROCFS view is:

```
CREATE TEMPORARY TABLE `PROCFS` (
  `FILE` varchar(1024) NOT NULL DEFAULT '',
  `CONTENTS` longtext NOT NULL
) ENGINE=InnoDB DEFAULT CHARSET=utf8;
```

Status variables provide the basic metrics:

Name	Description
procfs_access_violations	The number of attempted queries by privilege.
procfs_queries	The number of queries made against
procfs_files_read	The number of files read to provide c
procfs_bytes_read	The number of bytes read to provide

## 3.8.5 Variable

### procfs\_files\_spec

Option	Description
Scope:	Global
Dynamic:	Yes
Read, Write, or Read-Only:	Read-Only

The default value for `procfs_files_spec` is: `/proc/cpuinfo;/proc/irq//;/proc/loadavg;/proc/net/dev;/proc/net/sockstat;/proc/net/sockstat_rhe4;/proc/net/tcpstat;/proc/self/net/netstat;/proc/self/stat;/proc/self/io;/proc/self/numa_maps;/proc/softirqs;/proc/spl/kstat/zfs/arcstats;/proc/stat;/proc/sys/fs/file-nr;/proc/version;/proc/vmstat`

Enables access to the `/proc` and `/sys` directories and files. This variable is global, read only, and is set by using either the `mysqld` command line or by editing `my.cnf`.

### 3.8.6 Limitations

The following limitations are:

- Only first 60k of `/proc/` `/sys/` files are returned
- The file name size is limited to 1k
- The plugin cannot read files if path does not start from `/proc` or `/sys`
- Complex WHERE conditions may force the plugin to read all configured files.

### 3.8.7 Uninstall plugin

The following statement removes the `procfs` plugin.

```
UNINSTALL PLUGIN procfs;
```

### 3.8.8 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-31

## 3.9 Support for PROXY protocol

The proxy protocol allows an intermediate proxying server speaking proxy protocol (ie. HAProxy) between the server and the ultimate client (i.e. mysql client etc) to provide the source client address to the server, which normally would only see the proxying server address instead.

As the proxy protocol amounts to spoofing the client address, it is disabled by default, and can be enabled on per-host or per-network basis for the trusted source addresses where trusted proxy servers are known to run. Unproxied connections are not allowed from these source addresses.

#### Note

Ensure that proper firewall access control lists (ACL) are in place when this feature is enabled.

Proxying is supported only for TCP over IPv4 and IPv6 connections. The UNIX socket connections can not be proxied and do not fall under the effect of using the asterisk symbol (\*).

You cannot have a proxied IP address that is `127.0.0.1` or `::1`, even if the IP address is in the `proxy_protocol_networks`.

## 3.9.1 System variables

### proxy\_protocol\_networks

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	No
Default	(empty string)

This variable is a global-only, read-only variable, which is either an asterisk symbol(\*), or a list of comma-separated IPv4 and IPv6 network and host addresses. For security reasons we do not recommend using an asterisk symbol for the IP address. This symbol causes the server to accept the proxy protocol from any host. Network addresses are specified in CIDR notation, i.e. 192.168.0.0/24. To prevent source host spoofing, the setting of this variable must be as restrictive as possible to include only trusted proxy hosts.

## 3.9.2 Related reading

- [PROXY protocol specification](#)

## 3.9.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-09-25

## 3.10 SEQUENCE\_TABLE(n) function

Percona Server for MySQL supports the SEQUENCE\_TABLE() function.

A sequence of numbers can be defined as an arithmetic progression when the common difference between two consecutive terms is always the same.

The function is an inline table-valued function. A single SELECT statement generates a multi-row result set. In contrast, a scalar function (like EXP(x) or LOWER(str)) always returns a single value of a specific data type.

The JSON\_TABLE() is the only table function available in Oracle MySQL Server. JSON\_TABLE and SEQUENCE\_TABLE() are the only table functions available in Percona Server.

The basic syntax is the following:

- `SEQUENCE_TABLE(n) [AS] alias`

### 3.10.1 Usage

#### Expected output

```
SELECT ... FROM SEQUENCE_TABLE(n) [AS] alias
SEQUENCE_TABLE(n) [AS] alias
```

`n`: The number of generated values.

As with any [derived tables](#), a table function requires an [alias](#) in the `SELECT` statement.

The result set is a single column with the predefined column name `value` of type `BIGINT UNSIGNED`. You can reference the `value` column in `SELECT` statements.

The following statements are valid:

```
mysql> SELECT * FROM SEQUENCE_TABLE(n) AS tt;
mysql> SELECT <expr(value)> FROM SEQUENCE_TABLE(n) AS tt;
```

The first number in the series, the initial term, is defined as `0` and the series ends with a value less than `n`. In this example, enter the following statement to generate a sequence:

```
mysql> SELECT * FROM SEQUENCE_TABLE(3) AS tt;
```

#### Expected output

```
+-----+
| value |
+-----+
|     0 |
|     1 |
|     2 |
+-----+
```

You can define the initial term using the `WHERE` clause. The following example starts the sequence with `4`.

```
SELECT value AS result FROM SEQUENCE_TABLE(8) AS tt WHERE value >= 4;
```

**Expected output**

```
+-----+
| result |
+-----+
|      4 |
|      5 |
|      6 |
|      7 |
+-----+
```

Consecutive terms increase or decrease by a common difference. The default common difference value is 1. However, it is possible to filter the results using the WHERE clause to simulate common differences greater than 1.

The following example prints only even numbers from the 0..7 range:

```
SELECT value AS result FROM SEQUENCE_TABLE(8) AS tt WHERE value % 2 = 0;
```

**Expected output**

```
+-----+
| result |
+-----+
|      0 |
|      2 |
|      4 |
|      6 |
+-----+
```

The following is an example of using the function to populate a table with a set of random numbers:

```
mysql> SELECT FLOOR(RAND() * 100) AS result FROM SEQUENCE_TABLE(4) AS tt;
```

The output could be the following:

**Expected output**

```
+-----+
| result |
+-----+
|      24 |
|      56 |
|      70 |
|      25 |
+-----+
```

You can populate a table with a set of pseudo-random strings with the following statement:

```
mysql> SELECT MD5(value) AS result FROM SEQUENCE_TABLE(4) AS tt;
```

### Expected output

```
+-----+
| result |
+-----+
| f17d9c990f40f8ac215f2ecdfd7d0451 |
| 2e5751b7cfd7f053cd29e946fb2649a4 |
| b026324c6904b2a9cb4b88d6d61c81d1 |
| 26ab0db90d72e28ad0ba1e22ee510510 |
+-----+
```

You can add the sequence as a column to a new table or an existing table, as shown in this example:

```
mysql> CREATE TABLE t1 AS SELECT * FROM SEQUENCE_TABLE(4) AS tt;

mysql> SELECT * FROM t1;
```

### Expected output

```
+-----+
| value |
+-----+
|      0 |
|      1 |
|      2 |
|      3 |
+-----+
```

There are many uses for a sequence when populating tables.

## 3.10.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 3.11 Slow query log rotation and expiration

Percona has implemented two new variables, `max_slowlog_size` and `max_slowlog_files` to provide users with ability to control the slow query log disk usage. These variables have the same behavior as the `max_binlog_size` variable and the `max_binlog_files` variable used for controlling the binary log.

### 3.11.1 max\_slowlog\_size

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Data type	numeric
Default	0 (unlimited)
Range	0 - 1073741824

The server rotates the slow query log when the log's size reaches this value. The default value is 0. If you limit the size and this feature is enabled, the server renames the slow query log file to `slow_query_log_file.000001`.

### 3.11.2 max\_slowlog\_files

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Data type	numeric
Default	0 (unlimited)
Range	0 - 102400

This variable limits the total amount of slow query log files and is used with `max_slowlog_size`.

The server creates and adds slow query logs until reaching the range's upper value. When the upper value is reached, the server creates a new slow query log file with a higher sequence number and deletes the log file with the lowest sequence number maintaining the total amount defined in the range.

### 3.11.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-31

## 3.12 Thread pool

MySQL executes statements using one thread per client connection. Once the number of connections increases past a certain point performance will degrade.

This feature enables the server to keep the top performance even with a large number of client connections by introducing a dynamic thread pool. By using the thread pool server would decrease the number of threads, which will then reduce the context switching and hot locks contentions. Using the thread pool will have the most effect with OLTP workloads (relatively short CPU-bound queries).

In order to enable the thread pool variable `thread_handling` should be set up to `pool-of-threads` value. This can be done by adding:

```
thread_handling=pool-of-threads
```

Although the default values for the thread pool should provide good performance, additional tuning can be performed with the dynamic system variables.

### Note

Current implementation of the thread pool is built in the server, unlike the upstream version which is implemented as a plugin. Another significant implementation difference is that this implementation doesn't try to minimize the number of concurrent transactions like the MySQL Enterprise Threadpool. Because of these differences, this implementation is not compatible with the upstream version.

### 3.12.1 Priority connection scheduling

Even though thread pool puts a limit on the number of concurrently running queries, the number of open transactions may remain high, because connections with already started transactions are put to the end of the queue. Higher number of open transactions has a number of implications on the currently running queries. To improve the performance new `thread_pool_high_prio_tickets` variable has been introduced.

This variable controls the high priority queue policy. Each new connection is assigned this many tickets to enter the high priority queue. Whenever a query has to be queued to be executed later because no threads are available, the thread pool puts the connection into the high priority queue if the following conditions apply:

- The connection has an open transaction in the server.
- The number of high priority tickets of this connection is non-zero.

If both the above conditions hold, the connection is put into the high priority queue and its tickets value is decremented. Otherwise the connection is put into the common queue with the initial tickets value specified with this option.

Each time the thread pool looks for a new connection to process, first it checks the high priority queue, and picks connections from the common queue only when the high priority one is empty.

The goal is to minimize the number of open transactions in the server. In many cases it is beneficial to give short-running transactions a chance to commit faster and thus deallocate server resources

and locks without waiting in the same queue with other connections that are about to start a new transaction, or those that have run out of their high priority tickets.

The default thread pool behavior is to always put events from already started transactions into the high priority queue, as we believe that results in better performance in vast majority of cases.

With the value of 0, all connections are always put into the common queue, i.e. no priority scheduling is used as in the original implementation in *MariaDB*. The higher is the value, the more chances each transaction gets to enter the high priority queue and commit before it is put in the common queue.

In some cases it is required to prioritize all statements for a specific connection regardless of whether they are executed as a part of a multi-statement transaction or in the autocommit mode. Or vice versa, some connections may require using the low priority queue for all statements unconditionally. To implement this new `thread_pool_high_prio_mode` variable has been introduced in *Percona Server for MySQL*.

### Low priority queue throttling

One case that can limit thread pool performance and even lead to deadlocks under high concurrency is a situation when thread groups are oversubscribed due to active threads reaching the oversubscribe limit, but all/most worker threads are actually waiting on locks currently held by a transaction from another connection that is not currently in the thread pool.

What happens in this case is that those threads in the pool that have marked themselves inactive are not accounted to the oversubscribe limit. As a result, the number of threads (both active and waiting) in the pool grows until it hits `thread_pool_max_threads` value. If the connection executing the transaction which is holding the lock has managed to enter the thread pool by then, we get a large (depending on the `thread_pool_max_threads` value) number of concurrently running threads, and thus, suboptimal performance as a result. Otherwise, we get a deadlock as no more threads can be created to process those transaction(s) and release the lock(s).

Such situations are prevented by throttling the low priority queue when the total number of worker threads (both active and waiting ones) reaches the oversubscribe limit. That is, if there are too many worker threads, do not start new transactions and create new threads until queued events from the already started transactions are processed.

## 3.12.2 Handling of long network waits

Certain types of workloads (large result sets, BLOBs, slow clients) can have longer waits on network I/O (socket reads and writes). Whenever server waits, this should be communicated to the Thread Pool, so it can start new query by either waking a waiting thread or sometimes creating a new one. This implementation has been ported from *MariaDB* patch MDEV-156.

## 3.12.3 System variables

### `thread_pool_idle_timeout`

Option	Description
Command-line:	Yes
Config file:	Yes

Option	Description
Scope:	Global
Dynamic:	Yes
Data type:	Numeric
Default value:	60 (seconds)

This variable can be used to limit the time an idle thread should wait before exiting.

#### `thread_pool_high_prio_mode`

Option	Description
Command-line:	Yes
Config file:	Yes
Scope:	Global, Session
Dynamic:	Yes
Data type:	String
Default value:	<code>transactions</code>
Allowed values:	<code>transactions, statements, none</code>

This variable is used to provide more fine-grained control over high priority scheduling either globally or per connection.

The following values are allowed:

- `transactions` (the default). In this mode only statements from already started transactions may go into the high priority queue depending on the number of high priority tickets currently available in a connection (see `thread_pool_high_prio_tickets`).
- `statements`. In this mode all individual statements go into the high priority queue, regardless of connection's transactional state and the number of available high priority tickets. This value can be used to prioritize `AUTOCOMMIT` transactions or other kinds of statements such as administrative ones for specific connections. Note that setting this value globally essentially disables high priority scheduling, since in this case all statements from all connections will use a single queue (the high priority one)
- `none`. This mode disables high priority queue for a connection. Some connections (e.g. monitoring) may be insensitive to execution latency and/or never allocate any server resources that would otherwise impact performance in other connections and thus, do not really require high priority scheduling. Note that setting `thread_pool_high_prio_mode` to `none` globally has essentially the same effect as setting it to `statements` globally: all connections will always use a single queue (the low priority one in this case).

#### `thread_pool_high_prio_tickets`

Option	Description
Command-line:	Yes

Option	Description
Config file:	Yes
Scope:	Global, Session
Dynamic:	Yes
Data type:	Numeric
Default value:	4294967295

This variable controls the high priority queue policy. Each new connection is assigned this many tickets to enter the high priority queue. Setting this variable to 0 will disable the high priority queue.

#### `thread_pool_max_threads`

Option	Description
Command-line:	Yes
Config file:	Yes
Scope:	Global
Dynamic:	Yes
Data type:	Numeric
Default value:	100000

This variable can be used to limit the maximum number of threads in the pool. Once this number is reached no new threads will be created.

#### `thread_pool_oversubscribe`

Option	Description
Command-line:	Yes
Config file:	Yes
Scope:	Global
Dynamic:	Yes
Data type:	Numeric
Default value:	3

The higher the value of this parameter the more threads can be run at the same time, if the values is lower than 3 it could lead to more sleeps and wake-ups.

#### `thread_pool_size`

Option	Description
Command-line:	Yes

Option	Description
Config file:	Yes
Scope:	Global
Dynamic:	Yes
Data type:	Numeric
Default value:	Number of processors

This variable can be used to define the number of threads that can use the CPU at the same time.

#### `thread_pool_stall_limit`

Option	Description
Command-line:	Yes
Config file:	Yes
Scope:	Global
Dynamic:	No
Data type:	Numeric
Default value:	500 (ms)

The number of milliseconds before a running thread is considered stalled. When this limit is reached thread pool will wake up or create another thread. This is being used to prevent a long-running query from monopolizing the pool.

### 3.12.4 Status variables

#### `Threadpool_idle_threads`

Option	Description
Scope:	Global
Data type:	Numeric

This status variable shows the number of idle threads in the pool.

#### `Threadpool_threads`

Option	Description
Scope:	Global
Data type:	Numeric

This status variable shows the number of threads in the pool.

### 3.12.5 Other reading

- [Thread pool in MariaDB 5.5](#)
- [Thread pool implementation in Oracle MySQL](#)

### 3.12.6 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-23

## 3.13 Trigger updates

Clients can issue simultaneous queries for a table. To avoid scalability problems, each thread-handling query has its own table instance. The server uses a special cache, called the Table Cache, which contains open table instances. The use of the cache avoids paying the penalty in resources for opening and closing tables for each statement.

The [ `table_open_cache` ] system variable sets soft limits on the cache size. This limit can be temporarily exceeded if the currently executing queries require more open tables than specified. However, when these queries complete, the server closes the unused table instances from this cache using the least recently used (LRU) algorithm.

The [ `table_open_cache_instances` ] system variable shows the number of open tables cache instances.

For more information, see [How MySQL opens and closes tables](#).

Opening a table with triggers in Table Cache also parses the trigger definitions and associates the open table instance with its own instances of the defined trigger bodies. When a connection executes a DML statement and must run a trigger, that connection gets its own instance of the trigger body for that specific open table instance. As a result of this approach, caching open table instances and also caching an associated trigger body for each trigger can consume a surprising amount of memory.

### 3.13.1 Version specific information

Percona Server for MySQL has the following abilities:

- Avoid using table instances with fully-loaded and parsed triggers by read-only queries
- Show trigger CREATE statements even if the statement is unparseable

The additional system variable reduces the Table Cache memory consumption on the server when tables that contain trigger definitions also are part of a significant read-only workload.

### 3.13.2 System variables

#### table\_open\_cache\_triggers

Option	Description
Command-line	<code>--table-open-cache-triggers</code>
Dynamic	Yes
Scope	Global
Data type	Integer
Default	524288
Minimum value	1
Maximum value	524288

This variable allows you to set a soft limit on the maximum of open tables in the Table Cache, which contains fully-loaded triggers. By default, the value is the maximum value to avoid existing users observing a change in behavior.

If the number of open table instances with fully-loaded triggers exceeds the value, then unused table instances with fully-loaded triggers are removed. This operation uses the least recently used (LRU) method for managing storage areas.

The value can be a start-up option or changed dynamically.

### 3.13.3 Status variables

The following status variables are available:

Variable name	Description
<code>table_open_cache_triggers_hits</code>	A hit means the statement required <code>table_open_cache_triggers</code> triggers and was able to get one from the cache.
<code>table_open_cache_triggers_misses</code>	A miss means the statement requiring <code>table_open_cache_triggers</code> loaded triggers was not found one in the cache. You may find a table instance without fully loading for it.
<code>table_open_cache_triggers_overflows</code>	An overflow indicates the number of <code>table_open_cache_triggers</code> were expelled from the <code>table_open_cache_triggers</code> soft limit. <code>table_open_cache_triggers</code> value shows the number of overflows.

### 3.13.4 SHOW CREATE TRIGGER statement changes

The `SHOW CREATE TRIGGER` statement shows the `CREATE` statement used to create the trigger. The statement also shows definitions which can no longer be parsed. For example, you can show the definition of a trigger created before a server upgrade which changed the trigger syntax.

### 3.13.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 3.14 Percona Toolkit UDFs

Three Percona Toolkit UDFs that provide faster checksums are provided:

- `libfnv1a_udf`
- `libfnv_udf`
- `libmurmur_udf`

### 3.14.1 Other information

- Author/Origin: Baron Schwartz

### 3.14.2 Installation

These UDFs are part of the *Percona Server for MySQL* packages. To install one of the UDFs into the server, execute one of the following commands, depending on which UDF you want to install:

```
mysql -e "CREATE FUNCTION fnv1a_64 RETURNS INTEGER SONAME 'libfnv1a_udf.so'"
mysql -e "CREATE FUNCTION fnv_64 RETURNS INTEGER SONAME 'libfnv_udf.so'"
mysql -e "CREATE FUNCTION murmur_hash RETURNS INTEGER SONAME 'libmurmur_udf.so'"
```

Executing each of these commands will install its respective UDF into the server.

### 3.14.3 Troubleshooting

If you get the error:

 **Error message**

```
ERROR 1126 (HY000): Can't open shared library 'fnv_udf.so' (errno: 22 fnv_udf.so: cannot open shared object file: No such file or directory)
```

Then you may need to copy the `.so` file to another location in your system. Try both `/lib` and `/usr/lib`. Look at your environment's `$LD_LIBRARY_PATH` variable for clues. If none is set, and neither `/lib` nor `/usr/lib` works, you may need to set `LD_LIBRARY_PATH` to `/lib` or `/usr/lib`.

### 3.14.4 Other reading

- Percona Toolkit [documentation](#)

### 3.14.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 3.15 Utility user

*Percona Server for MySQL* has implemented ability to have a *MySQL* user who has system access to do administrative tasks but limited access to user schema. This feature is especially useful to those operating *MySQL As A Service*.

This user has a mixed and special scope of abilities and protection:

- Utility user does not appear in the `mysql.user` table and can not be modified by any other user, including `root`.
- Utility user does not appear in `INFORMATION_SCHEMA.USER_STATISTICS`, `INFORMATION_SCHEMA.CLIENT_STATISTICS` or `THREAD_STATISTICS` tables or in any `performance_schema` tables.
- Utility user's queries may appear in the general and slow logs.
- Utility user does not have the ability create, modify, delete or see any schemas or data not specified, except for `information_schema`.
- Utility user may modify all visible, non-read-only system variables (see *expanded\_option\_modifiers* functionality).
- Utility user may see, create, modify and delete other system users only if given access to the `mysql` schema.
- Regular users may be granted proxy rights to the utility user but attempts to impersonate the utility user fail. The utility user may not be granted proxy rights on any regular user.

For example, `GRANT PROXY ON utility_user TO regular_user;` does not fail, but any actual attempt to impersonate as the utility user fails.

`GRANT PROXY ON regular_user TO utility_user;` fails when `utility_user` is an exact match or is more specific than the utility user specified.

At server start, the server notes in the log output that the utility user exists and the schemas that the utility user can access.

### 3.15.1 System variables

In order to have the ability for a special type of MySQL user, which will have a very limited and special amount of control over the system and can not be see or modified by any other user including the root user, three new options have been added.

#### utility\_user

Option	Description
Command Line:	Yes
Config file	utility_user=<user@host>
Scope:	Global
Dynamic:	No
Data type	String
Default	NULL

Specifies a MySQL user that will be added to the internal list of users and recognized as the utility user.

Option `utility_user` specifies the user which the system creates and recognizes as the utility user. The host in the utility user specification follows conventions described in the [MySQL manual](#). For example, the conventions allow wildcards and IP masks. Anonymous user names are not permitted to be used for the utility user name.

This user must not be an exact match to any other user that exists in the `mysql.user` table. If the server detects that the user specified with this option exactly matches any user within the `mysql.user` table on start up, the server reports an error and exits gracefully.

If host name wildcards are used and a more specific user specification is identified on start up, the server reports a warning and continues.

 **Error message**

```
utility_user=frank@% and [frank@localhost](mailto:frank@localhost) exists within the mysql.user table.
```

If a client attempts to create a MySQL user that matches this user specification exactly or if host name wildcards are used for the utility user and the user being created has the same name and a more specific host, the creation attempt fails with an error.

 **Error message**

```
utility_user=frank@% and CREATE USER ['frank@localhost](mailto:'frank@localhost');
```

As a result of these requirements, it is strongly recommended that a very unique user name and reasonably specific host be used.

Verify the script or tools test they are running within the correct user by executing `SELECT CURRENT_USER()` and comparing the result against the known utility user.

#### `utility_user_password`

Option	Description
Command Line:	Yes
Config file	<code>utility_user_password=password</code>
Scope:	Global
Dynamic:	No
Data type	String
Default	NULL

Specifies the password required for the utility user.

Option `utility_user_password` specifies the password for the utility user and must be specified or the server exits with an error.

Utility user password
<code>utility_user_password=Passw0rD</code>

#### `utility_user_schema_access`

Option	Description
Command Line:	Yes
Config file	<code>utility_user_schema_access=schem</code>
Scope:	Global
Dynamic:	No
Data type	String
Default	NULL

Specifies the schemas that the utility user has access to in a comma delimited list.

Option `utility_user_schema_access` specifies the name(s) of the schema(s) that the utility user will have access to read write and modify. If a particular schema named here does not exist on start up it will be ignored. If a schema by the name of any of those listed in this option is created after the server is started, the utility user will have full access to it.

Utility user schema access
<code>utility_user_schema_access=schema1,schema2,schema3</code>

**utility\_user\_privileges**

Option	Description
Command Line:	Yes
Config file	utility_user_privileges=privilege1,priv
Scope:	Global
Dynamic:	No
Data type	String
Default	NULL

This variable can be used to specify a comma-separated list of extra access privileges to grant to the utility user. Supported values for the privileges list are: SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, SHUTDOWN, PROCESS, FILE, GRANT, REFERENCES, INDEX, ALTER, SHOW DATABASES, SUPER, CREATE TEMPORARY TABLES, LOCK TABLES, EXECUTE, REPLICATION SLAVE, REPLICATION CLIENT, CREATE VIEW, SHOW VIEW, CREATE ROUTINE, ALTER ROUTINE, CREATE USER, EVENT, TRIGGER, CREATE TABLESPACE

Option [utility-user-privileges](#) allows a comma-separated list of extra access privileges to grant to the utility user.

 **Utility user privileges**

```
utility-user-privileges = "CREATE,DROP,LOCK TABLES";
```

**utility\_user\_dynamic\_privileges**

Option	Description
Command Line:	Yes
Config file	utility_user_dynamic_privileges=priv
Scope:	Global
Dynamic:	No
Data type	String
Default	NULL

This variable allows a comma-separated list of extra access dynamic privileges to grant to the utility user. The supported values for the dynamic privileges are:

- APPLICATION\_PASSWORD\_ADMIN
- AUDIT\_ADMIN
- BACKUP\_ADMIN
- BINLOG\_ADMIN
- BINLOG\_ENCRYPTION\_ADMIN
- CLONE\_ADMIN
- CONNECTION\_ADMIN
- ENCRYPTION\_KEY\_ADMIN
- FIREWALL\_ADMIN
- FIREWALL\_USER
- GROUP\_REPLICATION\_ADMIN
- INNODB\_REDO\_LOG\_ARCHIVE
- NDB\_STORED\_USER
- PERSIST\_RO\_VARIABLES\_ADMIN
- REPLICATION\_APPLIER
- REPLICATION\_SLAVE\_ADMIN
- RESOURCE\_GROUP\_ADMIN
- RESOURCE\_GROUP\_USER
- ROLE\_ADMIN
- SESSION\_VARIABLES\_ADMIN
- SET\_USER\_ID
- SHOW\_ROUTINE
- SYSTEM\_USER
- SYSTEM\_VARIABLES\_ADMIN
- TABLE\_ENCRYPTION\_ADMIN
- VERSION\_TOKEN\_ADMIN
- XA\_RECOVER\_ADMIN

Other dynamic privileges may be defined by plugins.

Option `utility_user_dynamic_privileges` allows a comma-separated list of extra-access dynamic privileges to grant to the utility user.

#### Utility user dynamic privileges

```
utility_user_dynamic_privileges = "SYSTEM_USER,AUDIT_ADMIN";
```

## 3.15.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 4. Quickstart guide

### 4.1 Quickstart guide for Percona Server for MySQL

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

#### 4.1.1 Install Percona Server for MySQL

You can install Percona Server for MySQL using different methods.

- [Use the Percona Repositories](#)
- [Use APT](#)
- [Use YUM](#)
- [Use binary tarballs](#)
- [Use Docker](#)

#### 4.1.2 For backups and restores

Percona XtraBackup (PXB) is a 100% open source backup solution for all versions of Percona Server for MySQL and MySQL® that performs online non-blocking, tightly compressed, highly secure full backups on transactional systems. Maintain fully available applications during planned maintenance windows with Percona XtraBackup.

[Install Percona XtraBackup](#)

#### 4.1.3 For Monitoring and Management

Percona Monitoring and Management (PMM) monitors and provides actionable performance data for MySQL variants, including Percona Server for MySQL, Percona XtraDB Cluster, Oracle MySQL Community Edition, Oracle MySQL Enterprise Edition, and MariaDB. PMM captures metrics and data for the InnoDB, XtraDB, and MyRocks storage engines, and has specialized dashboards for specific engine details.

[Install PMM and connect your MySQL instances to it.](#)

#### 4.1.4 For high availability

Percona XtraDB Cluster (PXC) is a 100% open source, enterprise-grade, highly available clustering solution for MySQL multi-master setups based on Galera. PXC helps enterprises minimize unexpected downtime and data loss, reduce costs, and improve performance and scalability of your database environments supporting your critical business applications in the most demanding public, private, and hybrid cloud environments.

[Percona XtraDB Cluster Quick start guide](#)

## 4.1.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-11-23

## 5. Install

### 5.1 Install Percona Server for MySQL

Before installing, read the [Percona Server for MySQL 8.1 Release notes](#).

We gather [Telemetry data](#) in the Percona packages and Docker images.

#### 5.1.1 Install Percona Server for MySQL from Repositories

Percona provides repositories for yum (RPM packages for Red Hat, CentOS and Amazon Linux AMI) and apt (.deb packages for Ubuntu and Debian) for software such as Percona Server for MySQL, Percona XtraBackup, and Percona Toolkit. This makes it easy to install and update your software and its dependencies through your operating system's package manager. This is the recommended way of installing where possible.

The following guides describe the installation process for using the official Percona repositories for the .deb and .rpm packages.

[Install Percona Server for MySQL on Debian and Ubuntu](#) [:material-arrow-right:](#)

[Install Percona Server for MySQL on Red Hat Enterprise Linux and CentOS](#) [:material-arrow-right:](#)

#### 5.1.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 5.2 Before you start

### 5.2.1 Percona Product Download Instructions

#### Select the software

Do the following steps to select the software:

1. Open [Percona Product Downloads](#)
2. Locate the Percona Software, for example, Percona Server for MySQL
3. In `Select Product`, select the which product, for example, Percona Server 8.1
4. In `Select Product Version`, select the version, for example, PERCONA-SERVER-8.1.0-1

5. In `Select Software`, select the operating system, for example, `RED HAT ENTERPRISE LINUX / CENTOS / ORACLE LINUX 9`

The easiest method is to download all packages.

### Download to a local computer

In `Package Download Options`, select a specific package or select the `DOWNLOAD ALL PACKAGES` button

The selected packages are downloaded to the local computer.

### Download to another computer

In `Package Download Options`, select a specific package or select the `DOWNLOAD ALL PACKAGES` button, and hover your cursor over the `DOWNLOAD` arrow. Right-click and in the drop-down menu, select `Copy Link`.

Paste the link in your terminal to download the selected package.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 5.3 Use APT

### 5.3.1 Use an APT repository to install Percona Server for MySQL 8.1

Ready-to-use packages are available from the Percona Server for MySQL software repositories and the [Percona downloads](#) page.

Specific information on the supported platforms, products, and versions is described in [Percona Software and Platform Lifecycle](#).

We gather [Telemetry data](#) in the Percona packages and Docker images.

#### Install Percona Server for MySQL using APT

To install Percona Server for MySQL using APT, do the following steps:

0.  Update the package repositories:

```
$ sudo apt update
```

0. 2 Install the `curl` download utility if needed:

```
$ sudo apt install curl
```

0. 3 Download the `percona-release` repository package:

```
$ curl -O https://repo.percona.com/apt/percona-release_latest.generic_all.deb
```

0. 4 Install the downloaded package with `apt` as root or with `sudo`:

```
$ sudo apt install gnupg2 lsb-release ./percona-release_latest.generic_all.deb
```

0. 5 Refresh the local cache to update the package information:

```
$ sudo apt update
```

0. 6 Use `percona-release` to set up the repository for the Percona Server for MySQL 8.1 version:

```
$ sudo percona-release enable-only ps-8x-innovation release
$ sudo percona-release enable tools release
```

0. 7 You can check the repository setup for the Percona original release list in `/etc/apt/sources.list.d/percona-original-release.list`.

0. 8 Install the server package with the `percona-release` command:

```
$ sudo apt install percona-server-server
```

See [Configuring Percona repositories with `percona-release`](#) for more information.

Percona Server for MySQL 8.1 comes with the MyRocks storage engine. This storage engine is installed as a plugin. For information on installing and configuring MyRocks, refer to the [Percona MyRocks Installation Guide](#).

Percona Server for MySQL contains user-defined functions from [Percona Toolkit](#). These user-defined functions provide faster checksums. For more details on the user-defined functions, see [Percona Toolkit UDF functions](#).

After the installation completes, run the following commands to create these functions:

```
mysql -e "CREATE FUNCTION fnv1a_64 RETURNS INTEGER SONAME 'libfnv1a_udf.so'"
mysql -e "CREATE FUNCTION fnv_64 RETURNS INTEGER SONAME 'libfnv_udf.so'"
mysql -e "CREATE FUNCTION murmur_hash RETURNS INTEGER SONAME 'libmurmur_udf.so'"
```

## Install the Percona Testing repository using APT

Percona offers pre-release builds from the testing repository. To enable it, run `percona-release` with the `testing` argument. Run the following command as root or use the `sudo` command:

```
$ sudo percona-release enable ps-8x-innovation testing
```

These builds should not be run in production. This build may not contain all of the features available in the final release. The features may change without notice.

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

### 5.3.2 Files in the DEB package built for Percona Server for MySQL 8.1

Package	Contains
percona-server-server	The database server itself, the mysql
percona-server-common	The files common to the server and c
percona-server-client	The command line client.
percona-server-dbg	Debug symbols for the server.
percona-server-test	The database test suite.
percona-server-source	The server source.
libperconaserverclient21-dev	Header files needed to compile softw
libperconaserverclient21	The client-shared library. The version change that requires software using source code modified.

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

### 5.3.3 Build APT packages

If you wish to build your own Debian/Ubuntu (dpkg) packages of Percona Server for MySQL, you first need to start with a source tarball, either from the Percona website or by generating your own by following the instructions above ([Installing Percona Server for MySQL from the Git Source Tree](#)).

Extract the source tarball:

```
$ tar xzf Percona-Server-8.1.0-1-Linux.x86_64.ssl102.tar.gz
$ cd Percona-Server-8.1.0-1
```

Copy the Debian packaging in the directory that Debian expects it to be in:

```
$ cp -ap build-ps/debian debian
```

Update the changelog for your distribution (here we update for the unstable distribution - sid), setting the version number appropriately. The trailing one in the version number is the revision of the Debian packaging.

```
$ dch -D unstable --force-distribution -v "8.0.13-3-1" "Update to 8.0.13-3"
```

Build the Debian source package:

```
$ dpkg-buildpackage -S
```

Use sbuild to build the binary package in a chroot:

```
$ sbuild -d sid percona-server-8.1_8.1.0-1.dsc
```

You can give different distribution options to `dch` and `sbuild` to build binary packages for all Debian and Ubuntu releases.

#### Note

[PAM Authentication Plugin](#) is not built with the server by default. In order to build the Percona Server for MySQL with PAM plugin, an additional option `-DWITH_PAM=ON` should be used.

#### Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

### 5.3.4 Install Percona Server for MySQL 8.1 using downloaded DEB packages

Download the packages from [Percona Product Downloads](#). If needed, [Instructions for the Percona Product Download](#) are available.

The following example downloads Percona Server for MySQL 8.1.0-1 release packages for Debian 10:

```
$ wget https://downloads.percona.com/downloads/Percona-Server-8.1/Percona-Server-8.1.0-1/binary/debian/buster/x86_64/Percona-Server-8.1.0-1-r71449379-buster-x86_64-bundle.tar
```

Unpack the download to get the packages:

```
$ tar xvf Percona-Server-8.1.0-1-r71449379-buster-x86_64-bundle.tar
```

#### Expected output

```
libperconaserverclient21_8.1.0-1-1.buster_amd64.deb
libperconaserverclient21-dev_8.1.0-1-1.buster_amd64.deb
percona-mysql-router_8.1.0-1-1.buster_amd64.deb
percona-server-client_8.1.0-1-1.buster_amd64.deb
percona-server-common_8.1.0-1-1.buster_amd64.deb
percona-server-dbg_8.1.0-1-1.buster_amd64.deb
percona-server-rocksdb_8.1.0-1-1.buster_amd64.deb
percona-server-server_8.1.0-1-1.buster_amd64.deb
percona-server-source_8.1.0-1-1.buster_amd64.deb
percona-server-test_8.1.0-1-1.buster_amd64.deb
```

Install Percona Server for MySQL using `dpkg`. Run this command as root or use the `sudo` command:

```
$ sudo dpkg -i *.deb
```

#### Warning

When installing packages manually like this, you'll need to resolve all the dependencies and install missing packages yourself. The following packages will need to be installed before you can manually install Percona Server: `mysql-common`, `libjemalloc1`, `libaio1`, and `libmecab2`.

#### Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

### 5.3.5 Apt pinning the Percona Server for MySQL 8.1 packages

Pinning allows you to stay on a release and get packages from a different version. In some cases, you can pin selected packages and avoid accidentally upgrading all the packages.

The pinning takes place in the `preference` file. To pin a package, set the `Pin-Priority` to higher numbers.

Make a new file `/etc/apt/preferences.d/00percona.pref`. For example, add the following to the `preference` file:

```
Package:
Pin: release o=Percona Development Team
Pin-Priority: 1001
```

For more information about the pinning, you can check the official [debian wiki](#).

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

### 5.3.6 Run Percona Server for MySQL 8.1 after APT repository installation

Percona Server for MySQL stores the data files in `/var/lib/mysql/` by default. You can find the configuration file that is used to manage Percona Server for MySQL in `/etc/mysql/my.cnf`.

#### Note

Debian and Ubuntu installation doesn't automatically create a special `debian-sys-maint` user which can be used by the control scripts to control the Percona Server for MySQL `mysqld` and `mysqld_safe` services which was the case with previous Percona Server for MySQL versions. If you still require this user you'll need to create it manually.

Run the following commands as root or by using the `sudo` command

#### 0. 1 Starting the service

Percona Server for MySQL is started automatically after it gets installed unless it encounters errors during the installation process. You can also manually start it by running: `service mysql start`

#### 0. 2 Confirming that service is running. You can check the service status by running: `service mysql status`

#### 0. 3 Stopping the service

You can stop the service by running: `service mysql stop`

0.4 Restarting the service. `service mysql restart`

#### Note

Debian 9.0 (stretch) and Ubuntu 18.04 LTS (bionic) come with `systemd` as the default system and service manager. You can invoke all the above commands with `systemctl` instead of `service`. Currently, both are supported.

### Working with AppArmor

For information on AppArmor, see [Working with AppArmor](#).

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)    [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 5.3.7 Uninstall Percona Server for MySQL 8.1 using the APT package manager

To uninstall Percona Server for MySQL you'll need to remove all the installed packages. Removing packages with `apt remove` does not remove the configuration and data files. Removing the packages with `apt purge` does remove the packages with configuration files and data files (all the databases). Depending on your needs you can choose which command better suits you.

1. Stop the Percona Server for MySQL service: `service mysql stop`
2. Remove the packages
  - a. Remove the packages. This will leave the data files (databases, tables, logs, configuration, etc.) behind. In case you don't need them you'll need to remove them manually: `apt remove percona-server\`
  - b. Purge the packages. This command removes all the packages and deletes all the data files (databases, tables, logs, and so on.): `apt purge percona-server\`

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[Community Forum](#) [Get a Percona Expert](#)

Last update: 2023-10-26

## 5.4 Use YUM

### 5.4.1 Install from Percona Software repository

Ready-to-use packages are available from the Percona Server for MySQL software repositories and the [download page](#). The Percona yum repository supports popular RPM-based operating systems. The easiest way to install the Percona RPM repository is to install an RPM that configures yum and installs the [Percona GPG key](#).

We gather [Telemetry data](#) in the Percona packages and Docker images.

#### Supported platforms

Specific information on the supported platforms, products, and versions are described in [Percona Software and Platform Lifecycle](#).

#### Red Hat Certified

Percona Server for MySQL is certified for Red Hat Enterprise Linux 8. This certification is based on common and secure best practices and successful interoperability with the operating system. Percona Server is listed in the [Red Hat Ecosystem Catalog](#).

#### Limitations

The RPM packages for Red Hat Enterprise Linux 7 and the compatible derivatives do not support TLSv1.3, as it requires OpenSSL 1.1.1, which is currently not available on this platform.

#### Install

Install from Percona Software Repository For more information on the Percona Software repositories and configuring Percona Repositories with percona-release, see the Percona Software Repositories Documentation. Run the following commands as a `root` user or with `sudo`.

**Install on Red Hat 7**    **Install on Red Hat 8 or later**

```
$ sudo yum install https://repo.percona.com/yum/percona-release-latest.noarch.rpm
$ sudo percona-release enable-only ps-8x-innovation release
$ sudo percona-release enable tools release
$ sudo yum install percona-server-server

$ sudo yum install https://repo.percona.com/yum/percona-release-latest.noarch.rpm
$ sudo percona-release enable-only ps-8x-innovation release
$ sudo percona-release enable tools release
$ sudo yum install percona-server-server
```

## Available storage engines

Percona Server for MySQL 8.1 comes with the `MyRocks` storage engine. This storage engine is installed as a plugin. For information on how to install and configure MyRocks, refer to the [Percona MyRocks Installation Guide](#).

### PERCONA YUM TESTING REPOSITORY

Percona offers pre-release builds from our testing repository. To subscribe to the testing repository, you enable the testing repository in `/etc/yum.repos.d/percona-release.repo`. To do so, set both `percona-testing-$basearch` and `percona-testing-noarch` to `enabled = 1` (Note that there are three sections in this file: `release`, `testing`, and `experimental` - in this case, it is the second section that requires updating).

You must install the Percona repository first if the installation has not been done already.

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 5.4.2 Files in the RPM package built for Percona Server for MySQL 8.1

Each of the Percona Server for MySQL RPM packages has a particular purpose.

Package	Contains
<code>percona-server-server</code>	Server itself (the <code>mysqld</code> binary)
<code>percona-server-debuginfo</code>	Debug symbols for the server
<code>percona-server-client</code>	Command line client
<code>percona-server-devel</code>	Header files needed to compile software
<code>percona-server-shared</code>	Client shared library.
<code>percona-server-shared-compat</code>	Shared libraries for software compilation. The following libraries are included: <code>libmysqlclient.so.12</code> , <code>libmysqlclient.so.14</code> , <code>libmysqlclient.so.16</code> , and <code>libmysqlclient.so.18</code> .  This package is not included in downstream derivatives.
<code>percona-server-test</code>	Includes the test suite for Percona Server

## Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

### 5.4.3 Install Percona Server for MySQL using downloaded RPM packages

Download the packages from [Percona Product Downloads](#). If needed, [Instructions for the Percona Product Download](#) are available.

The following example downloads *Percona Server for MySQL* 8.1.0-1 release packages for *RHEL 8*.

0. Use `wget` to download the tar file.

```
$ wget https://downloads.percona.com/downloads/Percona-Server-LATEST/Percona-Server-8.1.0-1/binary/redhat/8/x86_64/Percona-Server-8.1.0-1-rc59f87d2854-el8-x86_64-bundle.tar
```

0. 2. Unpack the bundle to get the packages: `tar xvf Percona-Server-8.1.0-1-rc59f87d2854-el8-x86_64-bundle.tar`

0. 3. To view a list of packages, run the following command:

```
$ ls *.rpm
```

The output should look like the following:

#### Expected output

```
percona-icu-data-files-8.1.0-1.1.el8.x86_64.rpm
percona-mysql-router-8.1.0-1.1.el8.x86_64.rpm
percona-mysql-router-debuginfo-8.1.0-1.1.el8.x86_64.rpm
percona-server-client-8.1.0-1.1.el8.x86_64.rpm
percona-server-client-debuginfo-8.1.0-1.1.el8.x86_64.rpm
percona-server-debuginfo-8.1.0-1.1.el8.x86_64.rpm
percona-server-debugsource-8.1.0-1.1.el8.x86_64.rpm
percona-server-devel-8.1.0-1.1.el8.x86_64.rpm
percona-server-rocksdb-8.1.0-1.1.el8.x86_64.rpm
percona-server-rocksdb-debuginfo-8.1.0-1.1.el8.x86_64.rpm
percona-server-server-8.1.0-1.1.el8.x86_64.rpm
percona-server-server-debuginfo-8.1.0-1.1.el8.x86_64.rpm
percona-server-shared-8.1.0-1.1.el8.x86_64.rpm
percona-server-shared-compat-8.1.0-1.1.el8.x86_64.rpm
percona-server-shared-debuginfo-8.1.0-1.1.el8.x86_64.rpm
percona-server-test-8.1.0-1.1.el8.x86_64.rpm
percona-server-test-debuginfo-8.1.0-1.1.el8.x86_64.rpm
```

0. 4 Install `jemalloc` with the following command, if needed:

```
$ wget https://repo.percona.com/yum/release/8/RPMS/x86_64/
jemalloc-3.6.0-1.el8.x86_64.rpm
```

0. 5 In EL8-based *RHEL* distribution or derivatives package installation requires the `mysql` module to be disabled before installing the packages:

```
$ sudo yum module disable mysql
```

0. 6 Install all the packages (for debugging, testing, etc.) with the following command:

```
$ sudo rpm -ivh *.rpm
```

#### Note

When installing packages manually, you must make sure to resolve all dependencies and install any missing packages yourself.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 5.4.4 Run Percona Server for MySQL

Percona Server for MySQL stores the data files in `/var/lib/mysql/` by default. The configuration file used to manage Percona Server for MySQL is the `/etc/my.cnf`.

The following commands start, provide the server status, stop the server, and restart the server.

#### Note

The RHEL distributions and derivatives come with `systemd` as the default system and service manager so you can invoke all of the commands with `systemctl` instead of `service`. Currently, both options are supported.

1. Percona Server for MySQL is not started automatically on the RHEL distributions and derivatives after installation. Start the server with the following command:

```
$ sudo service mysql start
```

2. Review the service status with the following command:

```
$ sudo service mysql status
```

3. Stop the service with the following command:

```
$ sudo service mysql stop
```

4. Restart the service with the following command:

```
$ sudo service mysql restart
```

### SELinux and security considerations

For information on working with SELinux, see [Working with SELinux](#).

The RHEL 8 distributions and derivatives have added [system-wide cryptographic policies component](#). This component allows the configuration of cryptographic subsystems.

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-04

## 5.4.5 Uninstall Percona Server for MySQL

To completely uninstall Percona Server for MySQL, remove all the installed packages and data files.

0. **1** Stop the Percona Server for MySQL service:

```
$ sudo service mysql stop
```

0. **2** Remove the packages:

```
$ sudo yum remove percona-server*
```

0. **3** Remove the data and configuration files:

**Warning**

This step removes all the packages and deletes all the data files (databases, tables, logs, etc.). Take a backup before this operation in case you need the data.

```
$ rm -rf /var/lib/mysql
$ rm -f /etc/my.cnf
```

## Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-04

## 5.5 Use binary tarballs

### 5.5.1 Install Percona Server for MySQL 8.1 from a binary tarball

A binary tarball contains a group of files, including the source code, bundled together into one file using the `tar` command and compressed using `gzip`.

See the list of the [binary tarball available based on the Percona Server for MySQL version](#) to select the right tarball for your environment.

You can download the binary tarballs from the [Linux - Generic section](#) on the download page.

Fetch and extract the correct binary tarball. For example for *Debian 10*:

```
$ wget https://downloads.percona.com/downloads/Percona-Server-8.1/Percona-Server-8.1.0-1/binary/tarball/Percona-Server-8.1.0-1-Linux.x86_64.glibc2.12.tar.gz
```

## Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 5.5.2 Binary tarball file names available based on the Percona Server for MySQL version

For later version of Percona Server for MySQL, the tar files are organized by the `glibc2` version. You can find this version on your operating system with the following command:

```
$ ldd --version
```

### Expected output

```
ldd (Ubuntu GLIBC 2.35-0ubuntu3.1) 2.35
Copyright (C) 2022 Free Software Foundation, Inc.
This is free software; see the source for copying conditions. There is NO
warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
Written by Roland McGrath and Ulrich Drepper.
```

If the `glibc2` version from your operating system is not listed, then this Percona Server for MySQL version does not support that operating system.

### Binary tarball file name organization

The following lists the platform and the associated full binary file name used by Percona Server for MySQL tar files 8.1.0-1.

Platform	Percona Server for MySQL tarball name	glibc2 version
Ubuntu 22.04	Percona-Server-8.1.0-1-Linux.x86_64.glibc2.35-zenfs.tar.gz	glibc2.35
Ubuntu 20.04	Percona-Server-8.1.0-1-Linux.x86_64.glibc2.31.tar.gz	glibc2.31
Ubuntu 18.04	Percona-Server-8.1.0-1-Linux.x86_64.glibc2.27.tar.gz	glibc2.27
Red Hat Enterprise 9	Percona-Server-8.1.0-1-Linux.x86_64.glibc2.34.tar.gz	glibc2.34
Red Hat Enterprise 8	Percona-Server-8.1.0-1-Linux.x86_64.glibc2.28.tar.gz	glibc2.28
Red Hat Enterprise 7	Percona-Server-8.1.0-1-Linux.x86_64.glibc2.17.tar.gz	glibc2.17

The types of files are as follows:

Type	Name	Description
Full	Percona-Server--Linux.x86_64..tar.gz	Contains all files
Minimal	Percona-Server-<version-number>-Linux.x86_64.<glibc2-version>.minimal.tar.gz	Contains minimal files
Debug	Percona-Server-<version-number>-Linux.x86_64.<glibc2-version>.debug.tar.gz	Contains debug symbols and debug files

Type	Name	Description
Zenfs	Percona-Server-<version-number>-Linux.x86_64.<glibc2-version>-zenfs.tar.gz	Contains the source code for Zenfs or minifuse.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 5.6 Compile from source

### 5.6.1 Install Percona Server for MySQL from a source tarball

Fetch and extract the source tarball. For example:

```
$ wget https://downloads.percona.com/downloads/Percona-Server-8.1/Percona-Server-8.1.0-1/binary/tarball/Percona-Server-8.1.0-1-Linux.x86_64.glibc2.12.tar.gz
$ tar xzf Percona-Server-8.1.0-1-Linux.x86_64.glibc2.12.tar.gz
```

To complete the installation, follow the instructions in [Compile Percona Server for MySQL from Source](#).

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

### 5.6.2 Compile Percona Server for MySQL from source

The following instructions install Percona Server for MySQL 8.1.

#### Install Percona Server for MySQL from the Git Source Tree

Percona uses the [Github](#) revision control system for development. To build the latest Percona Server for MySQL from the source tree, you will need `git` installed on your system.

You can now fetch the latest Percona Server for MySQL 8.1 sources.

```
$ git clone https://github.com/percona/percona-server.git
$ cd percona-server
$ git checkout 8.1
$ git submodule init
$ git submodule update
```

If you are going to be making changes to Percona Server for MySQL 8.1 and wanting to distribute the resulting work, you can generate a new source tarball (exactly the same way as we do for release):

```
$ cmake .
$ make dist
```

After either fetching the source repository or extracting a source tarball (from Percona or one you generated yourself), you will now need to configure and build Percona Server for MySQL.

First, run CMake to configure the build. Here you can specify all the normal build options as you do for a normal MySQL build. Depending on what options you wish to compile Percona Server for MySQL with, you may need other libraries installed on your system. Here is an example using a configure line similar to the options that Percona uses to produce binaries:

```
$ cmake . -DCMAKE_BUILD_TYPE=RelWithDebInfo -DBUILD_CONFIG=mysql_release -
DFEATURE_SET=community
```

## Compile from source

Now, compile using make:

```
$ make
```

Install:

```
$ make install
```

Percona Server for MySQL 8.1 is installed on your system.

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 5.7 Docker

### 5.7.1 Running Percona Server for MySQL in a Docker Container

Percona Server for MySQL has an official Docker image hosted on [Docker Hub](#). Download a specific version by adding the [Docker tag filter for the 8.1 versions](#).

We gather [Telemetry data](#) in the Percona packages and Docker images.

Make sure that you are using the latest version of Docker. The `APT` version or the `YUM` version may be outdated and cause errors.

#### Starting a detached container

Start a container with the `--detached` or `-d` option, which runs the container in the background. In `detached` mode, when the root process used to run the container exits, the container exits.

The following example starts a container named `ps` with the latest version of Percona Server for MySQL 8.1. This action also creates the `root` user and uses `root` as the password. Please note that `root` is not a secure password.

```
$ docker run -d \
  --name ps \
  -e MYSQL_ROOT_PASSWORD=root \
  percona/percona-server:8.1
```

#### Expected output

```
Unable to find image 'percona/percona-server:8.1' locally
8.1: Pulling from percona/percona-server
```

By default, Docker pulls the image from Docker Hub if it is not available locally.

To view the container's logs, use the following command:

```
docker logs ps --follow
```

#### Expected output

```
Initializing database
2022-09-07T15:20:03.158128Z 0 [System] [MY-013169] [Server] /usr/sbin/mysqld (mysqld 8.1.0-1)
initializing of server in progress as process 15
2022-09-07T15:20:03.167764Z 1 [System] [MY-013576] [InnoDB] InnoDB initialization has started.
2022-09-07T15:20:03.530600Z 1 [System] [MY-013577] [InnoDB] InnoDB initialization has ended.
2022-09-07T15:20:04.367600Z 0 [Warning] [MY-013829] [Server] Missing data directory for ICU
regular expressions: /usr/lib64/mysql/private/.
...
2022-09-07T15:20:13.706090Z 0 [System] [MY-011323] [Server] X Plugin ready for connections. Bind-
address: '::' port: 33060, socket: /var/lib/mysql/mysqlx.sock
2022-09-07T15:20:13.706136Z 0 [System] [MY-010931] [Server] /usr/sbin/mysqld: ready for
connections. Version: '8.1.0-1' socket: '/var/lib/mysql/mysql.sock' port: 3306 Percona Server
(GPL), Release 21, Revision c59f87d2854.
```

You can access the server when you see the `ready for connections` information in the log.

## Passing Options

You can pass options with the `docker run` command. For example, the following command uses UTF-8 as the default setting for character set and collation for all databases:

```
[root@docker-host] $ docker run -d \
--name ps \
-e MYSQL_ROOT_PASSWORD=root \
percona/percona-server:8.1 \
--character-set-server=utf8 \
--collation-server=utf8_general_ci
```

## Accessing the Percona Server Container

The `docker exec` command lets you have a shell inside the container. This command uses `it` which forwards your input stream as an interactive TTY.

An example of accessing the detached container:

```
[root@docker-host] $ docker exec -it ps /bin/bash
```

If you need to troubleshoot, the error log is found in `/var/log/` or `/var/log/mysql/`. The file name may be `error.log` or `mysqld.log`.

## Troubleshooting

You can view the error log with the following command:

```
[mysql@ps] $ more /var/log/mysql/error.log
```

### Expected output

```
...
2017-08-29T04:20:22.190474Z 0 [Warning] 'NO_ZERO_DATE', 'NO_ZERO_IN_DATE' and
'ERROR_FOR_DIVISION_BY_ZERO' sql modes should be used with strict mode. They will be merged with
strict mode in a future release.
2017-08-29T04:20:22.190520Z 0 [Warning] 'NO_AUTO_CREATE_USER' sql mode was not set.
...
```

## Accessing the database

You can access the database either with `Docker exec` or using the `mysql` command in the container's shell.

An example of using `Docker exec` to access the database:

```
$ docker exec -ti ps mysql -uroot -proot
```

**Expected output**

```
mysql: [Warning] Using a password on the command line interface can be insecure.
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 9
...
```

Exiting Percona Server also exits the container.

You can also run the MySQL command-line client within the container's shell to access the database:

```
[mysql@ps] $ mysql -uroot -proot
```

**Expected output**

```
mysql: [Warning] Using a password on the command line interface can be insecure.
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 8
Server version: 8.1.0-1 Percona Server (GPL), Release 21, Revision c59f87d2854

Copyright (c) 2009-2022 Percona LLC and/or its affiliates
Copyright (c) 2000, 2022, Oracle and/or its affiliates.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
```

**Accessing the server from an application in another container**

The image exposes the standard MySQL port 3306, so container linking makes the Percona Server instance available from other containers. To link a container running your application (in this case, from an image named `app/image`) with the Percona Server container, run it with the following command:

```
[root@docker-host] $ docker run -d \
  --name app \
  --link ps \
  app/image:latest
```

This application container will be able to access the Percona Server container via port 3306.

## Storing data

There are two ways to store data used by applications that run in Docker containers:

- Let Docker manage the storage of your data by writing the database files to disk on the host system using its internal volume management.
- Create a data directory on the host system on high-performance storage and mount it to a directory visible from the container. This method places the database files in a known location on the host system, and makes it easy for tools and applications on the host system to access the files. The user should ensure that the directory exists, that the user accounts have required permissions, and that any other security mechanisms on the host system are set up correctly.

For example, if you create a data directory on a suitable volume on your host system named `/local/datadir`, you run the container with the following command:

```
[root@docker-host] $ docker run -d \
  --name ps \
  -e MYSQL_ROOT_PASSWORD=root \
  -v /local/datadir:/var/lib/mysql \
  percona/percona-server:8.1
```

The `-v /local/datadir:/var/lib/mysql` option mounts the `/local/datadir` directory on the host to `/var/lib/mysql` in the container, which is the default data directory used by *Percona Server for MySQL*.

Do not add `MYSQL_ROOT_PASSWORD` to the `docker run` command if the data directory contains subdirectories, files, or data.

### Note

If you have SELinux enabled, assign the relevant policy type to the new data directory so that the container will be allowed to access it:

```
[root@docker-host] $ chcon -Rt svirt_sandbox_file_t /local/datadir
```

## Port forwarding

Docker allows mapping ports on the container to ports on the host system using the `-p` option. If you run the container with this option, you can connect to the database by connecting your client to a port on the host machine. This ability simplifies consolidating instances to a single host.

To map the standard MySQL port 3306 to port 6603 on the host:

```
[root@docker-host] $ docker run -d \
  --name ps \
  -e MYSQL_ROOT_PASSWORD=root \
  -p 6603:3306 \
  percona/percona-server:8.1
```

## Exiting the container

If you are in the interactive shell, use CTRL-D or `exit` to exit the session.

If you have a non-shell process running, interrupt the process with `CTRL-C` before using either `CTRL-D` or `exit`.

### Stopping the container

The `docker stop` container command sends a `TERM` signal, then waits 10 seconds and sends a `KILL` signal. The following example stops the `ps` container:

```
$ docker stop ps
```

The default length of time before stopping a container is 10 seconds. A very large instance cannot dump the data from memory to disk within that time. With this type of instance, add the `--time` or the `-t` option to `docker stop`:

```
$ docker stop ps -t 600
```

### Removing the container

To remove a stopped container, use the `docker rm` command.

```
$ docker rm ps
```

### For more information

Review the [Docker Docs](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 5.7.2 Docker environment variables

When running a Docker container with Percona Server, you can adjust the configuration of the instance. Add one or more environment variables to the `docker run` command.

These variables will not affect you if you start the container with a data directory that already contains a database. Any pre-existing database remains untouched on the container startup.

The variables are optional, but you must specify at least one of the following:

- `MYSQL_DATABASE` - the database schema name that is created when the container starts
- `MYSQL_USER` - create a user account when the container starts
- `MYSQL_PASSWORD` - used with `MYSQL_USER` to create a password for that user account.
- `MYSQL_ALLOW_EMPTY_PASSWORD` - creates a root user with an empty password. This option is insecure and only should be used for testing or proof of concept when the database can be removed afterward. Anyone can connect as `root`.
- `MYSQL_ROOT_PASSWORD` - this password is used for the `root` user account. This option is not recommended for production.
- `MYSQL_RANDOM_ROOT_PASSWORD` - set this variable instead of `MYSQL_ROOT_PASSWORD` when you want Percona Server to generate a password for you. The generated password is available in the container's logs only during the first start of the container. Use `docker logs`. You cannot retrieve the password after the first start.

To further secure your instance, use the `MYSQL_ONETIME_PASSWORD` variable.

These variables are visible to anyone able to run `Docker inspect`.

```
$ docker inspect ps
```

 **Expected output**

```
...
  "Env": [
    "MYSQL_ROOT_PASSWORD=root",
    "PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin",
    "PS_VERSION=8.1.0-1",
    "OS_VER=e18",
    "FULL_PERCONA_VERSION=8.1.0-1.e18"
  ]
  ...
```

You should use `Docker secrets` or volumes instead.

Percona Server for MySQL also allows adding the `_FILE` suffix to a variable name. This suffix lets you add the value in a path so that the value cannot be inspected from outside the container.

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 5.8 Upgrade

### 5.8.1 Upgrade from 8.0 to 8.1 overview

Upgrading your server to 8.1 has the following benefits:

Benefits	Description
Security fixes	These patches and updates protect your data and address vulnerabilities or bugs in the software.
New or improved features	You have access to new or improved functionality, performance, and availability.
Reduced labor	You can automate some routine tasks.
Relevance	Your customers and stakeholders have the latest version. Using the latest version can help to drive business growth.
Reduced operational costs	An upgraded database server can help reduce costs because the server has improved efficiency.

Not upgrading your database can have the following risks:

Risks	Description
Security risks	Your database server is vulnerable to attacks and receive security fixes. These attacks can result in data loss and data corruption. These actions can result in downtime and lose money.
Service risks	You do not benefit from new or improved features, user experience, reduced productivity, and higher costs.
Support risks	You are limited in support access. This can result in longer times, unresolved issues, and higher costs.
Compatibility risks	You can experience compatibility issues with applications since the older version is no longer supported. At some point, the database server is no longer supported.
Failure risk	A failure in either hardware, operating system, or software can result in downtime and higher costs if you upgrade at the wrong time.

Create a test environment to verify the upgrade before you upgrade the production servers. The test environment is crucial to the success of the upgrade. There is no supported [downgrade procedure](#). You can try to replicate from an 8.1 version to 8.0 version or restore a backup.

Several tools in the [Percona Toolkit](#) can help with the upgrade process.

We recommend upgrading to the latest version.

Review the documentation for other changes between 8.0 to 8.1.

Review [Upgrade Strategies](#) for an overview of the major strategies.

The following list summarizes a number of the changes in the 8.0 series and has useful guides that can help you perform a smooth upgrade. We strongly recommend reading this information:

- [Upgrading MySQL](#)
- [Before You Begin](#)
- [Upgrade Paths](#)
- [Changes in MySQL 8.0](#)
- [Preparing your Installation for Upgrade](#)
- [Percona Server for MySQL 8.1 Release notes](#)
- [Upgrade Troubleshooting](#)
- [Rebuilding or Repairing Tables or Indexes](#)

Review other [Percona blogs](#) that contain upgrade information.

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 5.8.2 Upgrade strategies

There are different strategies to consider when upgrading from 8.0 to 8.1.

### In-place upgrade

An upgrade to 8.1 does not allow a rollback. The in-place upgrade strategy is not recommended, and it should be used only as a last resort.

An in-place upgrade involves shutting down the 8.0 server, and replacing the server binaries, or packages, with new ones. At this point the new server version can be started on the existing data directory. Note that the server should be configured to perform a slow shutdown by setting `innodb_fast_shutdown=0` prior to shutdown.

The benefits are:

- Less additional infrastructure cost compared to a new environment, but nodes must be tested.
- An upgrade can be completed over weeks with cool-down periods between reader node upgrades.
- Requires a failover of production traffic, and for minimal downtime you must have good high-availability tools.

If you use XA transactions with InnoDB, running XA RECOVER before upgrading checks for uncommitted XA transactions. If results are returned, either commit or rollback the XA transactions by issuing an XA COMMIT or XA ROLLBACK statement.

### New environment with cut over

Upgrading with a new environment involves provisioning a duplicate environment with the same number of servers with the same hardware specs and same operating system as the current production nodes.

On the newly provided hardware, the target MySQL version will be installed. The new environment will be set up, and the production data will be recovered. Remember that you can use pt-config-diff to verify MySQL configurations.

Replication from the current source to the newly built environment will be established. At cutover time, all writes on the current source will be halted, and the application traffic will need to be redirected to the new source. The cutover can be done using a Virtual IP address or manually redirecting the application itself. Once writes are being received on the new environment, you are in a fail forward situation, and the old environment can be torn down.

The new environment strategy has the following pros and cons:

- Additional infrastructure cost since a new environment must be built.
- Ability to upgrade both the OS and the DBMS at the same time.
- Allows upgrade of hardware easily.
- Requires only a single cutover window.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 5.8.3 Upgrading from plugins to components

The following plugins have changed:

Plugin	8.0 information	8.1 changes
<code>keyring_vault</code>	Only available as a plugin	<code>component_keyring_vault</code>

Plugin	8.0 information	8.1 changes
<code>audit_log</code>	Only available as a plugin	removed
<code>audit_log_filter</code>	Only available as a plugin	<code>component_audit_log_filter</code>
<code>data_masking</code>	Available as a plugin and component	<code>component_masking_functions</code>
<code>binlog_utils_udf</code> user defined functions	Only available as plugin. Users must install the plugin and then run <code>CREATE FUNCTION ... SONAME...</code>	<code>component_binlog_utils_udf</code>
<code>percona-udf</code> user defined functions	Must create individual functions with <code>CREATE FUNCTION ... SONAME ...</code> .	<code>component_percona_udf</code>

We recommend if you use a plugin and the feature also available a component, switch to the component in 8.0 series before upgrading to 8.1.

### Transition from a plugin to a component

The operation to transition from a plugin to a component can be complicated. You should plan for downtime while you plan and test each step in the procedure.

Before you start, review the differences between the plugin and the component. A plugin configuration has plugin-specific system variables and uses the `--early-plugin-load` option. A component has a configuration file and loads using a manifest.

1. Setup the component's configuration file.
2. Use the manifest to load the component.
3. Confirm that the component works. Run queries or other operations and test the component in your environment.
4. After confirmation, remove the plugin.

Some plugins may require more configuration and setup during the transition to a component. For those plugins, you may have the following scenario:

1. Test the plugin in 8.0.
2. Stop the service
3. Upgrade the packages to 8.1
4. Review and edit configurations, as needed
5. Start 8.1
6. Test the component in 8.1

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 5.8.4 Upgrade using the Percona repositories

We recommend using the Percona repositories to upgrade your server.

Find the instructions on how to enable the repositories in the following documents:

- [Percona APT Repository](#)
- [Percona RPM Repository](#)

**DEB-based distributions**      **RPM-based distributions**

Run the following commands as root or use the `sudo` command.

1. Make a full backup (or dump if possible) of your database. Move the database configuration file, `my.cnf`, to another directory to save it. If the configuration file is not moved, it can be overwritten.
2. Stop the server with the appropriate command for your system:

```
systemctl stop mysql`
```

1. Modify the database configuration file, `my.cnf`, as needed.
2. Install Percona Server for MySQL:

```
$ sudo apt update
$ sudo apt install curl
$ curl -O https://repo.percona.com/apt/percona-release_latest.generic_all.deb
$ sudo apt install gnupg2 lsb-release ./percona-release_latest.generic_all.deb
$ sudo apt update
$ sudo percona-release setup ps-8x-innovation
$ sudo apt install percona-server-server
```

3. Install the storage engine packages.

If you used the MyRocks storage engine in Percona Server for MySQL 8.1, install the `percona-server-rocksdb` package:

```
$ sudo apt install percona-server-rocksdb
```

4. Running the upgrade:

The `mysqld` binary automatically runs the upgrade process if needed. To find more information, see [MySQL Upgrade Process Upgrades](#)

5. Restart the service

```
$ sudo systemctl restart mysqld
```

After the service has been successfully restarted you can use the new Percona Server for MySQL 8.1.

Run the following commands as root or use the `sudo` command.

1. Make a full backup (or dump if possible) of your database. Copy the database configuration file, for example, `my.cnf`, to another directory to save it.
2. Stop the server with the appropriate command for your system:

```
$ systemctl stop mysql`
```

3. Check your installed packages with `rpm -qa | grep Percona-Server`.
4. Remove only the packages without dependencies and leave dependent packages. The command does not prompt for confirmation:

```
$ rpm -qa | grep Percona-Server | xargs rpm -e --nodeps
```

5. Remove the `mysql`-related packages; run `rpm -qa | grep '^mysql-' | xargs rpm -e --nodeps`

```
$ rpm -qa | grep '^mysql-' | xargs rpm -e --nodeps
```

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 5.8.5 Upgrade using standalone packages

Make a full backup (or dump if possible) of your database. Move the database configuration file, `my.cnf`, to another direction to save it. Stop the server with `/etc/init.d/mysql stop`.

## Debian-derived distribution    Red Hat-derived distributions

1. Remove the installed packages with their dependencies: `sudo apt autoremove percona-server percona-client`
2. Do the required modifications in the database configuration file `my.cnf`.
3. Download the following packages for your architecture:
  - `percona-server-server`
  - `percona-server-client`
  - `percona-server-common`
  - `libperconaserverclient21`

The following example downloads Percona Server for MySQL 8.1.0-1 packages for Debian 11.0:

```
$ wget https://downloads.percona.com/downloads/Percona-Server-LATEST/Percona-Server-8.1.0-1/binary/debian/bullseye/x86_64/Percona-Server-8.1.0-1-bullseye-x86_64-bundle.tar
```

4. Unpack the bundle to get the packages: `tar xvf Percona-Server-8.1.0-1-x86_64-bundle.tar`.

After you unpack the bundle, you should see the following packages:

```
$ ls *.deb
```

### Expected output

```
libperconaserverclient21-dev_8.1.0-1.bullseye_amd64.deb
percona-server-dbg_8.1.0-1.bullseye_amd64.deb
libperconaserverclient21_8.1.0-1.bullseye_amd64.deb
percona-server-rocksdb_8.1.0-1.bullseye_amd64.deb
percona-mysql-router_8.1.0-1.bullseye_amd64.deb
percona-server-server_8.1.0-1.bullseye_amd64.deb
percona-server-client_8.1.0-1.bullseye_amd64.deb
percona-server-source_8.1.0-1.bullseye_amd64.deb
percona-server-common_8.1.0-1.bullseye_amd64.deb
percona-server-test_8.1.0-1.bullseye_amd64.deb
```

5. Install Percona Server for MySQL:

```
$ sudo dpkg -i *.deb
```

This command installs the packages from the bundle. Another option is to download or specify only the packages you need for running Percona Server for MySQL installation (`libperconaserverclient21-dev_8.1.0-1.bullseye_amd64.deb`, `percona-server-client-8.1.0-1.bullseye_amd64.deb`, `percona-server-common-8.1.0-1.bullseye_amd64.deb`, and `percona-server-server-8.1.0-1.bullseye_amd64.deb`).

### Warning

When installing packages manually, you must resolve all the dependencies and install missing packages yourself. At least the following packages should be installed before installing Percona Server for MySQL 8.1.0-1: `* libmecab2, * libjemalloc1, * zlib1g-dev, * libaio1`.

6. Running the upgrade:

The `mysqld` binary automatically runs the upgrade process. To find more information, see [MySQL Upgrade Process Upgrades](#)

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[Community Forum](#) [Get a Percona Expert](#)

Last update: 2023-11-27

## 5.8.6 Downgrade Percona Server for MySQL

Downgrading to a 5.7 or earlier series is not supported.

The following table lists the downgrade paths:

Path	Examples	Methods
Within a bugfix release series or an LTS release	8.0.35 to 8.0.34	in-place asynchron
A bugfix or LTS release to the last bugfix or LTS release series	8.4.x to 8.0.x	logical c
A bugfix or LTS release to an innovation release after the last LTS series	8.4.x to 8.3.0	logical c
An innovation release to another innovation release after the last LTS series	8.3.0 to 8.2.0	logical c

We don't support downgrades with any 8.0.x release below 8.0.34. A releases in the range above 8.0.34 can be downgraded to any release within that range, including 8.0.34.

### Downgrading risks

Downgrading has the following risks:

Risk	Description
Data loss	If the downgrade process has issues, you back up your data before attempt
Incompatibility	If you use any feature or improvement could result in incompatibility issues.
Performance	Downgrading may result in a loss of p
Security	Newer versions have security updates versions, which could lead to exposu

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)

[:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 6. Post-Installation

### 6.1 Binary logs and replication improvements

Due to continuous development, Percona Server for MySQL incorporated a number of improvements related to replication and binary logs handling. This resulted in replication specifics, which distinguishes it from MySQL.

#### 6.1.1 Safety of statements with a LIMIT clause

##### Summary of the fix

MySQL considers all `UPDATE/DELETE/INSERT ... SELECT` statements with `LIMIT` clause to be unsafe, no matter whether they are really producing non-deterministic result or not, and switches from statement-based logging to row-based one. *Percona Server for MySQL* is more accurate, it acknowledges such instructions as safe when they include `ORDER BY PK` or `WHERE` condition. This fix has been ported from the upstream bug report [#42415](#) ([#44](#)).

#### 6.1.2 Performance improvement on relay log position update

##### Relay log position fix

MySQL always updated relay log position in multi-source replications setups regardless of whether the committed transaction has already been executed or not. *Percona Server* omits relay log position updates for the already logged GTIDs.

##### Relay log position details

Particularly, such unconditional relay log position updates caused additional `fsync` operations in case of `relay-log-info-repository=TABLE`, and with the higher number of channels transmitting such duplicate (already executed) transactions the situation became proportionally worse. Bug fixed [#1786](#), (upstream [#85141](#)).

#### 6.1.3 Performance improvement on source and connection status updates

##### Source and connection status update fix

Replica nodes configured to update source status and connection information only on log file rotation did not experience the expected reduction in load. *MySQL* was additionally updating this information in case of multi-source replication when replica had to skip the already executed GTID event.

##### Source and connection status details

The configuration with `master_info_repository=TABLE` and `sync_master_info=0` makes replica to update source status and connection information in this table on log file rotation and not after each `sync_master_info` event, but it didn't work on multi-source replication setups. Heartbeats sent to the replica to skip GTID events which it had already executed previously, were evaluated as relay log rotation events and reacted with `mysql.slave_master_info` table sync. This inaccuracy could produce

huge (up to 5 times on some setups) increase in write load on the replica, before this problem was fixed in *Percona Server for MySQL*. Bug fixed [#1812](#) (upstream [#85158](#)).

## 6.1.4 Write FLUSH commands to the binary log

FLUSH commands, such as `FLUSH SLOW LOGS`, are not written to the binary log if the system variable `binlog_skip_flush_commands` is set to **ON**.

In addition, the following changes were implemented in the behavior of `read_only` and `super_read_only` modes:

- When `read_only` is set to **ON**, any `FLUSH ...` command executed by a normal user (without the `SUPER` privilege) are not written to the binary log regardless of the value of the `binlog_skip_flush_command` variable.
- When `super_read_only` is set to **ON**, any `FLUSH ...` command executed by any user (even by those with the `SUPER` privilege) are not written to the binary log regardless of the value of the `binlog_skip_flush_commands` variable.

An attempt to run a `FLUSH` command without either `SUPER` or `RELOAD` privileges results in the `ER_SPECIFIC_ACCESS_DENIED_ERROR` exception regardless of the value of the `binlog_skip_flush_commands` variable.

### `binlog_skip_flush_commands`

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Default	OFF

When `binlog_skip_flush_commands` is set to **ON**, `FLUSH ...` commands are not written to the binary log. See [Writing FLUSH Commands to the Binary Log](#) for more information about what else affects the writing of `FLUSH` commands to the binary log.

#### Note

`FLUSH LOGS`, `FLUSH BINARY LOGS`, `FLUSH TABLES WITH READ LOCK`, and `FLUSH TABLES ... FOR EXPORT` are not written to the binary log no matter what value the `binlog_skip_flush_commands` variable contains. The `FLUSH` command is not recorded to the binary log and the value of `binlog_skip_flush_commands` is ignored if the `FLUSH` command is run with the `NO_WRITE_TO_BINLOG` keyword (or its alias `LOCAL`).

## 6.1.5 Maintaining comments with DROP TABLE

When you issue a `DROP TABLE` command, the binary log stores the command but removes comments and encloses the table name in quotation marks. If you require the binary log to maintain the comments and not add quotation marks, enable `binlog_ddl_skip_rewrite`.

**binlog\_ddl\_skip\_rewrite**

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Default	OFF

If the variable is enabled, single table `DROP TABLE` DDL statements are logged in the binary log with comments. Multi-table `DROP TABLE` DDL statements are not supported and return an error.

```
SET binlog_ddl_skip_rewrite = ON;
/*comment at start*/DROP TABLE t /*comment at end*/;
```

**6.1.6 Binary log user defined functions**

To implement Point in Time recovery, we have added the `binlog_utils_udf`. The following user-defined functions are included:

Name	Returns	Description
<code>get_binlog_by_gtid()</code>	Binlog file name as STRING	Returns specified
<code>get_last_gtid_from_binlog()</code>	GTID as STRING	Returns binlog
<code>get_gtid_set_by_binlog()</code>	GTID set as STRING	Returns
<code>get_binlog_by_gtid_set()</code>	Binlog file name as STRING	Returns contains
<code>get_first_record_timestamp_by_binlog()</code>	Timestamp as INTEGER	Returns specified
<code>get_last_record_timestamp_by_binlog()</code>	Timestamp as INTEGER	Returns specified

**Note**

All functions returning timestamps return their values as microsecond precision UNIX time. In other words, they represent the number of microseconds since 1-JAN-1970.

All functions accepting a binlog name as the parameter accepts only short names, without a path component. If the path separator (`/`) is found in the input, an error is returned. This serves the purpose of restricting the locations from where binlogs can be read. They are always read from the current binlog directory (`@@log_bin_basename` system variable).

All functions returning binlog file names return the name in short form, without a path component.

The basic syntax for `get_binlog_by_gtid()` is the following:

```
* get_binlog_by_gtid(string) [AS] alias
```

Usage: `SELECT get_binlog_by_gtid(string) [AS] alias`

Example:

```
CREATE FUNCTION get_binlog_by_gtid RETURNS STRING SONAME 'binlog_utils_udf.so';
SELECT get_binlog_by_gtid("F6F54186-8495-47B3-8D9F-011DDB1B65B3:1") AS result;
```

#### Expected output

```
+-----+
| result |
+-----+
| binlog.00001 |
+-----+
```

```
DROP FUNCTION get_binlog_by_gtid;
```

The basic syntax for `get_last_gtid_from_binlog()` is the following:

```
* get_last_gtid_from_binlog(string) [AS] alias
```

Usage: `SELECT get_last_gtid_from_binlog(string) [AS] alias`

For example:

```
CREATE FUNCTION get_last_gtid_from_binlog RETURNS STRING SONAME
'binlog_utils_udf.so';
SELECT get_last_gtid_from_binlog("binlog.00001") AS result;
```

#### Expected output

```
+-----+
| result |
+-----+
| F6F54186-8495-47B3-8D9F-011DDB1B65B3:10 |
+-----+
```

```
DROP FUNCTION get_last_gtid_from_binlog;
```

The basic syntax for `get_gtid_set_by_binlog()` is the following:

```
* get_gtid_set_by_binlog(string) [AS] alias
```

Usage: `SELECT get_gtid_set_by_binlog(string) [AS] alias`

For example:

```
CREATE FUNCTION get_gtid_set_by_binlog RETURNS STRING SONAME 'binlog_utils_udf.so';
SELECT get_gtid_set_by_binlog("binlog.000001") AS result;
```

#### Expected output

```
+-----+
| result          |
+=====+
| 11ea-b9a7:7,11ea-b9a7:8 |
+-----+
```

```
DROP FUNCTION get_gtid_set_by_binlog;
```

The basic syntax for `get_binlog_by_gtid_set()` is the following:

- `get_binlog_by_gtid_set(string) [AS] alias`

Usage: `SELECT get_binlog_by_gtid_set(string) [AS] alias`

Example:

```
CREATE FUNCTION get_binlog_by_gtid_set RETURNS STRING SONAME 'binlog_utils_udf.so';
SELECT get_binlog_by_gtid_set("11ea-b9a7:7,11ea-b9a7:8") AS result;
```

#### Expected output

```
+-----+
| result          |
+=====+
| bin.000003      |
+-----+
```

```
DROP FUNCTION get_binlog_by_gtid_set;
```

The basic syntax for `get_first_record_timestamp_by_binlog()` is the following:

- \* `get_first_record_timestamp_by_binlog(TIMESTAMP) [AS] alias`

Usage: `SELECT get_first_record_timestamp_by_binlog(TIMESTAMP) [AS] alias`

For example:

```
CREATE FUNCTION get_first_record_timestamp_by_binlog RETURNS INTEGER SONAME
'binlog_utils_udf.so';
SELECT FROM_UNIXTIME(get_first_record_timestamp_by_binlog("bin.000003")) DIV 1000000
AS result;
```

**Expected output**

```
+-----+
| result          |
+-----+
| 2020-12-03 09:10:40 |
+-----+
```

```
DROP FUNCTION get_first_record_timestamp_by_binlog;
```

The basic syntax for `get_last_record_timestamp_by_binlog()` is the following:

```
* get_last_record_timestamp_by_binlog(TIMESTAMP) [AS] alias
```

Usage: `SELECT get_last_record_timestamp_by_binlog(TIMESTAMP) [AS] alias`

For example:

```
CREATE FUNCTION get_last_record_timestamp_by_binlog RETURNS INTEGER SONAME
'binlog_utils_udf.so';
SELECT FROM_UNIXTIME(get_last_record_timestamp_by_binlog("bin.00003")) DIV 1000000)
AS result;
```

**Expected output**

```
+-----+
| result          |
+-----+
| 2020-12-04 04:18:56 |
+-----+
```

```
DROP FUNCTION get_last_record_timestamp_by_binlog;
```

## 6.1.7 Limitations

For the following variables, do not define values with one or more dot (.) characters:

- `log_bin`
- `log_bin_index`

A value defined with these characters is handled differently in *MySQL* and *Percona XtraBackup* and can cause unpredictable behavior.

## 6.1.8 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 6.2 Post-installation

Depending on the type of installation, you may need to do the following tasks:

### 6.2.1 Installed using binary files or compiling from source

Task
<a href="#">Initialize the data dictionary</a>
<a href="#">Test the server</a>
<a href="#">Set service to start at boot time</a>

#### Initialize the data directory

If you install the server using either the source distribution or generic binary distribution files, the data directory is not initialized, and you must run the initialization process after installation.

Run `mysqld` with the `-initialize` option or the `initialize-insecure` option.

Executing `mysqld` with either option does the following:

- Verifies the existence of the data directory
- Initializes the system tablespace and related structures
- Creates system tables including grant tables, time zone tables, and server-side help tables
- Creates `root@localhost`

You should run the following steps with the `mysql` login.

4.  Navigate to the MySQL directory. The example uses the default location.

```
$ cd /usr/local/mysql
```

4.  Create a directory for the MySQL files. The `secure_file_priv` uses the directory path as a value.

```
$ mkdir mydata
```

The `mysql` user account should have the `drwxr-x---` permissions. Four sections define the permissions; file or directory, User, Group, and Others.

The first character designates if the permissions are for a file or directory. The first character is `d` for a directory.

The rest of the sections are specified in three-character sets.

Permission	User	Group
Read	Yes	Yes
Write	Yes	No
Execute	Yes	Yes

4. Run the command to initialize the data directory.

```
$ bin/mysqld --initialize
```

### Test the server

After you have initialized the data directory, and the server is started, you can run tests on the server.

This section assumes you have used the default installation settings. If you have modified the installation, navigate to the installation location. You can also add the location by [Setting the Environment Variables](#).

You can use the [mysqladmin](#) client to access the server.

If you have issues connecting to the server, use the `root` user and the root account password.

```
$ sudo mysqladmin -u root -p version
```

#### Expected output

```
Enter password:
mysql Ver 8.1.0-1 for debian-linux-gnu on x86_64 (Percona Server (GPL), Release '10', Revision
'f446c04')
...
Server version      8.1.0-1
Protocol version    10
Connection          Localhost via UNIX socket
UNIX socket         /var/run/mysqld/mysqld.sock
Uptime:             4 hours 58 min 10 section

Threads: 2 Questions: 16 Slow queries: 0 Opens: 139 Flush tables: 3
Open tables: 59 Queries per second avg: 0.0000
```

Use [mysqlshow](#) to display database and table information.

```
$ sudo mysqlshow -u root -p
```

```

Expected output

Enter password:

+-----+
|   Databases   |
+=====+
| information_schema |
+-----+
| mysql          |
+-----+
| performance_schema |
+-----+
| sys            |
+-----+

```

### Set service to run at boot time

After a generic binary installation, manually configure systemd support.

The following commands start, check the status, and stop the server:

```

$ sudo systemctl start mysqld
$ sudo systemctl status mysqld
$ sudo systemctl stop mysqld

```

Run the following command to start the service at boot time:

```

$ sudo systemctl enable mysqld

```

Run the following command to prevent a service from starting at boot time:

```

$ sudo systemctl disable mysqld

```

## 6.2.2 All installations

Task
<a href="#">Update the root password</a>
<a href="#">Secure the server</a>
<a href="#">Populate the time zone tables</a>

### Update the root password

During an installation on Debian/Ubuntu, you are prompted to enter a root password. On Red Hat Enterprise Linux and derivatives, you update the root password after installation.

Restart the server with the `--skip-grant-tables` option to allow access without a password. This option is insecure. This option also disables remote connections.

```

$ sudo systemctl stop mysqld
$ sudo systemctl set-environment MYSQLD_OPTS="--skip-grant-tables"

```

```
$ sudo systemctl start mysqld
$ mysql
```

Reload the grant tables to be able to run the `ALTER USER` statement. Enter a password that satisfies the current policy.

```
mysql> FLUSH PRIVILEGES;
mysql> ALTER USER 'root'@'localhost' IDENTIFIED BY 'rootPassword_12';
mysql> exit
```

If, when adding the password, MySQL returns `ERROR 1819 (HY000) Your password does not satisfy the current policy`, run the following command to see policy requirement.

```
mysql> SHOW VARIABLES LIKE 'validate_password%';
```

Redo your password to satisfy the requirements.

Stop the server, remove the `--skip-grant-tables` option, start the server, and log into the server with the updated password.

```
$ sudo systemctl stop mysqld
$ sudo systemctl unset-environment MYSQLD_OPTS
$ sudo systemctl start mysqld
$ mysql -u root -p
```

## Secure the server

The `mysql_secure_installation` script improves the security of the instance.

The script does the following:

- Changes the `root` password
- Disallows remote login for `root` accounts
- Removes anonymous users
- Removes the `test` database
- Reloads the privilege tables

The following statement runs the script:

```
$ mysql_secure_installation
```

## 6.2.3 Populate the time zone tables

The time zone system tables are the following:

- `time_zone`
- `time_zone_leap_second`
- `time_zone_name`
- `time_zone_transition`
- `time_zone_transition_type`

If you install the server using either the source distribution or the generic binary distribution files, the installation creates the time zone tables, but the tables are not populated.

The `mysql_tzinfo_to_sql` program populates the tables from the `zoneinfo` directory data available in Linux.

A common method to populate the tables is to add the `zoneinfo` directory path to `mysql_tzinfo_to_sql` and then send the output into the `mysql` system schema.

The example assumes you are running the command with the `root` account. The account must have the privileges for modifying the `mysql` system schema.

```
$ mysql_tzinfo_to_sql /usr/share/zoneinfo | mysql -u root -p -D mysql
```

## 6.2.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 6.3 Working with AppArmor

The operating system has a Discretionary Access Controls (DAC) system. AppArmor supplements the DAC with a Mandatory Access Control (MAC) system. AppArmor is the default security module for Ubuntu or Debian systems and uses profiles to define how programs access resources.

AppArmor is path-based and restricts processes by using profiles. Each profile contains a set of policy rules. Some applications may install their profile along with the application. If an installation does not also install a profile, then that application is not part of the AppArmor subsystem. You can also create profiles since they are simple text files stored in the `/etc/apparmor.d` directory.

A profile is in one of the following modes:

- Enforce – the default setting, applications are prevented from taking actions restricted by the profile rules.
- Complain – applications are allowed to take restricted actions, and the actions are logged.
- Disabled – applications are allowed to take restricted actions, and the actions are not logged.

You can mix enforce profiles and complain profiles in your server.

### 6.3.1 Install the utilities used to control AppArmor

Install the `apparmor-utils` package to work with profiles. Use these utilities to create, update, enforce, switch to complain mode, and disable profiles, as needed:

```
$ sudo apt install apparmor-utils
```

#### Expected output

```
Reading package lists... Done
Building dependency tree
...
The following additional packages will be installed:
  python3-apparmor python3-libapparmor
...
```

### 6.3.2 Check the current status

As root or using `sudo`, you can check the AppArmor status:

```
$ sudo aa-status
```

#### Expected output

```
apparmor module is loaded.
34 profiles are loaded.
32 profiles in enforce mode.
...
  /usr/sbin/mysqld
...
2 profiles in complain mode.
...
3 profiles have profiles defined.
...
0 processes are in complain mode.
0 processes are unconfined but have a profile defined.
```

### 6.3.3 Switch a profile to complain mode

Switch a profile to complain mode when the program is in your path with this command:

```
$ sudo aa-complain <program>
```

If needed, specify the program's path in the command:

```
$ sudo aa-complain /sbin/<program>
```

If the profile is not stored in `/etc/apparmor.d/`, use the following command:

```
$ sudo aa-complain /path/to/profiles/<program>
```

### 6.3.4 Switch a profile to enforce mode

Switch a profile to the enforce mode when the program is in your path with this command:

```
$ sudo aa-enforce <program>
```

If needed, specify the program's path in the command:

```
$ sudo aa-enforce /sbin/<program>
```

If the profile is not stored in `/etc/apparmor.d/`, use the following command:

```
$ sudo aa-enforce /path/to/profile
```

### 6.3.5 Disable one profile

You can disable a profile but it is recommended to Switch a Profile to Complain mode.

Use either of the following methods to disable a profile:

```
$ sudo ln -s /etc/apparmor.d/usr.sbin.mysqld /etc/apparmor.d/disable/  
$ sudo apparmor_parser -R /etc/apparmor.d/usr.sbin.mysqld
```

or

```
$ aa-disable /etc/apparmor.d/usr.sbin.mysqld
```

### 6.3.6 Reload all profiles

Run either of the following commands to reload all profiles:

```
$ sudo service apparmor reload
```

or

```
$ sudo systemctl reload apparmor.service
```

### 6.3.7 Reload one profile

To reload one profile, run the following:

```
$ sudo apparmor_parser -r /etc/apparmor.d/<profile>
```

For some changes to take effect, you may need to restart the program.

### 6.3.8 Disable AppArmor

AppArmor provides security and disabling the system is not recommended. If AppArmor must be disabled, run the following commands:

- 3.1 Check the status.

```
$ sudo apparmor_status
```

- 3.2 Stop and disable AppArmor.

```
$ sudo systemctl stop apparmor
$ sudo systemctl disable apparmor
```

### 6.3.9 Add the mysql profile

Add the mysql profile with the following procedure:

- 3.1 Download the current version of the AppArmor:

```
$ wget https://raw.githubusercontent.com/mysql/mysql-server/trunk/packaging/debian/extra/apparmor-profile
```

The expected output:

```
...
Saving to 'apparmor-profile`
...`
```

- 3.2 Move the file to `/etc/apparmor.d/usr.sbin.mysqld`

```
$ sudo mv apparmor-profile /etc/apparmor.d/usr.sbin.mysqld
```

- 3.3 Create an empty file for editing:

```
$ sudo touch /etc/apparmor.d/local/usr.sbin.mysqld
```

- 3.4 Load the profile:

```
$ sudo apparmor_parser -r -T -W /etc/apparmor.d/usr.sbin.mysqld
```

### 3. Restart *Percona Server for MySQL*:

```
$ sudo systemctl restart mysql
```

### 3. Verify the profile status:

```
$ sudo aa-status
```

#### Expected output

```
...
processes are in enforce mode
...
/usr/sbin/mysqld (100840)
...
```

## 6.3.10 Edit the mysqld profile

Only edit `/etc/apparmor.d/local/usr.sbin.mysqld`. We recommend that you Switch a Profile to Complain mode before editing the file. Edit the file in any text editor. When your work is done, Reload one profile and Switch a Profile to Enforce mode.

## 6.3.11 Configure a custom data directory location

You can change the data directory to a non-default location, like `/var/lib/mysqlcustom`. You should enable audit mode, to capture all of the actions, and edit the profile to allow access for the custom location.

```
$ cat /etc/mysql/mysql.conf.d/mysqld.cnf
```

#### Expected output

```
The Percona Server 8.1 configuration file.

For explanations see
https://dev.mysql.com/doc/mysql/en/server-system-variables.html

[mysqld]
pid-file      = /var/run/mysqld/mysqld.pid
socket        = /var/run/mysqld/mysqld.sock
*datadir      = /var/lib/mysqlcustom*
log-error     = /var/log/mysql/error.log
```

Enable audit mode for mysqld. In this mode, the security policy is enforced and all access is logged.

```
$ aa-audit mysqld
```

Restart Percona Server for MySQL.

```
$ sudo systemctl mysql restart
```

The restart fails because AppArmor has blocked access to the custom data directory location. To diagnose the issue, check the logs for the following:

- ALLOWED - A log event when the profile is in complain mode and the action violates a policy.
- DENIED - A log event when the profile is in enforce mode and the action is blocked.

For example, the following log entries show DENIED :

 **Expected output**

```
...
Dec 07 12:17:08 ubuntu-s-4vcpu-8gb-nyc1-01-aa-ps audit[16013]: AVC apparmor="DENIED"
operation="mknod" profile="/usr/sbin/mysqld" name="/var/lib/mysqlcustom/binlog.index" pid=16013
comm="mysqld" requested_mask="c" denied_mask="c" fsuid=111 ouid=111
Dec 07 12:17:08 ubuntu-s-4vcpu-8gb-nyc1-01-aa-ps kernel: audit: type=1400
audit(1607343428.022:36): apparmor="DENIED" operation="mknod" profile="/usr/sbin/mysqld" name="/
var/lib/mysqlcustom/mysqld_tmp_file_case_insensitive_test.lower-test" pid=16013 comm="mysqld"
requested_mask="c" denied_mask="c" fsuid=111 ouid=111
...
```

Open `/etc/apparmor.d/local/usr.sbin.mysqld` in a text editor and edit the following entries in the `Allow data dir access` section.

```
Allow data dir access
/var/lib/mysqlcustom/ r,
/var/lib/mysqlcustom/** rwk,
```

In `etc/apparmor.d/local/usr.sbin.mysqld`, comment out, using the `#` symbol, the current entries in the `Allow data dir access` section. This step is optional. If you skip this step, `mysqld` continues to access the default data directory location.

#### Note

Edit the local version of the file instead of the main profile. Separating the changes makes maintenance easier.

Reload the profile:

```
$ apparmor_parser -r -T /etc/apparmor.d/usr.sbin.mysqld
```

Restart `mysql`:

```
$ systemctl restart mysqld
```

## 6.3.12 Set up a custom log location

To move your logs to a custom location, you must edit the `my.cnf` configuration file and then edit the local profile to allow access:

```
cat /etc/mysql/mysql.conf.d/mysqld.cnf
```

### Expected output

```
The Percona Server 8.1 configuration file.

For explanations see
https://dev.mysql.com/doc/mysql/en/server-system-variables.html

[mysqld]
pid-file      = /var/run/mysqld/mysqld.pid
socket        = /var/run/mysqld/mysqld.sock
datadir       = /var/lib/mysql
log-error     = /*custom-log-dir*/mysql/error.log
```

Verify the custom directory exists.

```
$ ls -la /custom-log-dir/
```

### Expected output

```
total 12
drwxrwxrwx 3 root root 4096 Dec  7 13:09 .
drwxr-xr-x 24 root root 4096 Dec  7 13:07 ..
drwxrwxrwx 2 root root 4096 Dec  7 13:09 mysql
```

Restart Percona Server.

```
$ service mysql start
```

### Expected output

```
Job for mysql.service failed because the control process exited with error code.
See "systemctl status mysql.service" and "journalctl -xe" for details.
```

```
$ journalctl -xe
```

**Expected output**

```
...
AVC apparmor="DENIED" operation="mknod" profile="/usr/sbin/mysqld" name="/custom-log-dir/mysql/
error.log"
...
```

The access has been denied by AppArmor. Edit the local profile in the `Allow log file access` section to allow access to the custom log location.

```
$ cat /etc/apparmor.d/local/usr.sbin.mysqld
```

**Expected output**

```
Site-specific additions and overrides for usr.sbin.mysqld..
For more details, please see /etc/apparmor.d/local/README.

Allow log file access
/custom-log-dir/mysql/ r,
/custom-log-dir/mysql/** rw,
```

Reload the profile:

```
$ apparmor_parser -r -T /etc/apparmor.d/usr.sbin.mysqld
```

Restart Percona Server:

```
$ systemctl restart mysqld
```

### 6.3.13 Set `secure_file_priv` directory location

By default, `secure_file_priv` points to the following location:

```
mysql> mysqlshow variables like 'secure_file_priv';
```

**Expected output**

```
+-----+-----+
| Variable_name | Value |
+-----+-----+
| secure_file_priv | /var/lib/mysql-files/ |
+-----+-----+
```

To allow access to another location, in a text editor, open the local profile. Review the settings in the `Allow data dir access` section:

```
Allow data dir access
/var/lib/mysql/ r,
/var/lib/mysql/** rwk,
```

Edit the local profile in a text editor to allow access to the custom location.

```
$ cat /etc/apparmor.d/local/usr.sbin.mysqld
```

#### Expected output

```
Site-specific additions and overrides for usr.sbin.mysqld..
For more details, please see /etc/apparmor.d/local/README.
```

```
Allow data dir access
/var/lib/mysqlcustom/ r,
/var/lib/mysqlcustom/** rwk,
```

Reload the profile:

```
$ apparmor_parser -r -T /etc/apparmor.d/usr.sbin.mysqld
```

Restart Percona Server for MySQL:

```
$ systemctl restart mysqld
```

## 6.3.14 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-11-23

## 6.4 Working with SELinux

The Linux kernel, through the Linux Security Module (LSM), supports Security-Enhanced Linux (SELinux). This module provides a way to support mandatory access control policies. SELinux defines how confined processes interact with files, network ports, directories, other processes, and additional server components.

An SELinux policy defines the set of rules, the `types` for files, and the `domains` for processes. Rules determine how a process interacts with another type. SELinux decides whether to allow or deny an action based on the subject's context, what object initiates the action and what object is the action's target.

A label represents the context for administrators and users.

CentOS 7 and CentOS 8 contain a MySQL SELinux policy. *Percona Server for MySQL* is a drop-in replacement for MySQL and can use this policy without changes.

### 6.4.1 SELinux context example

To view the SELinux context, add the `-Z` switch to many of the utilities. Here is an example of the context for `mysqld`:

```
$ ps -eZ | grep mysqld_t
```

#### Expected output

```
system_u:system_r:mysqld_t:s0 3356 ? 00:00:01 mysqld
```

The context has the following properties:

- User - `system_u`
- Role - `system_r`
- Type or domain - `mysqld_t`
- Sensitivity level - `s0 3356`

Most SELinux policy rules are based on the type or domain.

### 6.4.2 List SELinux types or domains associated with files

The security property that SELinux relies on is the Type security property. The type name often end with a `_t`. A group of objects with the same type security value belongs to the same domain.

To view the `mysqldb_t` types associated with the MySQL directories and files, run the following command:

```
$ ls -laZ /var/lib/ | grep mysql
```

#### Expected output

```
drwxr-x--x. mysql mysql system_u:object_r:mysqldb_t:s0 mysql
drwxr-x---. mysql mysql system_u:object_r:mysqldb_t:s0 mysql-files
drwxr-x---. mysql mysql system_u:object_r:mysqldb_t:s0 mysql-keyring
```

#### Note

If a policy type does not define the type property for an object, the default value is `unconfined_t`.

### 6.4.3 SELinux modes

SELinux has the following modes:

- Disabled – No SELinux policy modules loaded, which disables policies. Nothing is reported.
- Permissive – SELinux is active, but policy modules are not enforced. A policy violation is reported but does not stop the action.
- Enforcing – SELinux is active, and violations are reported and denied. If there is no rule to allow access to a confined resource, SELinux denies the access.

### 6.4.4 Policy types

SELinux has several policy types:

- Targeted – Most processes operate without restriction. Specific services are contained in security domains and defined by policies.
- Strict – All processes are contained in security domains and defined by policies.

SELinux has confined processes that run in a domain and restricts everything unless explicitly allowed. An unconfined process in an unconfined domain is allowed almost all access.

MySQL is a confined process, and the policy module defines which files are read, which ports are opened, and so on. SELinux assumes the *Percona Server for MySQL* installation uses the default file locations and default ports.

If you change the default, you must also edit the policy. If you do not update the policy, SELinux, in enforcing mode, denies access to all non-default resources.

### 6.4.5 Check the SELinux mode

To check the current SELinux mode, use either of the following commands:

```
$ sestatus
```

#### Expected output

```
SELinux status:                enabled
SELinuxfs mount:              /sys/fs/selinux
SELinux root directory:       /etc/selinux
Loaded policy name:            targeted
Current mode:                  enforcing
Mode from config file:         enforcing
Policy MLS status:             enabled
Policy deny_unknown status:    allowed
Memory protection checking:    actual (secure)
Max kernel policy version:     31
```

or

```
$ grep ^SELINUX= /etc/selinux/config
```

**Expected output**

```
SELINUX=enforcing
```

**Note**

Add the `-b` parameter to `sestatus` to display the Policy booleans. The boolean values for each parameter is shown. An example of using the `b` parameter is the following:

```
$ sestatus -b | grep mysql
```

**Expected output**

```
mysql_connect_any          off
selinuxuser_mysql_connect_enabled
```

The `/etc/selinux/config` file controls if SELinux is disabled or enabled, and if enabled, whether SELinux operates in enforcing mode or permissive mode.

## 6.4.6 Disable SELinux

If you plan to use the enforcing mode at another time, use the permissive mode instead of disabling SELinux. During the time that SELinux is disabled, the system may contain mislabeled objects or objects with no label. If you re-enable SELinux and plan to set SELinux to enforcing, you must follow the steps to Relabel the entire file system.

On boot, to disable SELinux, set the `selinux=0` kernel option. The kernel does not load the SELinux infrastructure. This option has the same effect as changing the `SELINUX=disabled` instruction in the configuration file and then rebooting the system.

## 6.4.7 Additional SELinux tools

Install the SELinux management tools, such as `semanage` or `sesearch`, if needed.

On RHEL 7 or compatible operating systems, use the following command as root:

```
$ yum -y install policycoreutils-python
```

On RHEL 8 or compatible operating systems, use the following command as root:

```
$ yum -y install policycoreutils-python-utils
```

**Note**

You may need root privileges to run SELinux management commands.

## 6.4.8 Switch the mode in the configuration file

Switching between modes may help when troubleshooting or when modifying rules.

To permanently change the mode, edit the `/etc/selinux/config` file and change the `SELINUX=` value. You should also verify the change.

```
$ cat /etc/selinux/config | grep SELINUX= | grep -v ^#
```

### Expected output

```
SELINUX=enforcing
SELINUX=enforcing
```

```
$ sudo sed -i 's/^SELINUX=.*SELINUX=permissive/g' /etc/selinux/config
```

```
$ cat /etc/selinux/config | grep SELINUX= | grep -v ^#
```

### Expected output

```
SELINUX=permissive
SELINUX=permissive
```

Reboot your system after the change.

If switching from either disabled mode or permissive mode to enforcing, see [Relabel the entire file system](#).

## 6.4.9 Switch the mode until the next reboot

To change the mode until the next reboot, use either of the following commands as root:

```
$ setenforce Enforcing
```

or

```
$ setenforce 1
```

The following `setenforce` parameters are available:

setenforce parameters	Also Permitted
0	Permissive
1	Enforcing

You can view the current mode by running either of the following commands:

```
$ getenforce
```

#### Expected output

```
Enforcing
```

or

```
$ sestatus | grep -i mode
```

#### Expected output

```
Current mode:                permissive
Mode from config file:       enforcing
```

## 6.4.10 Switch the mode for a service

You can move one or more services into a permissive domain. The other services remain in enforcing mode.

To add a service to the permissive domain, run the following as root:

```
$ sudo semanage permissive -a mysqld_t
```

To list the current permissive domains, run the following command:

```
$ sudo semanage permissive -l
```

#### Expected output

```
...
Customized Permissive Types

mysqld_t

Builtin Permissive Types
...
```

To delete a service from the permissive domain, run the following:

```
$ sudo semanage permissive -d mysqld_t
```

The service returns to the system's SELinux mode. Be sure to follow the steps to Relabel the entire file system.

## 6.4.11 Relabel the entire file system

Switching from disabled or permissive to enforcing requires additional steps. The enforcing mode requires the correct contexts, or labels, to function. The permissive mode allows users and processes to label files and system objects incorrectly. The disabled mode does not load the SELinux infrastructure and does not label resources or processes.

RHEL and compatible systems, use the `fixfiles` application for relabeling. You can relabel the entire file system or the file contexts of an application.

For one application, run the following command:

```
$ fixfiles -R mysqld restore
```

To relabel the file system without rebooting the system, use the following command:

```
$ fixfiles -f -F relabel
```

Another option relabels the file system during a reboot. You can either add a touch file, read during the reboot operation, or configure a kernel boot parameter. The completion of the relabeling operation automatically removes the touch file.

Add the touch file as root:

```
$ touch /.autorelabel
```

To configure the kernel, add the `autorelabel=1` kernel parameter to the boot parameter list. The parameter forces a system relabel. Reboot in permissive mode to allow the process to complete before changing to enforcing.

### Note

Relabeling an entire filesystem takes time. When the relabeling is complete, the system reboots again.

## 6.4.12 Set a custom data directory

If you do not use the default settings, SELinux, in enforcing mode, prevents access to the system.

For example, during installation, you have used the following configuration:

```
datadir=/var/lib/mysqlcustom  
socket=/var/lib/mysqlcustom/mysql.sock
```

Restart the service.

```
$ service mysqld restart
```

**Expected output**

```
Redirecting to /bin/systemctl restart mysqld.service
Job for mysqld.service failed because the control process exited with error code.
See "systemctl status mysqld.service" and "journalctl -xe" for details.
```

Check the journal log to see the error code.

```
$ journalctl -xe
```

**Expected output**

```
...
SELinux is preventing mysqld from getattr access to the file /var/lib/mysqlcustom/ibdata1.
...
```

Check the SELinux types in `/var/lib/mysqlcustom`.

```
ls -laZ /var/lib/mysqlcustom
```

**Expected output**

```
total 164288
drwxr-x--x. 6 mysql mysql system_u:object_r:var_lib_t:s0      4096 Dec  2 07:58 .
drwxr-xr-x. 38 root  root  system_u:object_r:var_lib_t:s0      4096 Dec  1 14:29 ..
...
-rw-r-----. 1 mysql mysql system_u:object_r:var_lib_t:s0 12582912 Dec  1 14:29 ibdata1
...
```

To solve the issue, use the following methods:

- Set the proper labels for `mysqlcustom` files
- Change the `mysqld` SELinux policy to allow `mysqld` access to `var_lib_t` files.

The recommended solution is to set the proper labels. The following procedure assumes you have already created and set ownership to the custom data directory location:

- 2.1 To change the SELinux context, use `semanage fcontext`. In this step, you define how SELinux deals with the custom paths:

```
$ semanage fcontext -a -e /var/lib/mysql /var/lib/mysqlcustom
```

SELinux applies the same labeling schema, defined in the `mysqld` policy, for the `/var/lib/mysql` directory to the custom directory. Files created within the custom directory are labeled as if they were in `/var/lib/mysql`.

- 2.2 `restorecon` command applies the change.

```
$ restorecon -R -v /var/lib/mysqlcustom
```

2. Restart the mysqld service:

```
$ service mysqld start
```

### 6.4.13 Set a custom log location

If you do not use the default settings, SELinux, in enforcing mode, prevents access to the location. Change the log location to a custom location in my.cnf:

```
log-error=/logs/mysqld.log
```

Verify the log location with the following command:

```
$ ls -laZ /
```

#### Expected output

```
...
drwxrwxrwx.  2 root root unconfined_u:object_r:default_t:s0  6 Dec  2 09:16 logs
...
```

Starting MySQL returns the following message:

```
$ service mysql start
```

#### Expected output

```
Redirecting to /bin/systemctl start mysql.service
Job for mysqld.service failed because the control process exited with error code.
See "systemctl status mysqld.service" and "journalctl -xe" for details.

$ journalctl -xe
...
SELinux is preventing mysqld from write access to the directory logs.
...
```

The default SELinux policy allows mysqld to write logs into a location tagged with `var_log_t`, which is the `/var/log` location. You can solve the issue with either of the following methods:

- Tag the `/logs` location properly
- Edit the SELinux policy to allow mysqld access to all directories.

To tag the custom `/logs` location is the recommended method since it locks down access. Run the following commands to tag the custom location:

```
$ semanage fcontext -a -t var_log_t /logs
$ restorecon -v /logs
```

You may not be able to change the `/logs` directory label. For example, other applications, with their own rules, use the same directory.

To adjust the SELinux policy when a directory is shared, follow these steps:

2. Create a local policy:

```
ausearch -c 'mysqld' --raw | audit2allow -M my-mysqld
```

2. This command generates the `my-mysqld.te` and the `my-mysqld.pp` files. The `mysqld.te` is the type enforcement policy file. The `my-mysqld.pp` is the policy module loaded as a binary file into the SELinux subsystem.

An example of the `my-mysqld.te` file:

```
module my-mysqld 1.0;

require {
    *type mysqld_t*;
    type var_lib_t;
    *type default_t*;
    class file getattr;
    *class dir write*;
}

===== mysqld_t =====
*allow mysqld_t default_t:dir write*;
allow mysqld_t var_lib_t:file getattr;
```

The policy contains rules for the custom data directory and the custom logs directory. We have set the proper labels for the data directory location, and applying this auto-generated policy would loosen our hardening by allowing `mysqld` to access `var_lib_t` tags.

2. SELinux-generated events are converted to rules. A generated policy may contain rules for recent violations and include unrelated rules. Unrelated rules are generated from actions, such as changing the data directory location, that are not related to the logs directory. Add the `--start` parameter to use log events after a specific time to filter out the unwanted events. This parameter captures events when the time stamp is equal to the specified time or later. SELinux generates a policy for the current actions.

```
$ ausearch --start 10:00:00 -c 'mysqld' --raw | audit2allow -M my-mysqld
```

2. This policy allows `mysqld` writing into the tagged directories. Open the `my_mysqld` file:

```
module my-mysqld 1.0;

require {
    type mysqld_t;
    type default_t;
    class dir write;
```

```

}

===== mysql_d_t =====
allow mysql_d_t default_t:dir write;

```

- 2.5 Install the SELinux policy module:

```
$ semodule -i my-mysqld.pp
```

Restart the service. If you have a failure, check the journal log and follow the same procedure.

If SELinux prevents mysql from creating a log file inside the directory. You can view all the violations by changing the SELinux mode to `permissive` and then running `mysqld`. All violations are logged in the journal log. After this run, you can generate a local policy module, install it, and switch SELinux back to `enforcing` mode.

Follow this procedure:

- 2.1 Unload the current local my-mysqld policy module:

```
$ semodule -r my-mysqld
```

- 2.2 You can put a single domain into permissive mode. Other domains on the system to remain in enforcing mode. Use `semanage permissive` with the `-a` parameter to change `mysql_d_t` to permissive mode:

```
$ semanage permissive -a mysql_d_t
```

- 2.3 Verify the mode change:

```
$ semodule -l | grep permissive
```

#### Expected output

```

...
permissive_mysql_d_t
...

```

- 2.4 To make searching the log easier, return the time:

```
$ date
```

- 2.5 Start the service.

```
$ service mysqld start
```

- 2.6 MySQL starts, and SELinux logs the violations in the journal log. Check the journal log:

```
$ journalctl -xe
```

2. 7 Stop the service:

```
$ service mysqld stop
```

2. 8 Generate a local mysqld policy, using the time returned from step 4:

```
$ ausearch --start <date-c 'mysqld' --raw | audit2allow -M my-mysqld
```

2. 9 Review the policy (the policy you generate may be different):

```
$ cat my-mysqld.te
```

#### Expected output

```
module my-mysqld 1.0;

require {
  type default_t;
  type mysqld_t;
  class dir { add_name write };
  class file { append create open };
}

===== mysqld_t =====
allow mysqld_t default_t:dir { add_name write };
allow mysqld_t default_t:file { append create open };
```

2. 10 Install the policy:

```
$ semodule -i my-mysqld.pp
```

2. 11 Use `semanage permissive` with the `-d` parameter, which deletes the permissive domain for the service:

```
$ semanage permissive -d mysqld_t
```

2. 12 Restart the service:

```
$ service mysqld start
```

#### Note

Use this procedure to adjust the local mysqld policy module. You should review the changes which are generated to ensure the rules are not too tolerant.

## 6.4.14 Set secure\_file\_priv directory

Update the SELinux tags for the `/var/lib/mysql-files/` directory, used for `SELECT ... INTO OUTFILE` or similar operations, if required. The server needs only read/write access to the destination directory.

To set `secure_file_priv` to use this directory, run the following commands to set the context:

```
$ semanage fcontext -a -t mysqld_db_t "/var/lib/mysql-files/(.*)?"
$ restorecon -Rv /var/lib/mysql-files
```

Edit the path for a different location, if needed.

## 6.4.15 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-04

## 7. Manage

### 7.1 Extended SHOW GRANTS

In Oracle MySQL `SHOW GRANTS` displays only the privileges granted explicitly to the named account. Other privileges might be available to the account, but they are not displayed. For example, if an anonymous account exists, the named account might be able to use its privileges, but `SHOW GRANTS` will not display them. *Percona Server for MySQL* offers the `SHOW EFFECTIVE GRANTS` command to display all the effectively available privileges to the account, including those granted to a different account.

#### 7.1.1 Example

If we create the following users:

```
mysql> CREATE USER grantee@localhost IDENTIFIED BY 'grantee1';
```

##### Expected output

```
Query OK, 0 rows affected (0.50 sec)
```

```
mysql> CREATE USER grantee IDENTIFIED BY 'grantee2';
```

##### Expected output

```
Query OK, 0 rows affected (0.09 sec)
```

```
mysql> CREATE DATABASE db2;
```

##### Expected output

```
Query OK, 1 row affected (0.20 sec)
```

```
mysql> GRANT ALL PRIVILEGES ON db2.* TO grantee WITH GRANT OPTION;
```

##### Expected output

```
Query OK, 0 rows affected (0.12 sec)
```

- `SHOW EFFECTIVE GRANTS` output before the change:

```
mysql> SHOW EFFECTIVE GRANTS;
```

### Expected output

```
+-----+
+
| Grants for
| grantee@localhost
|
+-----+
+
| GRANT USAGE ON *.* TO 'grantee'@'localhost' IDENTIFIED BY PASSWORD
| '*9823FF338D44DAF02422CF24DD1F879FB4F6B232' |
+-----+
+
1 row in set (0.04 sec)
```

Although the grant for the `db2` database isn't shown, `grantee` user has enough privileges to create the table in that database:

```
user@trusty:~$ mysql -ugrantee -pgrantee1 -h localhost
```

```
mysql> CREATE TABLE db2.t1(a int);
```

### Expected output

```
Query OK, 0 rows affected (1.21 sec)
```

- The output of `SHOW EFFECTIVE GRANTS` after the change shows all the privileges for the `grantee` user:

```
mysql> SHOW EFFECTIVE GRANTS;
```

### Expected output

```
+-----+
| Grants for grantee@localhost
|
+-----+
| GRANT USAGE ON *.* TO 'grantee'@'localhost' IDENTIFIED BY PASSWORD|
| '*9823FF338D44DAF02422CF24DD1F879FB4F6B232' |
| GRANT ALL PRIVILEGES ON `db2`.* TO 'grantee'@'%' WITH GRANT OPTION|
+-----+
2 rows in set (0.00 sec)
```

## 7.1.2 Other reading

- [#53645](#) - `SHOW GRANTS` not displaying all the applicable grants

### 7.1.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 7.2 Restrict dynamic log file locations

The `secure_log_path` system variable restricts the dynamic log file locations.

### 7.2.1 `secure_log_path`

The variable is read-only and must be set up in a configuration file or the command line.

The accepted value is the directory name as a string. The default value is an empty string. When the value is an empty string, the variable only adds a warning to the error log and does nothing. If the value contains a directory name, then the slow query log and the general log must be located in that directory. An attempt to move either of these files outside of the specified directory results in an error.

### 7.2.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-23

## 7.3 MySQL Clone plugin

The MySQL Clone plugin lets you clone data from either a local server or from a remote server. The plugin creates a physical snapshot of the data stored in InnoDB, which includes schemas, tables, tablespaces, and data dictionary metadata. The cloned data is a functional data directory and can be used for provisioning a server .

The following table lists the cloning operation types:

Cloning operation type	Description
Local	Clones the data from the server when the data directory is located in the same directory on the same server or a server

Cloning operation type	Description
Remote	Clones the data from the donor to th

When replicating a large number of transactions, the Clone plugin may be a more efficient solution.

### 7.3.1 Install the Clone plugin

The Clone plugin must be installed on both the donor and the joiner servers at either server startup or at runtime. To install the plugin at runtime, run the following command:

```
mysql> INSTALL PLUGIN clone SONAME 'mysql_clone.so';
```

Review the INFORMATION\_SCHEMA.PLUGINS table or run the `SHOW PLUGINS` command to verify the installation. The following is an example of querying the PLUGINS table.

```
mysql> SELECT PLUGIN_NAME, PLUGIN_STATUS FROM INFORMATION_SCHEMA.PLUGINS WHERE
PLUGIN_NAME='clone';
```

The result lists the Clone plugin and the status.

### 7.3.2 Clone data

The SQL statement used to clone data depends on if the operation is local or remote. The following code is an example of cloning data from a remote server:

```
mysql> CLONE INSTANCE FROM `root@remote.server:13336` IDENTIFIED BY `user`;
```

Replace the user name, host name, and port number with the settings from the donor server.

### 7.3.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 7.4 Manage components

### 7.4.1 INSTALL COMPONENT

The `INSTALL COMPONENT` does the following:

- Installs the component
- Activates the component

If an error, such as a misspelled component name, occurs, the statement fails and nothing happens.

You can install multiple components at the same time.

#### Example

The following is an example of the `INSTALL COMPONENT` statement.

```
mysql> INSTALL COMPONENT 'file://componentA';
```

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

### 7.4.2 UNINSTALL COMPONENT

The `UNINSTALL COMPONENT` does the following:

- Deactivates the component
- Uninstalls the component

If the statement does not undo any persisted variables.

If an error, such as a misspelled component name, occurs, the statement fails and nothing happens.

You can uninstall multiple components at the same time.

#### Required privilege

The statement requires the `DELETE` privilege for the `mysql.component` system table. Executing the statement removes the registration row from this table.

## Example

The following is an example of the `UNINSTALL COMPONENT` statement.

```
mysql > UNINSTALL COMPONENT 'file://componentA' ;
```

Find more information in the [UNINSTALL COMPONENT](#) document.

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-11-27

## 8. Back up and restore

### 8.1 Backup and restore overview

Backups are data snapshots that are taken at a specific time and are stored in a common location in a common format. A backup is only useful for a defined time.

The following scenarios require a backup to recover:

Reason	Description
Hardware or host failure	Issues with disks, such as stalls or broken instances can be inaccessible or broken.
Corrupted data	This issue can be caused by power outages or not saving files correctly and close the file.
User mistake	Deleting data or an update overwriting data.
Natural disaster or data center failure	Power outage, flooding, or internet issues.
Compliance	Required to comply with regulations.

#### 8.1.1 Strategies

Define a backup and restore strategy for each of your databases. The strategies should have the following practices:

Practice	Description
Retention	How long should you keep the backups? This is determined by your organization's data governance policies and the retention policy for backups.
Document	Document the strategy and any related procedures. This should include information about the process of backup and restore.
Encrypt	Encrypt the backup and secure the storage.
Test	Test the backups on a timely basis.

The backup strategy defines type and the backup frequency, the hardware required, how the backups are verified, and storing the backups, which also includes the backup security. The strategy uses the following metrics:

Metric	Description
Recovery Time Objective (RTO)	How long can the system be down?
Recovery Point Objective (RPO)	How much data can the organization afford to lose?

The restore strategy defines which user account has the restore responsibility and how and frequency of testing the restore process.

These strategies require planning, implementation, and rigorous testing. You must test your restore process with each type of backup used to validate the backup and measure the recovery time. Automate this testing as much as possible. You should also document the process. In case of disaster, you can follow the procedures in the document without wasting time.

If you are using replication, consider using a dedicated replica for backups because the operation can cause a high CPU load.

## 8.1.2 Physical backup or logical backup

A backup can be either a physical backup or a logical backup.

### Physical backups

A physical backup copies the files needed to store and recover the database. They can be data files, configuration files, logs, and other types of files. The physical database can be stored in the cloud, in offline storage, on disc, or tape.

[Percona XtraBackup](#) takes a physical backup. You can also use RDS/LVM Snapshots or the MySQL Enterprise Backup.

If the server is stopped or down, you can copy the datadir with the `cp` command or the `rsync` command.

### Logical backups

A logical backup contains the structural details. This type of backup contains tables, views, procedures, and functions.

Tools like `[ mysqldump ]`, `[ mydumper ]`, `[ mysqlpump ]`, and `[ mysql shell ]` take a logical backup.

### Comparison

Comparison	Physical backup	Logical backup
Content	The physical database files	The table metadata
Restore speed	Restore can be quick	Restore is slower
Storage	Can take more space	Based on the size of the database, it can be smaller

## 8.1.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-11-23

## 8.2 Backup locks

Percona Server for MySQL offers the `LOCK TABLES FOR BACKUP` statement as a lightweight alternative to `FLUSH TABLES WITH READ LOCK` for both physical and logical backups.

### 8.2.1 LOCK TABLES FOR BACKUP

`LOCK TABLES FOR BACKUP` uses a new MDL lock type to block updates to non-transactional tables and DDL statements for all tables. If there is an active `LOCK TABLES FOR BACKUP` lock then all DDL statements and all updates to MyISAM, CSV, MEMORY, ARCHIVE, and MyRocks tables will be blocked in the `Waiting for backup lock` status, visible in `PERFORMANCE_SCHEMA` or `PROCESSLIST`.

`LOCK TABLES FOR BACKUP` has no effect on `SELECT` queries for all mentioned storage engines. Against *InnoDB*, *MyRocks*, *Blackhole* and *Federated* tables, the `LOCK TABLES FOR BACKUP` is not applicable to the `INSERT`, `REPLACE`, `UPDATE`, `DELETE` statements: *Blackhole* tables obviously have no relevance to backups, and *Federated* tables are ignored by both logical and physical backup tools.

Unlike `FLUSH TABLES WITH READ LOCK`, `LOCK TABLES FOR BACKUP` does not flush tables, i.e. storage engines are not forced to close tables and tables are not expelled from the table cache. As a result, `LOCK TABLES FOR BACKUP` only waits for conflicting statements to complete (i.e. DDL and updates to non-transactional tables). It never waits for `SELECT`s, or `UPDATE`s to *InnoDB* or *MyRocks* tables to complete, for example.

If an “unsafe” statement is executed in the same connection that is holding a `LOCK TABLES FOR BACKUP` lock, it fails with the following error:

```

Expected output

ERROR 1880 (HY000): Can't execute the query because you have a conflicting backup lock

UNLOCK TABLES releases the lock acquired by LOCK TABLES FOR BACKUP.
```

The intended use case for *Percona XtraBackup* is:

```

LOCK TABLES FOR BACKUP
... copy .frm, MyISAM, CSV, etc. ...
UNLOCK TABLES
... get binlog coordinates ...
... wait for redo log copying to finish ...
```

### 8.2.2 Privileges

The `LOCK TABLES FOR BACKUP` requires the `BACKUP_ADMIN` privilege.

### 8.2.3 Interaction with other global locks

The `LOCK TABLES FOR BACKUP` has no effect if the current connection already owns a `FLUSH TABLES WITH READ LOCK` lock, as it is a more restrictive lock. If `FLUSH TABLES WITH READ LOCK` is executed in a

connection that has acquired `LOCK TABLES FOR BACKUP`, `FLUSH TABLES WITH READ LOCK` fails with an error.

If the server is operating in the read-only mode (i.e. `read_only` set to `1`), statements that are unsafe for backups will be either blocked or fail with an error, depending on whether they are executed in the same connection that owns `LOCK TABLES FOR BACKUP` lock, or other connections.

## 8.2.4 MyISAM index and data buffering

MyISAM key buffering is normally write-through, i.e. by the time each update to a MyISAM table is completed, all index updates are written to disk. The only exception is delayed key writing feature which is disabled by default.

When the global system variable `delay_key_write` is set to `ALL`, key buffers for all MyISAM tables are not flushed between updates, so a physical backup of those tables may result in broken MyISAM indexes. To prevent this, `LOCK TABLES FOR BACKUP` will fail with an error if `delay_key_write` is set to `ALL`. An attempt to set `delay_key_write` to `ALL` when there's an active backup lock will also fail with an error.

Another option to involve delayed key writing is to create MyISAM tables with the `DELAY_KEY_WRITE` option and set the `delay_key_write` variable to `ON` (which is the default). In this case, `LOCK TABLES FOR BACKUP` will not be able to prevent stale index files from appearing in the backup. Users are encouraged to set `delay_key_writes` to `OFF` in the configuration file, `my.cnf`, or repair MyISAM indexes after restoring from a physical backup created with backup locks.

MyISAM may also cache data for bulk inserts, e.g. when executing multi-row `INSERTs` or `LOAD DATA` statements. Those caches, however, are flushed between statements, so have no effect on physical backups as long as all statements updating MyISAM tables are blocked.

## 8.2.5 The mysqldump Command

`mysqldump` has also been extended with a new option, `lock-for-backup` (disabled by default). When used together with the `--single-transaction` option, the option makes `mysqldump` issue `LOCK TABLES FOR BACKUP` before starting the dump operation to prevent unsafe statements that would normally result in an inconsistent backup.

When used without the `--single-transaction` option, `lock-for-backup` is automatically converted to `lock-all-tables`.

The option `lock-for-backup` is mutually exclusive with `lock-all-tables`, i.e. specifying both on the command line will lead to an error.

If the backup locks feature is not supported by the target server, but `lock-for-backup` is specified on the command line, `mysqldump` aborts with an error.

## 8.2.6 System Variables

### have\_backup\_locks

Option	Description
Command Line:	Yes
Config file	No

Option	Description
Scope:	Global
Dynamic:	No
Data type	Boolean
Default value	YES

This is a server variable implemented to help other utilities decide what locking strategy can be implemented for a server. When available, the backup locks feature is supported by the server and the variable value is always `YES`.

## 8.2.7 Status variables

### `Com_lock_tables_for_backup`

Option	Description
Scope:	Global/Session
Data type	Numeric

This status variable indicates the number of times the corresponding statements have been executed.

## 8.2.8 Client command line parameter

### `Lock-for-backup`

Option	Description
Command Line:	Yes
Scope:	Global
Dynamic:	No
Data type	String
Default value	Off

When used together with the `--single-transaction` option, the option makes `mysqldump` issue `LOCK TABLES FOR BACKUP` before starting the dump operation to prevent unsafe statements that would normally result in an inconsistent backup.

## 8.2.9 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 8.3 Extended mysqldump

### 8.3.1 Backup locks support

When used together with the `--single-transaction` option, the `lock-for-backup` option makes `mysqldump` issue `LOCK TABLES FOR BACKUP` before starting the dump operation to prevent unsafe statements that would normally result in an inconsistent backup.

More information can be found on the Backup Locks feature documentation.

### 8.3.2 Compressed columns support

`mysqldump` supports the Compressed columns with dictionaries feature. More information about the relevant options can be found on the Compressed columns with dictionaries feature page.

### 8.3.3 Taking backup by descending primary key order

`--order-by-primary-desc` tells `mysqldump` to take the backup by descending primary key order (`PRIMARY KEY DESC`) which can be useful if the storage engine is using the reverse order column for a primary key.

### 8.3.4 RocksDB support

`mysqldump` detects when MyRocks is installed and available. If there is a session variable named `rocksdb_skip_fill_cache` `mysqldump` sets it to `1`.

`mysqldump` will now automatically enable session the variable `rocksdb_bulk_load` if it is supported by the target server.

### 8.3.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-09-25

## 8.4 Start transaction with consistent snapshot

Percona Server for MySQL has ported *MariaDB enhancement* for `START TRANSACTION WITH CONSISTENT SNAPSHOT` feature to MySQL 5.6 group commit implementation. This enhancement makes binary log positions consistent with *InnoDB* transaction snapshots.

This feature is quite useful to obtain logical backups with correct positions without running a `FLUSH TABLES WITH READ LOCK`. Binary log position can be obtained by two newly implemented status variables: `Binlog_snapshot_file` and `Binlog_snapshot_position`. After starting a transaction using the `START TRANSACTION WITH CONSISTENT SNAPSHOT`, these two variables will provide you with the binlog position corresponding to the state of the database of the consistent snapshot so taken, irrespectively of which other transactions have been committed since the snapshot was taken.

## 8.4.1 Snapshot cloning

The *Percona Server for MySQL* implementation extends the `START TRANSACTION WITH CONSISTENT SNAPSHOT` syntax with the optional `FROM SESSION` clause:

```
START TRANSACTION WITH CONSISTENT SNAPSHOT FROM SESSION <session_id>;
```

When specified, all participating storage engines and binary log instead of creating a new snapshot of data (or binary log coordinates), create a copy of the snapshot which has been created by an active transaction in the specified session. `session_id` is the session identifier reported in the `Id` column of `SHOW PROCESSLIST`.

Currently snapshot cloning is only supported by *XtraDB* and the binary log. As with the regular `START TRANSACTION WITH CONSISTENT SNAPSHOT`, snapshot clones can only be created with the `REPEATABLE READ` isolation level.

For *XtraDB*, a transaction with a cloned snapshot will only see data visible or changed by the donor transaction. That is, the cloned transaction will see no changes committed by transactions that started after the donor transaction, not even changes made by itself. Note that in case of chained cloning the donor transaction is the first one in the chain. For example, if transaction A is cloned into transaction B, which is in turn cloned into transaction C, the latter will have read view from transaction A (i.e., the donor transaction). Therefore, it will see changes made by transaction A, but not by transaction B.

## 8.4.2 mysqldump

`mysqldump` has been updated to use new status variables automatically when they are supported by the server and both `--single-transaction` and `--master-data` are specified on the command line. Along with the `mysqldump` improvements introduced in [Backup Locks](#) there is now a way to generate `mysqldump` backups that are guaranteed to be consistent without using `FLUSH TABLES WITH READ LOCK` even if `--master-data` is requested.

## 8.4.3 System variables

### have\_snapshot\_cloning

Option	Description
Command Line:	Yes
Config file	No
Scope:	Global
Dynamic:	No

Option	Description
Data type	Boolean

This server variable is implemented to help other utilities detect if the server supports the `FROM SESSION` extension. When available, the snapshot cloning feature and the syntax extension to `START TRANSACTION WITH CONSISTENT SNAPSHOT` are supported by the server, and the variable value is always `YES`.

## 8.4.4 Status variables

### `Binlog_snapshot_file`

Option	Description
Scope:	Global
Data type	String

### `Binlog_snapshot_position`

Option	Description
Scope:	Global
Data type	Numeric

These status variables are only available when the binary log is enabled globally.

## 8.4.5 Other reading

- [MariaDB Enhancements for START TRANSACTION WITH CONSISTENT SNAPSHOT](#)

### 8.4.6 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 9. Secure

### 9.1 LDAP features

#### 9.1.1 Using LDAP authentication plugins

LDAP (Lightweight Directory Access Protocol) provides an alternative method to access existing directory servers, which maintain information about individuals, groups, and organizations.

Percona Server for MySQL supports the simple LDAP authentication. The Percona simple LDAP authentication plugin is a free and Open Source implementation of the MySQL Enterprise Simple LDAP authentication plugin. Percona Server for MySQL also supports an SASL-based LDAP authentication plugin. This plugin only supports the SCRAM-SHA-1 SASL mechanism.

#### Plugin names and file names

The following tables show the plugin names and the file name for simple LDAP authentication and SASL-based LDAP authentication.

Simple LDAP authentication plugin names and library files and SASL-based LDAP authentication plugin names and li

Plugin or file	Plugin name or file name
Server-side plugin	authentication_ldap_simple
client-side plugin	mysql_clear_password
library file	authentication_ldap_simple.so
Plugin or file	Plugin name or file name
Server-side plugin	authentication_ldap_sasl
client-side plugin	authentication_ldap_sasl_client
library files	authentication_ldap_sasl.so authentication_ldap_sasl_client.so

#### How does the authentication work

The server-side LDAP plugins work only with the specific client-side plugin:

- The `authentication_ldap_simple` plugin, on the server, performs the simple LDAP authentication. The client, using `mysql_clear_password`, connects to the server. The client plugin sends the password to the server as cleartext. For this method, use a secure connection between the client and server.
- The `authentication_ldap_sasl` plugin, on the server, performs the SASL-based LDAP authentication. The client must use the `authentication_ldap_sasl_client` plugin. The method does not send the password to the server in cleartext. The server-side and client-side plugins use Simple Authentication and Security Layer (SASL) to send secure messages within the LDAP protocol.

For either method, the database server rejects the connection if the client user name and the host name do not match a server account.

If a database server LDAP authentication is successful, the LDAP server searches for an entry. The LDAP server matches the user and authenticates using the LDAP password. If the database server account names the LDAP user distinguished name (DN), added by the `IDENTIFIED WITH <plugin-name> BY '<auth-string>'` clause, the LDAP server uses that value and the LDAP password provided by the client. This method fails if the DN and password have incorrect values

If the LDAP server finds multiple matches or no match, authentication fails.

If the password is correct, and the LDAP server finds a match, then LDAP authentication succeeds. The LDAP server returns the LDAP entry and the authentication plugin determines the authenticated user's name based on the entry. If the LDAP entry has no group attribute, the plugin returns the client user name as the authenticated name. If the LDAP entry has a group attribute, the plugin returns the group value as the authenticated name.

The database server compares the client user name to the authenticated user name. If these names are the same, the database server uses the client user name to check for privileges. If the name differs, then the database server looks for an account that matches the authenticated name.

### Prerequisites for authentication

The LDAP authentication plugins required the following:

- An available LDAP server
- The LDAP server must contain the LDAP user accounts to be authenticated
- The OpenLDAP client library must be available on the same system as the plugin

The SASL-based LDAP authentication additionally requires the following:

- Configure the LDAP server to communicate with a SASL server
- Available SASL client library on the same system as the client plugin.
- Services are configured to use the supported SCRAM-SHA-1 SASL mechanism

### Install the plugins

You can use either of the following methods to install the plugins.

## LOAD THE PLUGINS AT SERVER START

Use either of the following methods to load the plugin at server start.

**Load the simple LDAP authentication**      **Load the SASL-based LDAP authentication plugin**

Add the following statements to your `my.cnf` file to load simple LDAP authentication:

```
[mysqld]
plugin-load-add=authentication_ldap_simple.so
authentication_ldap_simple_server_host=127.0.0.1
authentication_ldap_simple_bind_base_dn='dc=percona, dc=com'
```

Restart the server for the changes to take effect.

Add the following statements to your `my.cnf` file to load the SASL-based LDAP authentication:

```
[mysqld]
plugin-load-add=authentication_ldap_sasl.so
authentication_ldap_sasl_server_host=127.0.0.1
authentication_ldap_sasl_bind_base_dn='dc=percona, dc=com'
```

## LOAD THE PLUGINS AT RUNTIME

Install the plugin with the following statements.

**Load the simple LDAP authentication plugin**      **Load the SASL-based LDAP authentication plugin**

```
mysql> INSTALL PLUGIN authentication_ldap_simple SONAME
'authentication_ldap_simple.so';
```

To set and persist values at runtime, use the following statements:

```
mysql> SET PERSIST authentication_ldap_simple_server_host='127.0.0.1';
mysql> SET PERSIST authentication_ldap_simple_bind_base_dn='dc=percona, dc=com';
```

```
mysql> INSTALL PLUGIN authentication_ldap_sasl SONAME 'authentication_ldap_sasl.so';
```

To set and persist values at runtime, use the following statements:

```
mysql> SET PERSIST authentication_ldap_sasl_server_host='127.0.0.1';
mysql> SET PERSIST authentication_ldap_sasl_bind_base_dn='dc=percona, dc=com';
```

## Create a user using simple LDAP authentication

There are several methods to add or modify a user.

**Use authentication\_ldap\_simple plugin**    **Use the authentication string in simple LDAP**

In the `CREATE USER` statement or the `ALTER USER` statement, for simple LDAP authentication, you can specify the `authentication_ldap_simple` plugin in the `IDENTIFIED WITH` clause:

```
mysql> CREATE USER ... IDENTIFIED WITH authentication_ldap_simple;
```

Using the `IDENTIFIED WITH` clause, the database server assigns the specified plugin.

If you provide the optional authentication string clause, 'cn,ou,dc,dc' in the example, the string is stored along with the password.

```
mysql> CREATE USER ... IDENTIFIED WITH authentication_ldap_simple BY 'cn=[user
name],ou=[organization unit],dc=[domain component],dc=com'
```

Unless the `authentication_ldap_simple_group_role_mapping` variable is used, creating a user with an authentication string does not use the following system variables:

- `authentication_ldap_simple_bind_base_dn`
- `authentication_ldap_simple_bind_root_dn`
- `authentication_ldap_simple_bind_root_pwd`
- `authentication_ldap_simple_user_search_attr`
- `authentication_ldap_simple_group_search_attr`

Creating the user with `IDENTIFIED BY authentication_ldap_simple` uses the variables.

Creating the user with the `authentication_ldap_simple_group_role_mapping` variable also adds the `authentication_ldap_simple_bind_root_dn` and `authentication_ldap_simple_bind_root_pwd` variables.

## Create a user using SASL-based LDAP authentication

There are several methods to add or modify a user.

[Use authentication\\_ldap\\_sasl plugin](#)    [Use the authentication string in SASL-based LDAP](#)

For SASL-based LDAP authentication, in the `CREATE USER` statement or the `ALTER USER` statement, you can specify the `authentication_ldap_sasl` plugin:

```
mysql> CREATE USER ... IDENTIFIED WITH authentication_ldap_sasl;
```

If you provide the optional authentication string clause, 'cn,ou,dc,dc' in the example, the string is stored along with the password.

```
mysql> CREATE USER ... IDENTIFIED WITH authentication_ldap_sasl BY 'cn=[user
name],ou=[organization unit],dc=[domain component],dc=com'
```

Unless the `authentication_ldap_sasl_group_role_mapping` variable is used, creating a user with an authentication string does not use the following system variables:

- `authentication_ldap_sasl_bind_base_dn`
- `authentication_ldap_sasl_bind_root_dn`
- `authentication_ldap_sasl_bind_root_pwd`
- `authentication_ldap_sasl_user_search_attr`
- `authentication_ldap_sasl_group_search_attr`

Creating the user with `IDENTIFIED BY authentication_ldap_sasl` uses the variables.

Creating the user with the `authentication_ldap_sasl_group_role_mapping` variable also adds the `authentication_ldap_sasl_bind_root_dn` and `authentication_ldap_sasl_bind_root_pwd` variables.

## Examples

The following sections are examples of using simple LDAP authentication and SASL-based LDAP authentication.

For the purposes of this example, we use the following LDAP user:

```
uid=ldapuser,ou=testusers,dc=percona,dc=com
```

### Simple LDAP authentication      SASL-based LDAP authentication

The following example configures an LDAP user and connects to the database server.

Create a database server account for `ldapuser` with the following statement:

```
mysql> CREATE USER 'ldapuser'@'localhost' IDENTIFIED WITH authentication_ldap_simple
BY 'uid=ldapuser,ou=testusers,dc=percona,dc=com';
```

The authentication string does not include the LDAP password. This password must be provided by the client user when they connect.

```
mysql> mysql --user=ldapuser --password --enable-cleartext-plugin
```

The user enters the `ldapuser` password. The client sends the password as cleartext, which is necessary when using a server-side LDAP library without SASL. The following actions may minimize the risk:

- Require that the database server clients explicitly enable the `mysql_clear_password` plugin with `--enable-cleartext-plugin`.
- Require that the database server clients connect to the database server using an encrypted connection

The following example configures an LDAP user and connect to the database server.

Create a database server account for `ldapuser` with the following statement:

```
mysql> CREATE USER 'ldapuser'@'localhost' IDENTIFIED WITH authentication_ldap_sasl
AS 'uid=ldapuser,ou=testusers,dc=percona,dc=com';
```

The authentication string does not include the LDAP password. This password must be provided by the client user when they connect.

Clients connect to the database server by providing the database server user name and LDAP password:

```
mysql> mysql --user=ldapuser --password
```

The authentication is similar to the authentication method used by simple LDAP authentication, except that the client and the database server SASL LDAP plugins use SASL messages. These messages are secure within the LDAP protocol.

**UNINSTALL THE PLUGINS**

If you installed either plugin at [server startup](#), remove those options from the `my.cnf` file, remove any startup options that set LDAP system variables, and restart the server.

**Uninstall the simple LDAP authentication plugin**    **Uninstall the SASL-based LDAP authentication plugin**

If you installed the plugins at [runtime](#), run the following statements:

```
mysql> UNINSTALL PLUGIN authentication_ldap_simple;
```

If you used `SET_PERSIST`, use `RESET PERSIST` to remove the settings.

If you installed the plugins at [runtime](#), run the following statements:

```
mysql> UNINSTALL PLUGIN authentication_ldap_sasl;
```

If you used `SET_PERSIST`, use `RESET PERSIST` to remove the settings.

**Get expert help**

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-31

## 9.1.2 LDAP authentication plugin system variables

**Authentication system variables**

The installation adds the following variables:

Variable name	Description
<a href="#">authentication_ldap_sasl_bind_base_dn</a>	Base distinguished name
<a href="#">authentication_ldap_sasl_bind_root_dn</a>	Root distinguished name
<a href="#">authentication_ldap_sasl_bind_root_dn_pwd</a>	Password for the root distinguished name
<a href="#">authentication_ldap_sasl_ca_path</a>	Absolute path of the certificate authority
<a href="#">authentication_ldap_sasl_fallback_server_host</a>	If the primary server is unavailable, the host to connect to the fallback server
<a href="#">authentication_ldap_sasl_fallback_server_port</a>	The port number for the fallback server
<a href="#">authentication_ldap_sasl_group_role_mapping</a>	A list of LDAP group names - MySQL roles
<a href="#">authentication_ldap_sasl_group_search_attr</a>	

Variable name	Description
	Name of the attribute that specifies the search filter for user entries
<code>authentication_ldap_sasl_group_search_filter</code>	Custom group search filter
<code>authentication_ldap_sasl_init_pool_size</code>	Initial size of the connection pool to the LDAP server
<code>authentication_ldap_sasl_log_status</code>	logging level
<code>authentication_ldap_sasl_max_pool_size</code>	Maximum size of the pool of connections to the LDAP server
<code>authentication_ldap_sasl_server_host</code>	LDAP server host
<code>authentication_ldap_sasl_server_port</code>	LDAP server TCP/IP port number
<code>authentication_ldap_sasl_ssl</code>	If plugin connections to the LDAP server use SSL
<code>authentication_ldap_sasl_tls</code>	If plugin connections to the LDAP server use TLS
<code>authentication_ldap_sasl_user_search_attr</code>	Name of the attribute that specifies the search filter for user entries
<code>authentication_ldap_simple_bind_base_dn</code>	Base distinguished name
<code>authentication_ldap_simple_bind_root_dn</code>	Root distinguished name
<code>authentication_ldap_simple_bind_root_dn_pwd</code>	Password for the root distinguished name
<code>authentication_ldap_simple_ca_path</code>	Absolute path of the certificate authority
<code>authentication_ldap_simple_fallback_server_host</code>	If the primary server is unavailable, the plugin will attempt to connect to the fallback server
<code>authentication_ldap_simple_fallback_server_port</code>	The port number for the fallback server
<code>authentication_ldap_simple_group_role_mapping</code>	A list of LDAP group names - MySQL roles
<code>authentication_ldap_simple_group_search_attr</code>	Name of the attribute that specifies the search filter for group entries
<code>authentication_ldap_simple_group_search_filter</code>	Custom group search filter
<code>authentication_ldap_simple_init_pool_size</code>	Initial size of the connection pool to the LDAP server
<code>authentication_ldap_simple_log_status</code>	logging level
<code>authentication_ldap_simple_max_pool_size</code>	Maximum size of the pool of connections to the LDAP server
<code>authentication_ldap_simple_server_host</code>	LDAP server host
<code>authentication_ldap_simple_server_port</code>	LDAP server TCP/IP port number
<code>authentication_ldap_simple_ssl</code>	If plugin connections to the LDAP server use SSL
<code>authentication_ldap_simple_tls</code>	If plugin connections to the LDAP server use TLS
<code>authentication_ldap_simple_user_search_attr</code>	Name of the attribute that specifies the search filter for user entries

The following variables are described in detail:

#### AUTHENTICATION\_LDAP\_SASL\_BIND\_BASE\_DN

Option	Description
Command-line	<code>-authentication-ldap-sasl-bind-base-dn</code>

Option	Description
Scope	Global
Dynamic	Yes
Data type	String
Default	NULL

The base distinguished name (DN) for SASL-based LDAP authentication. You can limit the search scope by using the variable as the base of the search.

#### AUTHENTICATION\_LDAP\_SASL\_BIND\_ROOT\_DN

Option	Description
Command-line	-authentication-ldap-sasl-bind-root-dn
Scope	Global
Dynamic	Yes
Data type	String
Default	NULL

The `root` distinguished name (DN) used to authenticate SASL-based LDAP. When performing a search, this variable is used with `authentication_ldap_sasl_bind_root_pwd` as the authenticating credentials to the LDAP server.

#### AUTHENTICATION\_LDAP\_SASL\_BIND\_ROOT\_PWD

Option	Description
Command-line	-authentication-ldap-sasl-bind-root-pwd
Scope	Global
Dynamic	Yes
Data type	String
Default	NULL

The `root` password used to authenticate against SASL-based LDAP server. This variable is used with `authentication_ldap_sasl_bind_root_dn`.

#### AUTHENTICATION\_LDAP\_SASL\_CA\_PATH

Option	Description
Command-line	-authentication-ldap-sasl-ca_path
Scope	Global
Dynamic	Yes
Data type	String
Default	NULL

The certificate authority's absolute path used to verify the LDAP certificate.

#### AUTHENTICATION\_LDAP\_SASL\_FALLBACK\_SERVER\_HOST

Option	Description
Command-line	-authentication-ldap-sasl-fallback-
Scope	Global
Dynamic	Yes
Type	String
Default	NULL

Use with `authentication_ldap_sasl_fallback_server_port`.

If the primary server is unavailable, the authentication plugin attempts to connect to the fallback server and authenticate using that server.

#### AUTHENTICATION\_LDAP\_SASL\_FALLBACK\_SERVER\_PORT

Option	Description
Command-line	-authentication-ldap-sasl-fallback-
Scope	Global
Dynamic	Yes
Type	Integer
Default	NULL

Use with `authentication_ldap_sasl_fallback_server_host`.

If the primary server is unavailable, the authentication plugin attempts to connect to the fallback server and authenticate using that server.

If the fallback server host has a value, and the fallback port is 0, users can specify multiple fallback servers.

Use this format to specify multiple fallback servers:  
`authentication_ldap_sasl_fallback_server_host="ldap(s)://host:port,ldap(s)://host2:port2,` for example.

#### AUTHENTICATION\_LDAP\_SASL\_GROUP\_ROLE\_MAPPING

Option	Description
Command-line	-authentication-ldap-sasl-group-ra
Scope	Global
Dynamic	Yes
Data type	String
Default	Null

When an LDAP user logs in, the server checks if the LDAP user is a member of the specified group. If the user is, then the server automatically grants the database server roles to the user.

The variable has this format: `<ldap_group>=<mysql_role>,<ldap_group2>=<mysql_role2>,. . .`

#### AUTHENTICATION\_LDAP\_SASL\_GROUP\_SEARCH\_ATTR

Option	Description
Command-line	<code>--authentication-ldap-sasl-group-search-attr</code>
Scope	Global
Dynamic	Yes
Data type	String
Default	cn

The attribute name that specifies group names in the LDAP directory entries for SASL-based LDAP authentication.

#### AUTHENTICATION\_LDAP\_SASL\_GROUP\_SEARCH\_FILTER

Option	Description
Command-line	<code>--authentication-ldap-sasl-group-search-filter</code>
Scope	Global
Dynamic	Yes
Data type	String
Default	<code>(!((&amp;(objectClass=posixGroup)(member=*))(&amp;(member=%s))))</code>

The custom group search filter for SASL-based LDAP authentication.

#### AUTHENTICATION\_LDAP\_SASL\_INIT\_POOL\_SIZE

Option	Description
Command-line	<code>--authentication-ldap-sasl-init-pool-size</code>
Scope	Global
Dynamic	Yes
Data type	Integer
Default	10
Minimum value	0
Maximum value	32767
Unit	connections

The initial size of the connection pool to the LDAP server for SASL-based LDAP authentication.

**AUTHENTICATION\_LDAP\_SASL\_LOG\_STATUS**

Option	Description
Command-line	-authentication-ldap-sasl-log-status
Scope	Global
Dynamic	Yes
Data type	Integer
Default	1
Minimum value	1
Maximum value	6

The logging level for messages written to the error log for SASL-based LDAP authentication.

**AUTHENTICATION\_LDAP\_SASL\_MAX\_POOL\_SIZE**

Option	Description
Command-line	-authentication-ldap-sasl-max-pool-size
Scope	Global
Dynamic	Yes
Data type	Integer
Default	1000
Minimum value	0
Maximum value	32767
Unit	connections

The maximum connection pool size to the LDAP server in SASL-based LDAP authentication. The variable is used with `authentication_ldap_sasl_init_pool_size`.

**AUTHENTICATION\_LDAP\_SASL\_SERVER\_HOST**

Option	Description
Command-line	-authentication-ldap-sasl-server-host
Scope	Global
Dynamic	Yes
Data type	String
Default	NULL

The LDAP server host used for SASL-based LDAP authentication. The LDAP server host can be an IP address or a host name.

**AUTHENTICATION\_LDAP\_SASL\_SERVER\_PORT**

Option	Description
Command-line	<code>-authentication-ldap-sasl-server-p</code>
Scope	Global
Dynamic	Yes
Data type	Integer
Default	389
Minimum value	1
Maximum value	32376

The LDAP server TCP/IP port number used for SASL-based LDAP authentication.

**AUTHENTICATION\_LDAP\_SASL\_SSL**

Option	Description
Command-line	<code>-authentication-ldap-sasl-ssl=value</code>
Scope	Global
Dynamic	Yes
Data type	Boolean
Default	OFF

If this variable is enabled, the plugin connects to the server with SSL.

**AUTHENTICATION\_LDAP\_SASL\_TLS**

Option	Description
Command-line	<code>-authentication-ldap-sasl-tls=value</code>
Scope	Global
Dynamic	Yes
Data type	Boolean
Default	OFF

If this variable is enabled, the plugin connects to the server with TLS.

**AUTHENTICATION\_LDAP\_SASL\_USER\_SEARCH\_ATTR**

Option	Description
Command-line	<code>-authentication-ldap-sasl-user-sea</code>
Scope	Global
Dynamic	Yes
Data type	String

Option	Description
Default	uid

The attribute name that specifies the user names in LDAP directory entries in SASL-based LDAP authentication.

#### AUTHENTICATION\_LDAP\_SIMPLE\_BIND\_BASE\_DN

Option	Description
Command-line	-authentication-ldap-simple-bind-b
Scope	Global
Dynamic	Yes
Data type	String
Default	NULL

The base distinguished name (DN) for simple LDAP authentication. You can limit the search scope by using the variable as the base of the search.

#### AUTHENTICATION\_LDAP\_SIMPLE\_BIND\_ROOT\_DN

Option	Description
Command-line	-authentication-ldap-simple-bind-r
Scope	Global
Dynamic	Yes
Data type	String
Default	NULL

The `root` distinguished name (DN) used to authenticate simple LDAP. When performing a search, this variable is used with `authentication_ldap_simple_bind_root_pwd` as the authenticating credentials to the LDAP server.

#### AUTHENTICATION\_LDAP\_SIMPLE\_BIND\_ROOT\_PWD

Option	Description
Command-line	-authentication-ldap-simple-bind-r
Scope	Global
Dynamic	Yes
Data type	String
Default	NULL

The `root` password used to authenticate against simple LDAP server. This variable is used with `authentication_ldap_simple_bind_root_dn`.

**AUTHENTICATION\_LDAP\_SIMPLE\_CA\_PATH**

Option	Description
Command-line	-authentication-ldap-simple-ca_path
Scope	Global
Dynamic	Yes
Data type	String
Default	NULL

The certificate authority's absolute path used to verify the LDAP certificate.

**AUTHENTICATION\_LDAP\_SIMPLE\_FALLBACK\_SERVER\_HOST**

Option	Description
Command-line	-authentication-ldap-simple-fallback-server-host
Scope	Global
Dynamic	Yes
Type	String
Default	NULL

Use with [authentication\\_ldap\\_simple\\_fallback\\_server\\_port](#).

If the primary server is unavailable, the authentication plugin attempts to connect to the fallback server and authenticate using that server.

**AUTHENTICATION\_LDAP\_SIMPLE\_FALLBACK\_SERVER\_PORT**

Option	Description
Command-line	-authentication-ldap-simple-fallback-server-port
Scope	Global
Dynamic	Yes
Type	Integer
Default	NULL

Use with [authentication\\_ldap\\_simple\\_fallback\\_server\\_host](#).

If the primary server is unavailable, the authentication plugin attempts to connect to the fallback server and authenticate using that server.

If the fallback server host has a value, and the fallback port is 0, users can specify multiple fallback servers.

Use this format to specify multiple fallback servers:  
`authentication_ldap_simple_fallback_server_host="ldap(s)://host:port,ldap(s)://host2:port2,` for example.

## AUTHENTICATION\_LDAP\_SIMPLE\_GROUP\_ROLE\_MAPPING

Option	Description
Command-line	-authentication-ldap-simple-group
Scope	Global
Dynamic	Yes
Data type	String
Default	Null

When an LDAP user logs in, the server checks if the LDAP user is a member of the specified group. If the user is, then the server automatically grants the database server roles to the user.

The variable has this format: <ldap\_group>=<mysql\_role>,<ldap\_group2>=<mysql\_role2> , .

## AUTHENTICATION\_LDAP\_SIMPLE\_GROUP\_SEARCH\_ATTR

Option	Description
Command-line	-authentication-ldap-simple-group
Scope	Global
Dynamic	Yes
Data type	String
Default	cn

The attribute name that specifies group names in the LDAP directory entries for simple LDAP authentication.

## AUTHENTICATION\_LDAP\_SIMPLE\_GROUP\_SEARCH\_FILTER

Option	Description
Command-line	-authentication-ldap-simple-group
Scope	Global
Dynamic	Yes
Data type	String
Default	(((&(objectClass=posixGroup)(member=*))

The custom group search filter for simple LDAP authentication.

## AUTHENTICATION\_LDAP\_SIMPLE\_INIT\_POOL\_SIZE

Option	Description
Command-line	-authentication-ldap-simple-init-po
Scope	Global
Dynamic	Yes

Option	Description
Data type	Integer
Default	10
Minimum value	0
Maximum value	32767
Unit	connections

The initial size of the connection pool to the LDAP server for simple LDAP authentication.

#### AUTHENTICATION\_LDAP\_SIMPLE\_LOG\_STATUS

Option	Description
Command-line	-authentication-ldap-simple-log-st
Scope	Global
Dynamic	Yes
Data type	Integer
Default	1
Minimum value	1
Maximum value	6

The logging level for messages written to the error log for simple LDAP authentication.

#### AUTHENTICATION\_LDAP\_SIMPLE\_MAX\_POOL\_SIZE

Option	Description
Command-line	-authentication-ldap-simple-max-p
Scope	Global
Dynamic	Yes
Data type	Integer
Default	1000
Minimum value	0
Maximum value	32767
Unit	connections

The maximum connection pool size to the LDAP server in simple LDAP authentication. The variable is used with `authentication_ldap_simple_init_pool_size`.

#### AUTHENTICATION\_LDAP\_SIMPLE\_SERVER\_HOST

Option	Description
Command-line	-authentication-ldap-simple-server
Scope	Global

Option	Description
Dynamic	Yes
Data type	String
Default	NULL

The LDAP server host used for simple LDAP authentication. The LDAP server host can be an IP address or a host name.

#### AUTHENTICATION\_LDAP\_SIMPLE\_SERVER\_PORT

Option	Description
Command-line	-authentication-ldap-simple-server
Scope	Global
Dynamic	Yes
Data type	Integer
Default	389
Minimum value	1
Maximum value	32376

The LDAP server TCP/IP port number used for simple LDAP authentication.

#### AUTHENTICATION\_LDAP\_SIMPLE\_SSL

Option	Description
Command-line	-authentication-ldap-simple-ssl=va
Scope	Global
Dynamic	Yes
Data type	Boolean
Default	OFF

If this variable is enabled, the plugin connects to the server with SSL.

#### AUTHENTICATION\_LDAP\_SIMPLE\_TLS

Option	Description
Command-line	-authentication-ldap-simple-tls=va
Scope	Global
Dynamic	Yes
Data type	Boolean
Default	OFF

If this variable is enabled, the plugin connects to the server with TLS.

## AUTHENTICATION\_LDAP\_SIMPLE\_USER\_SEARCH\_ATTR

Option	Description
Command-line	-authentication-ldap-simple-user-s
Scope	Global
Dynamic	Yes
Data type	String
Default	uid

The attribute name that specifies the user names in LDAP directory entries in simple LDAP authentication.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-31

## 9.2 Data masking

### 9.2.1 Data masking overview

Data masking helps to limit the exposure of sensitive data by preventing access to non-authorized users. Masking provides a way to create a version of the data in situations, such as a presentation, sales demo, or software testing, when the real data should not be used. Data masking changes the data values while using the same format and cannot be reverse engineered. Masking reduces an organization's risk by making the data useless to an outside party.

#### Data masking techniques

The common data masking techniques are the following:

Technique	Description
Custom string	Replaces sensitive data with a specific string, such as XXX-XXX-XXXX
Data substitution	Replaces sensitive data with realistic data from a dictionary

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#) [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-25

## 9.2.2 Install the data masking component

The component has the following parts:

- A database server system table used to store the terms and dictionaries
- A `component_masking_functions` component that contains the loadable functions

The `MASKING_DICTIONARIES_ADMIN` privilege may be required by some functions.

### Install the component

The following steps install the component:

2.  Create `masking_dictionaries`.

```
mysql> CREATE TABLE IF NOT EXISTS
mysql.masking_dictionaries(
  Dictionary VARCHAR(256) NOT NULL,
  Term VARCHAR(256) NOT NULL,
  UNIQUE INDEX dictionary_term_idx (Dictionary, Term)
) ENGINE = InnoDB DEFAULT CHARSET=utf8mb4;
```

2.  Install the data masking components and the loadable functions.

```
mysql> INSTALL COMPONENT 'file://component_masking_functions';
```

2.  The `MASKING_DICTIONARIES_ADMIN` is required to use the the following functions:

- `masking_dictionary_term_add`
- `masking_dictionary_term_remove`
- `masking_dictionary_remove`

```
mysql> GRANT MASKING_DICTIONARIES_ADMIN ON *.* TO <user>;
```

## Useful links

[Uninstall the data masking component](#)

[Data masking component functions](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-04

## 9.2.3 Data masking component functions

The feature is in [tech preview](#).

Name	Usage
<code>gen_blocklist(str, from_dictionary_name, to_dictionary_name)</code>	Replace a term from a dictionary
<code>gen_dictionary(dictionary_name)</code>	Returns a random term from a dictionary
<code>gen_range(lower, upper)</code>	Returns a number from a range
<code>gen_rnd_canada_sin()</code>	Generates a Canadian Social Insurance Number
<code>gen_rnd_email([name_size, surname_size, domain])</code>	Generates an email address
<code>gen_rnd_iban([country, size])</code>	Generates an International Bank Account Number
<code>gen_rnd_pan()</code>	Generates a Primary account number
<code>gen_rnd_ssn()</code>	Generates a US Social Security number
<code>gen_rnd_uk_nin()</code>	Generates a United Kingdom National Insurance Number
<code>gen_rnd_us_phone()</code>	Generates a US phone number
<code>gen_rnd_uuid()</code>	Generates a Universally Unique Identifier
<code>mask_canada_sin(str [,mask_char])</code>	Masks the Canadian Social Insurance Number
<code>mask_iban(str [,mask_char])</code>	Masks the International Bank Account Number
<code>mask_inner(str, margin1, margin2 [,mask_char])</code>	Masks the inner part of a string
<code>mask_outer(str, margin1, margin2 [,mask_char])</code>	Masks the outer part of the string
<code>mask_pan(str [,mask_char])</code>	Masks the Primary Account Number
<code>mask_pan_relaxed(str [,mask_char])</code>	Partially masks the Primary Account Number
<code>mask_ssn(str [,mask_char])</code>	Masks the US Social Security number
<code>mask_uk_nin(str [,mask_char])</code>	Masks the United Kingdom National Insurance Number

Name	Usage
<code>mask_uuid(str [,mask_char])</code>	Masks the Universally Unique Identifier
<code>masking_dictionary_remove(dictionary_name)</code>	Removes the dictionary
<code>masking_dictionary_term_add(dictionary_name, term_name)</code>	Adds a term to the masking dictionary
<code>masking_dictionary_term_remove(dictionary_name, term_name)</code>	Removes a term from the masking dictionary

### `gen_blocklist(str, from_dictionary_name, to_dictionary_name)`

Replaces one term in a dictionary with a term, selected at random, in another dictionary.

#### PARAMETERS

Parameter	Optional	Description
<code>term</code>	No	The term to replace
<code>from_dictionary_name</code>	No	The dictionary that stores the term
<code>to_dictionary_name</code>	No	The dictionary that stores the replacement term

#### RETURNS

A term, selected at random, from the dictionary listed in `to_dictionary_name` that replaces the selected term. If the selected `term` is not listed in the `from_dictionary_name` or a dictionary is missing, then the `term` is returned. If the `to_dictionary_name` does not exist, then returns NULL. The character set of the returned string is the same character set of the `term` parameter.

Returns NULL if you invoke this function with NULL as the primary argument.

#### EXAMPLE

```
mysql> SELECT gen_blocklist('apple', 'fruit', 'nut');
```

#### Expected output

```
+-----+
| gen_blocklist('apple', 'fruit', 'nut') |
+-----+
| walnut                                |
+-----+
```

### `gen_dictionary(dictionary_name)`

Returns a term from a dictionary selected at random.

#### PARAMETERS

Parameter	Optional	Description
<code>dictionary_name</code>	No	Select the random term from the dictionary

## RETURNS

A random term from the dictionary listed in `dictionary_name` in the `utf8mb4` character set. Returns NULL if the `dictionary_name` does not exist.

## EXAMPLE

```
mysql> SELECT gen_dictionary('trees');
```

## Expected output

```
+-----+
| gen_dictionary('trees') |
+-----+
| Norway spruce          |
+-----+
```

**gen\_range(lower, upper)**

Returns a number from a defined range.

## PARAMETERS

Parameter	Optional	Description
<code>lower</code>	No	The lower boundary of the range
<code>upper</code>	No	The upper boundary of the range

The `upper` parameter value must be an integer either greater than or equal to the `lower` parameter value.

## RETURNS

An integer, selected at random, from an inclusive range defined by the `lower` parameter value and the `upper` parameter value, or NULL if the `upper` boundary is less than the `lower` boundary.

## EXAMPLE

```
mysql> SELECT gen_range(10, 100);
```

## Expected output

```
+-----+
| gen_range(10,100)      |
+-----+
| 56                     |
+-----+
```

**gen\_rnd\_canada\_sin()**

Generates a Canada Social Insurance Number (SIN).

 **Important**

Only use this function for testing because the result could be a legitimate SIN. Use `mask_canada_sin` to disguise the result if you must publish the result.

**PARAMETERS**

None.

**RETURNS**

Returns a Canada SIN formatted in three groups of three digits (for example, 123-456-789) in the `utf8mb4` character set. To ensure the number is consistent, the number is verified with the [Luhn algorithm](#).

**EXAMPLE**

```
mysql> SELECT gen_rnd_canada_sin();
```

 **Expected output**

```
+-----+
| gen_rnd_canada_sin() |
+-----+
| 506-948-819         |
+-----+
```

**gen\_rnd\_email([name\_size, surname\_size, domain])**

Generates a random email address in the `name.surname@domain` format.

**PARAMETERS**

Parameter	Optional	Description
<code>name_size</code>	Yes	Specifies the number of characters in the name part. The default number is five. The minimum number is one. The maximum number is 1024.
<code>surname_size</code>	Yes	Specifies the number of characters in the surname part. The default number is seven. The minimum number is one. The maximum number is 1024.
<code>domain</code>	Yes	Specifies the domain name used. The default value is <code>example.com</code> .

**RETURNS**

A generated email address as a string in the same character set as `domain`. If the `domain` value is not specified, then the string is in the `utf8mb4` character set. The `name` and `surname` are random lower-case letters (a - z).

## EXAMPLE

```
mysql> SELECT gen_rnd_email(name_size=4, surname_size=5, domain='mydomain.edu');
```

 Expected output

```
+-----+
| gen_rnd_email(4, 5, 'mydomain.edu') |
+-----+
| qwer.asdfg@mydomain.edu           |
+-----+
```

**gen\_rnd\_iban([country, size])**

Generates an Internal Bank Account Number (IBAN).

 Important

Generating an IBAN with a valid country code should only be used for testing. The function does not check if the generated value is a legitimate bank account. If you must publish the result, consider using `mask_iban` to disguise the result. The function does not perform a checksum on the bank account number.

## PARAMETERS

Parameter	Optional	Description
country	Yes	A two-character country code
size	Yes	Number of characters

If the `country` is not specified, the default value is `ZZ`. The value must be two upper-case characters (A-Z) or an error is returned.

The default value for `size` is 16. The minimum value is 15. The maximum value is 34.

## RETURNS

The function returns a string that is the length of the `size` value. The string consists of `country` (two characters) followed by the  $(size - 2)$  random digits.

The character set is the same as the `country` parameter or if that parameter is not specified, the character set is `utf8mb4`.

## EXAMPLE

```
mysql> SELECT gen_rnd_iban();
```

**Expected output**

```
+-----+
| gen_rnd_iban() |
+-----+
| ZZ78959120078536 |
+-----+
```

**gen\_rnd\_pan()**

Generates a Primary Account Number (PAN) for a payment card that passes basic checksum validation.

The generated PAN can be one of the following:

- American Express
- Visa
- Mastercard
- Discover

**Important**

Generating the PAN should only be used for testing. The function does not check if the generated value is a legitimate primary account number. If you must publish the result, consider using `mask_pan` or `mask_pan_relaxed()` to disguise the result.

**PARAMETERS**

None

**RETURNS**

A random PAN string in `utf8mb4` character set.

**EXAMPLE**

```
mysql> SELECT gen_rnd_pan();
```

**Expected output**

```
+-----+
| gen_rnd_pan() |
+-----+
| 1234567898765432 |
+-----+
```

**gen\_rnd\_ssn()**

Generates a United States Social Security Account Number (SSN).

## PARAMETERS

None

## RETURNS

A SSN string in a nine-digit number format “AAA-GG-SSSS” in the `utf8mb4` character set. The number has three parts, the first three digits are the area number, the group number, and the serial number. The generated SSN uses ‘900’ or greater numbers for the area number. These numbers are not legitimate because they are outside the approved range.

## EXAMPLE

```
mysql> SELECT gen_rnd_ssn();
```

 Expected output

```
+-----+
| gen_rnd_ssn() |
+-----+
| 970-03-0370   |
+-----+
```

**gen\_rnd\_uk\_nin()**

Generates a United Kingdom National Insurance Number (NIN).

 Important

This function should only be used for testing. The function does not check if the generated value is a legitimate United Kingdom National Insurance number. If you must publish the result, consider masking the result with `mask_uk_nin`.

## PARAMETERS

None.

## RETURNS

A NIN string in the `utf8mb4` character set. The string is nine (9) characters in length, always starts with ‘AA’ and ends with ‘C’.

## EXAMPLE

```
mysql> SELECT gen_rnd_uk_nin();
```

 Expected output

```
+-----+
| gen_rnd_uk_nin() |
+-----+
| AA123456C        |
+-----+
```

**gen\_rnd\_us\_phone()**

Generates a United States phone number with the 555 area code. The '555' area code represents fictional numbers.

**PARAMETERS**

None

**RETURNS**

Returns a United States phone number in the `utf8mb4` character set.

**EXAMPLE**

```
mysql> SELECT gen_rnd_us_phone();
```

**Expected output**

```
+-----+
| gen_rnd_us_phone() |
+-----+
| 1-555-249-2029     |
+-----+
```

**gen\_rnd\_uuid()**

Generates a version 4 Universally Unique Identifier (UUID).

**PARAMETERS**

None.

**RETURNS**

Returns a UUID as a string in the `utf8mb4` character set.

**EXAMPLE**

```
mysql> SELECT gen_rnd_uuid();
```

**Expected output**

```
+-----+
| gen_rnd_uuid()          |
+-----+
| 9a3b642c-06c6-11ee-be56-0242ac120002 |
+-----+
```

**mask\_canada\_sin(str [,mask\_char])**

Masks a Canada Social Insurance Number (SIN).

## PARAMETERS

Parameter	Optional	Description
<code>str</code>	No	The string to be masked
<code>mask_char</code>	Yes	The masking character

The `str` accepts an alphanumeric string.

If you do not specify a `mask_char`, the default character is `X`. The `mask_char` value can be a multibyte character in any character set and may not be same character set as `str`.

## RETURNS

A string with the selected characters masked by a specified `mask_char` or the default value for that parameter. The function supports multibyte characters in any character set. The character set of the return value is the same as `str`.

An error is reported if `str` length is an incorrect length.

Returns a NULL if you invoke this function with NULL as the primary argument.

## EXAMPLE

```
mysql> SELECT mask_canada_sin('555-555-555');
```

## Expected output

```
+-----+
| mask_canada_sin('555-555-555') |
+-----+
| XXX-XXX-XXX |
+-----+
```

**mask\_iban(str [,mask\_char])**

Masks an Internal Bank Account Number (IBAN).

## PARAMETERS

Parameter	Optional	Description
<code>str</code>	No	The string to be masked
<code>mask_char</code>	Yes	Character used for masking

The `str` accepts either of the following:

- No separator symbol
- Groups of four characters. These groups can be separated by a space or any separator character.

The default value for `mask_char` is `*`. The value can be a multibyte character in any character set and may not be same character set as `str`.

## RETURNS

Returns the masked string. The character set of the result is the same as the character set of `str`.

An error is reported if the `str` length is incorrect.

Returns NULL if you invoke this function with NULL as the primary argument.

## EXAMPLE

```
mysql> SELECT mask_iban('DE27 1002 02003 77495 4156');
```

## Expected output

```
+-----+
| mask_iban('DE27 1002 02003 77495 4156') |
+-----+
| DE** **** **** **** ****             |
+-----+
```

**mask\_inner(str, margin1, margin2 [,mask\_char])**

Returns the string where a selected inner portion is masked with a substitute character.

## PARAMETERS

Parameter	Optional	Description
<code>string</code>	No	The string to be masked
<code>margin1</code>	No	The number of characters on the end of the string to remain unmasked
<code>margin2</code>	No	The number of characters on the right end of the string to remain unmasked
<code>mask_char</code>	Yes	The masking character

The `margin1` value cannot be a negative number. A value of 0 (zero) masks all characters.

The `margin2` value cannot be a negative number. A value of 0 (zero) masks all characters.

If the sum of `margin1` and `margin2` is greater than or equal to the string length, no masking occurs.

If the `mask_char` is not specified, the default is 'X'. The `mask_char` value can be a multibyte character in any character set and may not be same character set as `str`.

## RETURNS

A string with the selected characters masked by a specified `mask_char` or that parameter's default value in the character set of the `string` parameter.

Returns NULL if you invoke this function with NULL as the primary argument.

## EXAMPLE

```
mysql> SELECT mask_inner('123456789', 1, 2);
```

#### Expected output

```
+-----+
| mask_inner('123456789', 1, 2) |
+-----+
| 1XXXXX89                       |
+-----+
```

### mask\_outer(str, margin1, margin2 [,mask\_char])

Returns the string where a selected outer portion is masked with a substitute character.

#### PARAMETERS

Parameter	Optional	Description
string	No	The string to be masked
margin1	No	On the left end of the string, mask this designated number of characters
margin2	No	On the right end of the string, mask this designated number of characters
mask_char	Yes	The masking character

The `margin1` cannot be a negative number. A value of 0 (zero) does not mask any characters.

The `margin2` cannot be a negative number. A value of 0 (zero) does not mask any characters.

If the sum of `margin1` and `margin2` is greater than or equal to the string length, the string is masked.

If the `mask_char` is not specified, the default is 'X'. The `mask_char` value can be a multibyte character in any character set and may not be same character set as `str`.

#### RETURNS

A string with the selected characters masked by a specified `mask_char` or that parameter's default value in the same character set as `string`.

Returns NULL if you invoke this function with NULL as the primary argument.

#### EXAMPLE

```
mysql> SELECT mask_outer('123456789', 2, 2);
```

**Expected output**

```
+-----+
| mask_outer('123456789', 2, 2). |
+-----+
| XX34567XX |
+-----+
```

**mask\_pan(str [,mask\_char])**

Returns a masked payment card Primary Account Number (PAN). The mask replaces the PAN number with the specified character except for the last four digits.

**PARAMETERS**

Parameter	Optional	Description
str	No	The string to be masked
mask_char	Yes	The masking character

The `str` contains a minimum of 14 or a maximum of 19 alphanumeric characters.

If the `mask_char` is not specified, the default value is 'X'. The `mask_char` value can be a multibyte character in any character set and may not be same character set as `str`.

**RETURNS**

A string with the selected characters masked by a specified `mask_char` or that parameter's default value. The character set of the result is the same character set as `str`.

An error occurs if the `str` parameter is not the correct length.

Returns NULL if you invoke this function with NULL as the primary argument.

**EXAMPLE**

```
mysql> SELECT mask_pan (gen_rnd_pan());
```

**Expected output**

```
+-----+
| mask_pan(gen_rnd_pan()) |
+-----+
| XXXXXXXXXXXX2345 |
+-----+
```

**mask\_pan\_relaxed(str [,mask\_char])**

Returns a masked payment card Primary Account Number (PAN). The first six numbers and the last four numbers and the rest of the string masked by specified character or X.

## PARAMETERS

Parameter	Optional	Description
<code>str</code>	No	The string to be masked
<code>mask_char</code>	Yes	The specified character for mask

The `str` must contain a minimum of 14 or a maximum of 19 alphanumeric characters.

If the `mask_char` is not specified, the default value is 'X'.

## RETURNS

A string with the first six numbers and the last four numbers and the rest of the string masked by a specified `mask_char` or that parameter's default value (X). The character set of the result is the same character set as `str`.

The `mask_char` value can be a multibyte character in any character set and may not be same character set as `str`.

Reports an error is the `str` parameter is not the correct length.

Returns NULL if you invoke this function with NULL as the primary argument.

## EXAMPLE

```
mysql> SELECT mask_pan_relaxed(gen_rnd_pan());
```

## Expected output

```
+-----+
| mask_pan_relaxed(gen_rnd_pan()) |
+-----+
| 520754XXXXXX4848              |
+-----+
```

**mask\_ssn(str [,mask\_char])**

Returns a masked United States Social Security Number(SSN). The mask replaces the SSN number with the specified character except for the last four digits.

## PARAMETERS

Parameter	Optional	Description
<code>str</code>	No	The string to be masked
<code>mask_char</code>	Yes	The masking character

The `str` accepts either of the following:

- Nine integers, no separator symbol
- Nine integers in the AAA-GG-SSS pattern. The - (dash symbol) is the separator character.

If the `mask_char` is not specified, the default value is `*`. The `mask_char` value can be a multibyte character in any character set and may not be same character set as `str`.

#### RETURNS

A string with the selected characters masked by a specified `mask_char` or that parameter's default value in the same character set of `str`.

Reports an error if the value of the `str` is an incorrect length.

Returns a NULL value if you invoke this function with NULL as the primary argument.

#### EXAMPLE

```
mysql> SELECT mask_ssn('555-55-5555', 'X');
```

#### Expected output

```
+-----+
| mask_ssn('555-55-5555', 'X') |
+-----+
| XXX-XX-5555                    |
+-----+
```

### mask\_uk\_nin(`str` [,`mask_char`])

Returns a masked a United Kingdom National Insurance Number (NIN). The mask replaces the NIN number with the specified character except for the first two digits.

#### PARAMETERS

Parameter	Optional	Description
<code>str</code>	No	The string to be masked
<code>mask_char</code>	Yes	The masking character

The `str` accepts an alpha-numeric string and does not check format and the `str` can use any separator character.

If the `mask_char` is not specified, the default value is `*`. The `mask_char` value can be a multibyte character in any character set and may not be same character set as `str`.

#### RETURNS

Returns a string with the selected characters masked by a specified `mask_char` or that parameter's default value in the same character set as `str`.

An error occurs if the `str` parameter is not the correct length.

Returns a NULL value if you invoke this function with NULL as the primary argument.

#### EXAMPLE

```
mysql> SELECT mask_uk_nin ('CT 26 46 83 D');
```

```

Expected output
+-----+
| mask_uk_nin('CT 26 46 83 D') |
+-----+
| CT ** ** ** * |
+-----+
    
```

### mask\_uuid(str [,mask\_char])

Masks a Universally Unique Identifier (UUID).

PARAMETERS

Parameter	Optional	Description
str	No	The string to be masked
mask_char	Yes	The masking character

The `str` format is `*****_****_****_****_*****`.

If the `mask_char` is not specified, the default value is `*`. The `mask_char` value can be a multibyte character in any character set and may not be same character set as `str`.

RETURNS

A string with the characters masked by a specified `mask_char` or that parameter’s default value in the same character set as `str`.

Returns an error if the length of `str` is incorrect.

Returns `NULL` if you invoke this function with `NULL` as the primary argument.

EXAMPLE

```
mysql> SELECT mask_uuid('9a3b642c-06c6-11ee-be56-0242ac120002');
```

```

Expected output
+-----+
| mask_uuid('9a3b642c-06c6-11ee-be56-0242ac120002') |
+-----+
| *****_****_****_****_***** |
+-----+
    
```

### masking\_dictionary\_remove(dictionary\_name)

Removes all of the terms and then removes the dictionary.

Requires the `MASKING_DICTIONARIES_ADMIN` privilege.

## PARAMETERS

Parameter	Optional	Description
<code>dictionary_name</code>	No	The dictionary to be removed

## RETURNS

Returns a string value of `1` (one) in the `utf8mb4` character set if the operation is successful or `NULL` if the operation could not find the `dictionary_name`.

## EXAMPLE

```
mysql> SELECT masking_dictionary_remove('trees');
```

## Expected output

```
+-----+
| masking_dictionary_remove('trees') |
+-----+
|                                     1 |
+-----+
```

**masking\_dictionary\_term\_add(dictionary\_name, term\_name)**

Adds a term to the dictionary and requires the `MASKING_DICTIONARIES_ADMIN` privilege.

## PARAMETERS

Parameter	Optional	Description
<code>dictionary_name</code>	No	The dictionary where the term is added
<code>term_name</code>	No	The term added to the selected dictionary

## RETURNS

Returns a string value of `1` (one) in the `utf8mb4` character set if the operation is successful. If the `dictionary_name` does not exist, the operation creates the dictionary.

Returns `NULL` if the operation fails. An operation can fail if the `term_name` is already available in the dictionary specified by `dictionary_name`.

The operation uses `INSERT IGNORE` and can have the following outcomes:

- The `term_name` is truncated if the `term_name` length is greater than maximum length of the `Term` field in the `mysql.masking_dictionaries` table.
- The character of the `dictionary_name` is not supported by the `Dictionary` field in `mysql.masking_dictionaries` table, the character is implicitly converted to '?'.
- If the character of the `term_name` is not supported by the `Term` field in the `mysql.masking_dictionaries` table, the character is implicitly converted to '?'.

The following command returns the table information:

```
mysql> DESCRIBE mysql.masking_dictionaries;
```

The result returns the table structure.

#### Expected output

```
+-----+-----+-----+-----+-----+-----+
| Field      | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| Dictionary | varchar(256)  | NO   | PRI | NULL    |       |
| Term       | varchar(256)  | NO   | PRI | NULL    |       |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.02 sec)
```

Modify the table with an `ALTER TABLE` statement, if needed.

#### EXAMPLE

```
mysql> SELECT masking_dictionary_term_add('trees','pine');
```

#### Expected output

```
+-----+-----+
| masking_dictionary_term_add('trees', 'pine') |
+-----+-----+
|   | 1 |
+-----+-----+
```

### masking\_dictionary\_term\_remove(dictionary\_name, term\_name)

Removes the selected term from the dictionary.

Requires the `MASKING_DICTIONARIES_ADMIN` privilege.

#### PARAMETERS

Parameter	Optional	Description
<code>dictionary_name</code>	No	The dictionary that contains the <code>term_name</code>
<code>term_name</code>	No	The term to be removed

#### RETURNS

Returns a string value of `1` (one) in the `utf8mb4` character set if the operation is successful.

Returns `NULL` if the operation fails. An operation can fail if the following occurs:

- The `term_name` is not available in the dictionary specified by `dictionary_name`
- The `dictionary_name` could not be found

## PARAMETERS

Parameter	Optional	Description
dictionary_name	No	The dictionary for the term
term_name	No	The term to be added

## EXAMPLE

```
mysql> SELECT masking_dictionary_term_remove('trees','pine');
```

 Expected output

```
+-----+
| masking_dictionary_term_remove('trees', 'pine') |
+-----+
|                                     1 |
+-----+
```

## Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-25

## 9.2.4 Uninstall the component

The following steps uninstall the component:

1. Uninstall the component with `UNINSTALL_COMPONENT` and the loadable functions.

```
mysql> UNINSTALL COMPONENT 'file://component_masking_functions';
```

2. Drop `masking_dictionaries`.

```
mysql> DROP TABLE mysql.masking_dictionaries;
```

### Useful links

[Install the data masking component](#)

[Data masking component functions](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#) [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-04

## 9.3 PAM authentication plugin

Percona PAM Authentication Plugin is a free and Open Source implementation of the *MySQL's* authentication plugin. This plugin acts as a mediator between the *MySQL* server, the *MySQL* client, and the PAM stack. The server plugin requests authentication from the PAM stack, forwards any requests and messages from the PAM stack over the wire to the client (in cleartext) and reads back any replies for the PAM stack.

PAM plugin uses dialog as its client side plugin. Dialog plugin can be loaded to any client application that uses `libperconaserverclient / libmysqlclient` library.

Here are some of the benefits that Percona dialog plugin offers over the default one:

- It correctly recognizes whether PAM wants input to be echoed or not, while the default one always echoes the input on the user's console.
- It can use the password which is passed to *MySQL* client via “-p” parameter.
- Dialog client [installation bug](#) has been fixed.

Percona offers two versions of this plugin:

- Full PAM plugin called `auth_pam`. This plugin uses `dialog.so`. It fully supports the PAM protocol with arbitrary communication between client and server.
- Oracle-compatible PAM called `auth_pam_compat`. This plugin uses `mysql_clear_password` which is a part of Oracle *MySQL* client. It also has some limitations, such as, it supports only one password input. You must use `-p` option in order to pass the password to `auth_pam_compat`.

These two versions of plugins are physically different. To choose which one you want used, you must use `IDENTIFIED WITH 'auth_pam'` for `auth_pam`, and `IDENTIFIED WITH 'auth_pam_compat'` for `auth_pam_compat`.

### 9.3.1 Version specific information

A plugin may not be supported in later releases of *MySQL* or Percona Server for *MySQL* since version changes may introduce incompatible changes.

### 9.3.2 Installation

This plugin requires manual installation because it isn't installed by default.

```
mysql> INSTALL PLUGIN auth_pam SONAME 'auth_pam.so';
```

After the plugin has been installed it should be present in the plugins list. To check if the plugin has been correctly installed and active

```
mysql> SHOW PLUGINS;
```

#### Expected output

```
...
| auth_pam | ACTIVE | AUTHENTICATION | auth_pam.so | GPL |
```

### 9.3.3 Configuration

In order to use the plugin, authentication method should be configured. Simple setup can be to use the standard UNIX authentication method ( `pam_unix` ).

#### Note

To use `pam_unix`, mysql will need to be added to the shadow group in order to have enough privileges to read the `/etc/shadow`.

A sample `/etc/pam.d/mysql` file:

```
auth    required    pam_unix.so
account required    pam_unix.so
```

For added information in the system log, you can expand it to be:

```
auth    required    pam_warn.so
auth    required    pam_unix.so audit
account required    pam_unix.so audit
```

### 9.3.4 Creating a user

After the PAM plugin has been configured, users can be created with the PAM plugin as authentication method

```
mysql> CREATE USER 'newuser'@'localhost' IDENTIFIED WITH auth_pam;
```

This will create a user `newuser` that can connect from `localhost` who will be authenticated using the PAM plugin. If the `pam_unix` method is being used user will need to exist on the system.

### 9.3.5 Supplementary groups support

*Percona Server for MySQL* has implemented PAM plugin support for supplementary groups. Supplementary or secondary groups are extra groups a specific user is member of. For example user

`joe` might be a member of groups: `joe` (his primary group) and secondary groups `developers` and `dba`. A complete list of groups and users belonging to them can be checked with `cat /etc/group` command.

This feature enables using secondary groups in the mapping part of the authentication string, like `"mysql, developers=joe, dba=mark"`. Previously only primary groups could have been specified there. If user is a member of both `developers` and `dba`, PAM plugin will map it to the `joe` because `developers` matches first.

### 9.3.6 Known issues

Default `mysql` stack size is not enough to handle `pam_encryptfs` module. The workaround is to increase the `MySQL` stack size by setting the `thread-stack` variable to at least `512KB` or by increasing the old value by `256KB`.

PAM authentication can fail with `mysqld: pam_unix(mysqld:account): Fork failed: Cannot allocate memory` error in the `/var/log/secure` even when there is enough memory available. Current workaround is to set `vm.overcommit_memory` to `1`:

```
echo 1 /proc/sys/vm/overcommit_memory
```

and by adding the `vm.overcommit_memory = 1` to `/etc/sysctl.conf` to make the change permanent after reboot. Authentication of internal (i.e. non PAM) accounts continues to work fine when `mysqld` reaches this memory utilization level. *NOTE:* Setting the `vm.overcommit_memory` to `1` will cause kernel to perform no memory overcommit handling which could increase the potential for memory overload and invoking of OOM killer.

### 9.3.7 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 9.4 SSL improvements

Percona Server for MySQL passes Elliptic Curve Cryptography (ECC) ciphers to OpenSSL by default.

#### Note

Although documented as supported, elliptic-curve crypto-based ciphers do not work with `MySQL`.

## 9.4.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 9.5 FIDO authentication plugin

Percona Server for MySQL supports the Fast Identify Online (FIDO) authentication method that uses a plugin. The FIDO authentication provides a set of standards that reduces the reliance on passwords.

The server-side fido authentication plugin enables authentication using external devices. If this plugin is the only authentication plugin used by the account, this plugin allows authentication without a password. Multi-factor authentication can use non-FIDO MySQL authentication methods, the FIDO authentication method, or a combination of both.

All distributions include the client-side `authentication_fido_client` plugin. This plugin allows clients to connect to accounts that use `authentication_fido` and authenticate on a server that has that plugin loaded.

### 9.5.1 Plugin and library file names

The plugin and library file names are listed in the following table.

Plugin or file name	Plugin or library file name
Server-side plugin	<code>authentication_fido</code>
Client-side plugin	<code>authentication_fido_client</code>
Library file	<code>authentication_fido.so</code>

### 9.5.2 Install the FIDO authentication plugin

The library file must be stored in the directory named by the `[ plugin_dir ]` variable.

**At server startup**   **Edit my.cnf and restart the server**   **Load the plugin at runtime**

At server startup, use the `[ --plugin_load_add ]` option with the library name. The option must be added each time the server starts.

```
[mysqld]
...
plugin-load-add=authentication_fido.so
...

mysql> INSTALL PLUGIN authentication_fido SONAME `authentication_fido.so`;
```

## Verify installation

Use the [ `SHOW PLUGINS` ] statement or query the `INFORMATION_SCHEMA.PLUGINS` table to verify that the plugin was loaded successfully and is active.

Check the server error log if the plugin is not loaded.

## 9.5.3 FIDO authentication strategies

FIDO can be used with [non-FIDO authentication](#). FIDO can be used to [create IFA accounts that do not require passwords](#).

### Use FIDO authentication with non-FIDO authentication

A FIDO device is associated with the account using FIDO authentication. The FIDO device must be registered before the account can be used in a one-time process. This device must be available and the user must perform whatever FIDO device action required, such as adding a thumbprint, or the registration fails.

The registration can only be performed by the user named by the account. An error occurs if a user attempts the registration for another user.

The device registration can be performed on the `mysql` client or MySQL Shell. Use the `--fido-register-factor` option with the factor or factors for the device. For example, if you are using FIDO as a second authentication method, which is a common practice, the statement is `--fido-register-factor=2`.

Any authentication factors that proceed the FIDO registration must succeed before the registration continues.

The server checks the user account information to determine if the FIDO device requires registration. If the device must be registered, the server switches the client session to sandbox mode. The registration must be completed before any other activity. In this mode, only `ALTER USER` statements are permitted. If the session is started with `--fido-register-factor`, the client generates the statements required to register. After the registration is complete, the session is switched out of sandbox mode and the client can proceed as normal.

After the device is registered, the server updates the `mysql.user` system table for that account with the device registration status and stores the public key and credential ID.

The user must use the same FIDO device during registration and authentication. If the device is reset or the user attempts to use a different device, the authentication fails. To use a different device, the registered device must be unregistered and you must complete the registration process again.

### Use FIDO authentication as the only method

If FIDO is used as the only method of authentication, the method does not use a password. The authentication uses a method such as a biometric scan or a security key.

The user creates an account with the `PASSWORDLESS_USER_ADMIN` privilege and the `CREATE USER` privilege.

The first element of the `authentication_policy` value must be an asterisk(\*). Do not start with the plugin name. [Configuring the `authentication_policy` value] has more information.

You must include the `INITIAL AUTHENTICATION IDENTIFIED BY` clause in the `CREATE USER` statement. The server does not accept the statement without the clause but the account is unusable because the user cannot connect to the server to register the device.

The `CREATE USER` syntax is the following:

```
mysql> CREATE USER <username>@<hostname> IDENTIFIED WITH authentication_fido INITIAL
AUTHENTICATION IDENTIFIED BY '<password>';
```

During registration, the user must authenticate with the password. After the device is registered, the server deletes the password and modifies the account to make FIDO the only authentication method.

## 9.5.4 Unregister a FIDO device

If the FIDO device is replaced or lost, the following actions occur:

Action required	Who can perform the action
Unregister the previous device	The account owner or any user with access to the device
Register the new device	The user planning to use the device

Unregister a device with the following statement:

```
mysql> ALTER USER `username`@`hostname` {2|3} FACTOR UNREGISTER;
```

## 9.5.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-31

## 9.6 Encryption functions

Percona Server for MySQL adds encryption functions and variables to manage the encryption range. The functions may take an algorithm argument. Encryption converts plaintext into ciphertext using a key and an encryption algorithm.

You can also use the user-defined functions with the PEM format keys generated externally by the OpenSSL utility.

A digest uses plaintext and generates a hash value. This hash value can verify if the plaintext is unmodified. You can also sign or verify on digests to ensure that the original plaintext was not modified. You cannot decrypt the original text from the hash value.

When choosing key lengths, consider the following:

- Encryption strength increases with the key size and, also, the key generation time.
- If performance is important and the functions are frequently used, use symmetric encryption. Symmetric encryption functions are faster than asymmetric encryption functions. Moreover, asymmetric encryption has restrictions on the maximum length of a message being encrypted. For example, for RSA the algorithm maximum message size is the key length in bytes (key length in bits / 8) minus 11.

The following table and sections describe the functions. For examples, see function examples.

Function Name
<code>asymmetric_decrypt(algorithm, crypt_str, key_str)</code>
<code>asymmetric_derive(pub_key_str, priv_key_str)</code>
<code>asymmetric_encrypt(algorithm, str, key_str)</code>
<code>asymmetric_sign(algorithm, digest_str, priv_key_str, digest_type)</code>
<code>asymmetric_verify(algorithm, digest_str, sig_str, pub_key_str, digest_type)</code>
<code>create_asymmetric_priv_key(algorithm, (key_len   dh_parameters))</code>
<code>create_asymmetric_pub_key(algorithm, priv_key_str)</code>
<code>create_dh_parameters(key_len)</code>
<code>create_digest(digest_type, str)</code>

The following table describes the Encryption threshold variables which can be used to set the maximum value for a key length based on the type of encryption.

Variable Name
<code>encryption_udf.dh_bits_threshold</code>
<code>encryption_udf.dsa_bits_threshold</code>
<code>encryption_udf.rsa_bits_threshold</code>

## 9.6.1 Install component\_encryption\_udf

Use the `Install Component Statement` to add the `component_encryption_udf` component. The functions and variables are available. The user-defined functions and the Encryption threshold variables are auto-registered. There is no requirement to invoke `CREATE FUNCTION ... SONAME ...`.

The `INSERT` privilege on the `mysql.component` system table is required to run the `INSTALL COMPONENT` statement. To register the component, the operation adds a row to this table.

The following is an example of the installation command:

```
mysql> INSTALL COMPONENT 'file://component_encryption_udf';
```

 **Note**

If you are Compiling Percona Server for MySQL from Source, the Encryption UDF component is built by default when Percona Server for MySQL is built. Specify the `-DWITH_ENCRYPTION_UDF=OFF` cmake option to exclude it.

## 9.6.2 User-defined functions described

### 9.6.3 `asymmetric_decrypt(algorithm, crypt_str, key_str)`

Decrypts an encrypted string using the algorithm and a key string.

#### Returns

A plaintext as a string.

#### Parameters

The following are the function's parameters:

- `algorithm` - the encryption algorithm supports RSA to decrypt the string.
- `key_str` - a string in the PEM format. The key string must have the following attributes:
  - Valid
  - Public or private key string that corresponds with the private or public key string used with the `asymmetric_encrypt` function.

### 9.6.4 `asymmetric_derive(pub_key_str, priv_key_str)`

Derives a symmetric key using a public key generated on one side and a private key generated on another.

#### `asymmetric_derive` output

A key as a binary string.

#### `asymmetric_derive` parameters

The `pub_key_str` must be a public key in the PEM format and generated using the Diffie-Hellman (DH) algorithm.

The `priv_key_str` must be a private key in the PEM format and generated using the Diffie-Hellman (DH) algorithm.

### 9.6.5 `asymmetric_encrypt(algorithm, str, key_str)`

Encrypts a string using the algorithm and a key string.

**asymmetric\_encrypt output**

A ciphertext as a binary string.

**asymmetric\_encrypt parameters**

The parameters are the following:

- *algorithm* - the encryption algorithm supports RSA to encrypt the string.
- *str* - measured in bytes. The length of the string must not be greater than the *key\_str* modulus length in bytes - 11 (additional bytes used for PKCS1 padding)
- *key\_str* - a key (either private or public) in the PEM format

**9.6.6 asymmetric\_sign(*algorithm, digest\_str, priv\_key\_str, digest\_type*)**

Signs a digest string using a private key string.

**asymmetric\_sign output**

A signature is a binary string.

**asymmetric\_sign parameters**

The parameters are the following:

- algorithm - the encryption algorithm supports either RSA or DSA to encrypt the string.
- digest\_str - the digest binary string that is signed. Invoking create\_digest generates the digest.
- priv\_key\_str - the private key used to sign the digest string. The key must be in the PEM format.
- digest\_type - the OpenSSL version installed on your system determines the available hash functions. The following table lists these functions:

OpenSSL 1.0.2	OpenSSL 1.1.0	OpenSSL 1.1.1
md5	md5	md5
sha1	sha1	sha1
sha224	sha224	sha224
sha384	sha384	sha384
sha512	sha512	sha512
md4	md4	md4
sha	md5-sha1	md5-sha1
ripemd160	ripemd160	ripemd160
whirlpool	whirlpool	sha512-224
	blake2b512	sha512-256
	blake2s256	whirlpool
		sm3
		blake2b512
		blake2s256
		sha3-224
		sha3-384
		sha3-512
		shake128
		shake256

### 9.6.7 asymmetric\_verify(*algorithm*, *digest\_str*, *sig\_str*, *pub\_key\_str*, *digest\_type*)

Verifies whether the signature string matches the digest string.

**asymmetric\_verify output**

A 1 (success) or a 0 (failure).

**asymmetric\_verify parameters**

The parameters are the following:

- algorithm – supports either 'RSA' or 'DSA'.
- digest\_str – invoking create\_digest generates this digest binary string.
- sig\_str – the signature binary string. Invoking asymmetric\_sign generates this string.
- pub\_key\_str – the signer's public key string. This string must correspond to the private key passed to asymmetric\_sign to generate the signature string. The string must be in the PEM format.
- digest\_type – the supported values are listed in the digest type table of create\_digest

**9.6.8 create\_asymmetric\_priv\_key(algorithm, (key\_len | dh\_parameters))**

Generates a private key using the given algorithm and key length for RSA or DSA or Diffie-Hellman parameters for DH. For RSA or DSA, if needed, execute `KILL`

`[QUERY|CONNECTION] <id>` to terminate a long-lasting key generation. The DH key generation from existing parameters is a quick operation. Therefore, it does not make sense to terminate that operation with `KILL`.

**create\_asymmetric\_priv\_key output**

The key as a string in the PEM format.

**create\_asymmetric\_priv\_key parameters**

The parameters are the following:

- algorithm – the supported values are 'RSA', 'DSA', or 'DH'.
- key\_len – the supported key length values are the following:
  - RSA – the minimum length is 1,024. The maximum length is 16,384.
  - DSA – the minimum length is 1,024. The maximum length is 9,984.

**Note**

The key length limits are defined by OpenSSL. To change the maximum key length, use either `encryption_udf.rsa_bits_threshold` or `encryption_udf.dsa_bits_threshold`.

- dh\_parameters – Diffie-Hellman (DH) parameters. Invoking create\_dh\_parameter creates the DH parameters.

**9.6.9 create\_asymmetric\_pub\_key(algorithm, priv\_key\_str)**

Derives a public key from the given private key using the given algorithm.

**create\_asymmetric\_pub\_key output**

The key as a string in the PEM format.

**create\_asymmetric\_pub\_key parameters**

The parameters are the following:

- algorithm – the supported values are 'RSA', 'DSA', or 'DH'.
- priv\_key\_str – must be a valid key string in the PEM format.

**9.6.10 create\_dh\_parameters(key\_len)**

Creates parameters for generating a Diffie-Hellman (DH) private/public key pair. If needed, execute `KILL [QUERY|CONNECTION] <id>` to terminate the generation of long-lasting parameters.

Generating the DH parameters can take more time than generating the RSA keys or the DSA keys. OpenSSL defines the parameter length limits. To change the maximum parameter length, use `encryption_udf.dh_bits_threshold`.

**create\_dh\_parameters output**

A string in the PEM format and can be passed to `create_asymmetric_private_key`.

**create\_dh\_parameters parameters**

The parameters are the following:

- key\_len – the range for the key length is from 1024 to 10,000. The default value is 10,000.

**9.6.11 create\_digest(digest\_type, str)**

Creates a digest from the given string using the given digest type. The digest string can be used with `asymmetric_sign` and `asymmetric_verify`.

**create\_digest output**

The digest of the given string as a binary string

## create\_digest parameters

The parameters are the following:

- `digest_type` - the OpenSSL version installed on your system determines the available hash functions. The following table lists these functions:

OpenSSL 1.0.2	OpenSSL 1.1.0	OpenSSL 1.1.1
md5	md5	md5
sha1	sha1	sha1
sha224	sha224	sha224
sha384	sha384	sha384
sha512	sha512	sha512
md4	md4	md4
sha	md5-sha1	md5-sha1
ripemd160	ripemd160	ripemd160
whirlpool	whirlpool	sha512-224
	blake2b512	sha512-256
	blake2s256	whirlpool
		sm3
		blake2b512
		blake2s256
		sha3-224
		sha3-384
		sha3-512
		shake128
		shake256

- `str` - String used to generate the digest string.

## Encryption threshold variables

The maximum key length limits are defined by OpenSSL. Server administrators can limit the maximum key length using the encryption threshold variables.

The variables are automatically registered when `component_encryption_udf` is installed.

Variable Name
encryption_udf.dh_bits_threshold

**encryption\_udf.dh\_bits\_threshold**

The variable sets the maximum limit for the create\_dh\_parameters user-defined function and takes precedence over the OpenSSL maximum length value.

Option	Description
command-line	Yes
scope	Global
data type	unsigned integer
default	10000

The range for this variable is from 1024 to 10,000. The default value is 10,000.

**encryption\_udf.dsa\_bits\_threshold**

The variable sets the threshold limits for create\_asymmetric\_priv\_key user-defined function when the function is invoked with the DSA parameter and takes precedence over the OpenSSL maximum length value.

Option	Description
command-line	Yes
scope	Global
data type	unsigned integer
default	9984

The range for this variable is from 1,024 to 9,984. The default value is 9,984.

**encryption\_udf.rsa\_bits\_threshold**

The variable sets the threshold limits for the create\_asymmetric\_priv\_key user-defined function when the function is invoked with the RSA parameter and takes precedence over the OpenSSL maximum length value.

Option	Description
command-line	Yes
scope	Global
data type	unsigned integer
default	16384

The range for this variable is from 1,024 to 16,384. The default value is 16,384.

## Examples

Code examples for the following operations:

- set the threshold variables
- create a private key
- create a public key
- encrypt data
- decrypt data

```
-- Set Global variable
mysql> SET GLOBAL encryption_udf.dh_bits_threshold = 4096;

-- Set Global variable
mysql> SET GLOBAL encryption_udf.rsa_bits_threshold = 4096;
```

```
-- Create private key
mysql> SET @private_key = create_asymmetric_priv_key('RSA', 3072);

-- Create public key
mysql> SET @public_key = create_asymmetric_pub_key('RSA', @private_key);

-- Encrypt data using the private key (you can also use the public key)
mysql> SET @ciphertext = asymmetric_encrypt('RSA', 'This text is secret',
@private_key);

-- Decrypt data using the public key (you can also use the private key)
-- The decrypted value @plaintext should be identical to the original 'This text is
secret'
mysql> SET @plaintext = asymmetric_decrypt('RSA', @ciphertext, @public_key);
```

Code examples for the following operations:

- generate a digest string
- generate a digest signature
- verify the signature against the digest

```
-- Generate a digest string
mysql> SET @digest = create_digest('SHA256', 'This is the text for digest');

-- Generate a digest signature
mysql> SET @signature = asymmetric_sign('RSA', @digest, @private_key, 'SHA256');

-- Verify the signature against the digest
-- The @verify_signature must be equal to 1
mysql> SET @verify_signature = asymmetric_verify('RSA', @digest, @signature,
@public_key, 'SHA256');
```

Code examples for the following operations:

- generate a DH parameter
- generates two DH key pairs
- generate a symmetric key using the public\_1 and the private\_2
- generate a symmetric key using the public\_2 and the private\_1

```
-- Generate a DH parameter
mysql> SET @dh_parameter = create_dh_parameters(3072);

-- Generate DH key pairs
mysql> SET @private_1 = create_asymmetric_priv_key('DH', @dh_parameter);
mysql> SET @public_1 = create_asymmetric_pub_key('DH', @private_1);
mysql> SET @private_2 = create_asymmetric_priv_key('DH', @dh_parameter);
mysql> SET @public_2 = create_asymmetric_pub_key('DH', @private_2);

-- Generate a symmetric key using the public_1 and private_2
-- The @symmetric_1 must be identical to @symmetric_2
mysql> SET symmetric_1 = asymmetric_derive(@public_1, @private_2);

-- Generate a symmetric key using the public_2 and private_1
-- The @symmetric_2 must be identical to @symmetric_1
mysql> SET symmetric_2 = asymmetric_derive(@public_2, @private_1);
```

Code examples for the following operations:

- create a private key using a `SET` statement
- create a private key using a `SELECT` statement
- create a private key using an `INSERT` statement

```
mysql> SET @private_key1 = create_asymmetric_priv_key('RSA', 3072);
mysql> SELECT create_asymmetric_priv_key('RSA', 3072) INTO @private_key2;
mysql> INSERT INTO key_table VALUES(create_asymmetric_priv_key('RSA', 3072));
```

## 9.6.12 Uninstall component\_encryption\_udf

You can deactivate and uninstall the component using the [Uninstall Component](#) statement.

```
mysql> UNINSTALL COMPONENT 'file://component_encryption_udf';
```

### 9.6.13 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 9.7 Data-at-rest-encryption

### 9.7.1 Data at Rest Encryption

Data security is a concern for institutions and organizations. [Transparent Data Encryption \(TDE\)](#) or [Data at Rest Encryption](#) encrypts data files. Data at rest is any data that is not accessed or changed frequently, stored on different types of storage devices. Encryption ensures that if an unauthorized user accesses the data files from the file system, the user cannot read the contents.

If the user uses master key encryption, the MySQL keyring plugin stores the InnoDB master key, used for the master key encryption implemented by MySQL. The master key is also used to encrypt redo logs, and undo logs, along with the tablespaces.

The InnoDB tablespace encryption has the following components:

- The database instance has a master key for tablespaces and a master key for binary log encryption.
- Each tablespace has a tablespace key. The key is used to encrypt the Tablespace data pages. Encrypted tablespace keys are written on the tablespace header. In the master key implementation, the tablespace key cannot be changed unless you rebuild the table.

Two separate keys allow the master key to be rotated in a minimal operation. When the master key is rotated, each tablespace key is decrypted and re-encrypted with the new master key. The key rotation only reads and writes to the first page of each tablespace file (.ibd).

An InnoDB tablespace file is comprised of multiple logical and physical pages. Page 0 is the tablespace header page and keeps the metadata for the tablespace. The encryption information is stored on page 0 and the tablespace key is encrypted.

An encrypted page is decrypted at the I/O layer, added to the buffer pool, and used to access the data. A buffer pool page is not encrypted. The page is encrypted by the I/O layer before the page is flushed to disk.

#### Percona XtraBackup support

Percona XtraBackup version 8.1 supports the backup of encrypted general tablespaces.

Percona XtraBackup only supports features that are [Generally Available \(GA\)](#) in Percona Server for MySQL. Due to time constraints, a GA feature may be supported in a later Percona XtraBackup release. Review the [Percona XtraBackup release notes](#) for more information.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 9.7.2 Keyring components and plugins

### Keyring components and plugins overview

Percona Server supports a keyring that enables internal server components and plugins to store sensitive information securely for later retrieval.

#### Warning

Enable only one keyring plugin or one keyring component at a time for each server instance. Enabling multiple keyring plugins or keyring components or mixing keyring plugins or keyring components is not supported and may result in data loss.

Percona Server supports the following keyring components and plugins:

[Use the keyring file component or plugin](#) 

[Use the keyring vault component](#) 

[Use the Key Management Interoperability Protocol \(KMIP\)](#) 

[Use the Amazon Key Management Service \(AWS KMS\)](#) 

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[Community Forum](#)  [Get a Percona Expert](#) 

---

Last update: 2023-10-19

### Use the keyring file component or keyring file plugin

#### USE THE KEYRING FILE PLUGIN

Percona Server for MySQL supports the `keyring_file` plugin that stores the keyring data locally.

#### Warning

The `keyring_file` plugin should not be used for regulatory compliance.

To install the plugin, follow the [installing and uninstalling plugins](#) instructions.

### Load the keyring plugin

You should load the plugin at server startup with the `--early-plugin-load` option to enable keyrings.

We recommend that you load the plugin in the configuration file to facilitate recovery for encrypted tables. Also, the redo log encryption and the undo log encryption cannot be used without `--early-plugin-load`. The normal plugin load happens too late at startup.

To load the `keyring_file` plugin, run the following command:

```
$ mysqld --early-plugin-load="keyring_file=keyring_file.so"
```

If a server starts with different plugins loaded early, the `--early-plugin-load` option should contain the plugin names in a double-quoted list with each plugin name separated by a semicolon. The use of double quotes ensures the semicolons do not create issues when the list is executed in a script.

### USE THE KEYPING FILE COMPONENT

The `keyring_file` component is part of the component-based MySQL infrastructure which extends the server capabilities.

See the MySQL documentation on the [component installation](#) and on the [keyring\\_file component usage](#) for more information.

The component must be installed with a manifest. A keyring component is not loaded with the `--early-plugin-load` option on the server. The server uses a manifest and the component consults its configuration file during initialization. You should only load a keyring component with a manifest file. Do not use the `INSTALL_COMPONENT` statement, which loads the keyring components too late in the startup sequence of the server. For example, `InnoDB` requires the component, but because the components are registered in the `mysql.component` table, this table is loaded after `InnoDB` initialization.

You should create a global manifest file named `mysqld.my` in the installation directory and, optionally, create a local manifest file, also named `mysqld.my` in a data directory.

To install a keyring component, do the following:

1. Write a manifest in a valid JSON format
2. Write a configuration file

A manifest file indicates which component to load. If the manifest file does not exist, the server does not load the component associated with that file. During startup, the server reads the global manifest file from the installation directory. The global manifest file can contain the required information or point to a local manifest file located in the data directory. If you have multiple server instances that use different keyring components use a local manifest file in each data directory to load the correct keyring component for that instance.

#### Warning

Enable only one keyring plugin or one keyring component at a time for each server instance. Enabling multiple keyring plugins or keyring components or mixing keyring plugins or keyring components is not supported and may result in data loss.

An example of a manifest and a configuration file is the following:

An example of `./bin/mysqld.my`:

```
{
  "components": "file://component_keyring_file"
}
```

An example of `/lib/plugin/component_keyring_file.cnf`:

```
{
  "path": "/var/lib/mysql-keyring/keyring_file", "read_only": false
}
```

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-19

## Use the keyring vault component

The `keyring_vault` component extends the server capabilities and provides an interface for the database with a [HashiCorp Vault](#) server to store key and secure encryption keys.

The component must be installed with a manifest. A keyring component is not loaded with the `--early-plugin-load` option on the server. The server uses a manifest and the component consults its configuration file during initialization. You should only load a keyring component with a manifest file. Do not use the `INSTALL_COMPONENT` statement, which loads the keyring components too late in the startup sequence of the server. For example, `InnoDB` requires the component, but because the components are registered in the `mysql.component` table, this table is loaded after `InnoDB` initialization.

You should create a global manifest file named `mysqld.my` in the installation directory and, optionally, create a local manifest file, also named `mysqld.my` in a data directory.

To install a keyring component, do the following:

1. Write a manifest in a valid JSON format
2. Write a configuration file

A manifest file indicates which component to load. If the manifest file does not exist, the server does not load the component associated with that file. During startup, the server reads the global manifest file from the installation directory. The global manifest file can contain the required information or point to a local manifest file located in the data directory. If you have multiple server instances that use different keyring components use a local manifest file in each data directory to load the correct keyring component for that instance.

**Warning**

Enable only one keyring plugin or one keyring component at a time for each server instance. Enabling multiple keyring plugins or keyring components or mixing keyring plugins or keyring components is not supported and may result in data loss.

The following example is a global manifest file that does not use local manifests:

```
{
  "read_local_manifest": false,
  "components": "file:///component_keyring_vault"
}
```

The following is an example of a global manifest file that points to a local manifest file:

```
{
  "read_local_manifest": true
}
```

The following is an example of a local manifest file:

```
{
  "components": "file:///component_keyring_vault"
}
```

The configuration settings are either in a global configuration file or a local configuration file.

The `component_keyring_vault.cnf` file contains the following information:

- `read_local_config` [optional] - this option can be used only in the global configuration file. If you use this option, it indicates whether the component should read configuration information from the local configuration file. The allowed values are `true` or `false`. If you do not use this option, the component uses only the global configuration file.

If you use the `read_local_config` option in the global configuration file along with other items, the component checks the `read_local_config` item value first:

`false` - the component processes other items in the global configuration file and ignores the local configuration file.

`true` - the component ignores other items in the global configuration file and attempts to read the local configuration file.

- `timeout` - the duration in seconds for the Vault server connection timeout. The default value is 15. The allowed range is from 0 to 86400. The timeout can be also disabled to wait an infinite amount of time by setting this variable to 0.
- `vault_url` - the Vault server address.
- `secret_mount_point` - the mount point name where the `keyring_vault` stores the keys.
- `token` - a token generated by the Vault server.
- `secret_mount_point_version` [optional] - the KV Secrets Engine version (`kv` or `kv-v2`) used. The allowed values are `AUTO`, `1`, and `2`. The default value is `AUTO`.
- `vault_ca` [optional] - if the machine does not trust the Vault's CA certificate, this variable points to the CA certificate used to sign the Vault's certificates.

#### Example of a configuration file in JSON format

```
{
  "timeout": 15,
  "vault_url": "https://vault.public.com:8202",
  "secret_mount_point": "secret",
  "secret_mount_point_version": "AUTO",
  "token": "58a20c08-8001-fd5f-5192-7498a48eaf20",
  "vault_ca": "/data/keyring_vault_confs/vault_ca.crt"
}
```

#### Warning

Each `secret_mount_point` must be used by only one server. If multiple servers use the same `secret_mount_point`, the behavior is unpredictable.

The first time a key is fetched from a keyring, the `keyring_vault` communicates with the Vault server to retrieve the key type and data.

#### `secret_mount_point_version` information

The `secret_mount_point_version` can be either a `1`, `2`, `AUTO`, or the `secret_mount_point_version` parameter is not listed in the configuration file.

Value	Description
1	Works with KV Secrets Engine - Version 1 (kv). If the <code>secret_mount_point</code> is based on URLs, the <code>secret_mount_point</code> is always a URL. For example, to return a key named <code>skey</code> .
2	Works with KV Secrets Engine - Version 2 (kv-v2). The <code>secret_mount_point</code> parameter in the configuration file is a path. <ul style="list-style-type: none"> <li>The <code>mount_point_path</code> - the mount point path where the secret was created</li> <li>The <code>directory_path</code> - a virtual directory path for the secret. It is used to create virtual namespaces with the same name as the <code>mount_point_path</code>.</li> </ul> For example, both the <code>mount_point_path</code> and <code>directory_path</code> are <code>/v1/</code> to form key access URLs: <code>/v1/&lt;mount_point_path&gt;/&lt;secret_name&gt;</code> .
AUTO	An autodetection mechanism probes the <code>secret_mount_point</code> version is <code>kv</code> or <code>kv-v2</code> and based on the <code>secret_mount_point</code> as is, or split the <code>secret_mount_point</code> into <code>mount_point_path</code> and <code>directory_path</code> .
Not listed	If the <code>secret_mount_point_version</code> is not set, the behavior is the same as <code>AUTO</code> .

If you set the `secret_mount_point_version` to 2 but the path pointed by `secret_mount_point` is based on KV Secrets Engine - Version 1 (kv), an error is reported and the component fails to initialize.

If you set the `secret_mount_point_version` to 1 but the path pointed by `secret_mount_point` is based on KV Secrets Engine - Version 2 (kv-v2), the component initialization succeeds but any MySQL keyring-related operations fail.

#### Upgrade from Vault Secrets Engine Version 1 to Version 2

You can upgrade from the Vault Secrets Engine Version 1 to Version 2.

Use either of the following methods:

- Set the `secret_mount_point_version` to `AUTO` or the variable is not set in the `keyring_vault` component configuration files in all Percona Servers. The `AUTO` value ensures the autodetection mechanism is invoked during the component initialization.
- Set the `secret_mount_point_version` to 2 to ensure that components do not initialize unless the `kv` to `kv-v2` upgrade completes.

#### Note

The `keyring_vault` component that works with `kv-v2` secret engines does not use the built-in key versioning capabilities. The keyring key versions are encoded into key names.

#### See also

[Hashicorp Documentation: Installing Vault](#)

[Hashicorp Documentation: Production Hardening](#)

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)    [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-19

## Use the Key Management Interoperability Protocol (KMIP)

Percona Server for MySQL supports the [OASIS Key Management Interoperability Protocol \(KMIP\)](#). This implementation was tested with the [PyKMIP server](#) and the [HashiCorp Vault Enterprise KMIP Secrets Engine](#).

KMIP enables communication between key management systems and the database server. The protocol can do the following:

- Streamline encryption key management
- Eliminate redundant key management processes

### COMPONENT INSTALLATION

The component must be installed with a manifest. A keyring component is not loaded with the `--early-plugin-load` option on the server. The server uses a manifest and the component consults its configuration file during initialization. You should only load a keyring component with a manifest file. Do not use the `INSTALL_COMPONENT` statement, which loads the keyring components too late in the startup sequence of the server. For example, `InnoDB` requires the component, but because the components are registered in the `mysql.component` table, this table is loaded after `InnoDB` initialization.

You should create a global manifest file named `mysqld.my` in the installation directory and, optionally, create a local manifest file, also named `mysqld.my` in a data directory.

To install a keyring component, do the following:

1. Write a manifest in a valid JSON format
2. Write a configuration file

A manifest file indicates which component to load. If the manifest file does not exist, the server does not load the component associated with that file. During startup, the server reads the global manifest file from the installation directory. The global manifest file can contain the required information or point to a local manifest file located in the data directory. If you have multiple server instances that use different keyring components use a local manifest file in each data directory to load the correct keyring component for that instance.

**Warning**

Enable only one keyring plugin or one keyring component at a time for each server instance. Enabling multiple keyring plugins or keyring components or mixing keyring plugins or keyring components is not supported and may result in data loss.

For more information, see [Installing and Uninstalling Components](#).

The following is an example of a global manifest file that does not use local manifests:

```
{
  "read_local_manifest": false,
  "components": "file://component_keyring_kmip"
}
```

The following is an example of a global manifest file that points to a local manifest file:

```
{
  "read_local_manifest": true
}
```

The following is an example of a local manifest file:

```
{
  "components": "file://component_keyring_kmip"
}
```

The configuration settings are either in a global configuration file or a local configuration file. The settings are the same.

**Example of a configuration file in JSON format**

```
{
  "server_addr": "127.0.0.1",
  "server_port": "5696",
  "client_ca": "client_certificate.pem",
  "client_key": "client_key.pem",
  "server_ca": "root_certificate.pem"
}
```

For more information, see [Keyring Component installation](#).

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-19

## Use the Amazon Key Management Service (AWS KMS)

Percona Server for MySQL supports the [Amazon Key Management Service \(AWS KMS\)](#). Percona Server generates the keyring keys. Amazon Web Services (AWS) encrypts the keyring data.

The AWS KMS lets you create and manage cryptographic keys across AWS services. For more information, see the [AWS Key Management Service Documentation](#).

To use the AWS KMS component, do the following:

- Have an AWS user account. This account has an access key and a secret key.
- Create a KMS key ID. The KMS key can then be referenced in the configuration either by its ID, alias (the key can have any number of aliases), or ARN.

### COMPONENT INSTALLATION

The component must be installed with a manifest. A keyring component is not loaded with the `--early-plugin-load` option on the server. The server uses a manifest and the component consults its configuration file during initialization. You should only load a keyring component with a manifest file. Do not use the `INSTALL_COMPONENT` statement, which loads the keyring components too late in the startup sequence of the server. For example, `InnoDB` requires the component, but because the components are registered in the `mysql.component` table, this table is loaded after `InnoDB` initialization.

You should create a global manifest file named `mysqld.my` in the installation directory and, optionally, create a local manifest file, also named `mysqld.my` in a data directory.

To install a keyring component, do the following:

1. Write a manifest in a valid JSON format
2. Write a configuration file

A manifest file indicates which component to load. If the manifest file does not exist, the server does not load the component associated with that file. During startup, the server reads the global manifest file from the installation directory. The global manifest file can contain the required information or point to a local manifest file located in the data directory. If you have multiple server instances that use different keyring components use a local manifest file in each data directory to load the correct keyring component for that instance.

#### Warning

Enable only one keyring plugin or one keyring component at a time for each server instance. Enabling multiple keyring plugins or keyring components or mixing keyring plugins or keyring components is not supported and may result in data loss.

For more information, see [Installing and Uninstalling Components](#).

The following example is a global manifest file that does not use local manifests:

```
{
  "read_local_manifest": false,
```

```
"components": "file:///component_keyring_kmip"
}
```

The following is an example of a global manifest file that points to a local manifest file:

```
{
  "read_local_manifest": true
}
```

The following is an example of a local manifest file:

```
{
  "components": "file:///component_keyring_kmip"
}
```

The configuration settings are either in a global configuration file or a local configuration file. The settings are the same.

The KMS configuration file has the following options:

- `read_local_config`
- `path` - the location of the JSON keyring database file.
- `read_only` - if true, the keyring cannot be modified.
- `kms_key` - the identifier of an AWS KMS master key. This key must be created by the user before creating the manifest file. The identifier can be one of the following:
  - UUID
  - Alias
  - ARN

For more information, see [Finding the key ID and key ARN](#).

- `region` - the AWS where the KMS is stored. Any HTTP request connect to this region.
- `auth_key` - an AWS user authentication key. The user must have access to the KMS key.
- `secret_access_key` - the secret key (API "password") for the AWS user.

#### Note

The configuration file contains authentication information. Only the MySQL process should be able to read this file.

### Example of a configuration file in JSON format

```
{
  "read_local_config": "true/false",
  "path": "/usr/local/mysql/keyring-mysql/aws-keyring-data",
  "region": "eu-central-1",
  "kms_key": "UUID, alias or ARN as displayed by the KMS console",
  "auth_key": "AWS user key",
  "secret_access_key": "AWS user secret key"
}
```

For more information, see [Keyring Component installation](#).

GET EXPERT HELP

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-19

## Encrypt File-Per-Table Tablespace

The `file_per_table` tablespace inherits the default schema encryption setting. Use the `ENCRYPTION` clause in the `CREATE TABLE` statement to explicitly set the encryption.

```
mysql> CREATE TABLE ... ENCRYPTION='Y';
```

To change the encryption setting for an existing `file_per_table` tablespace, add the `ENCRYPTION` clause. The `ALTER TABLE` statement without the `ENCRYPTION` clause does not change the encryption state.

```
mysql> ALTER TABLE ... ENCRYPTION='Y';
```

If the `table_encryption_privilege_check` is enabled, the `TABLE_ENCRYPTION_ADMIN` privilege is required to change the `file_per_table` encryption setting from the default schema encryption.

GET EXPERT HELP

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-01

## Encrypt schema or general tablespace

Percona Server for MySQL uses the same encryption architecture as MySQL, a two-tier system consisting of a master key and tablespace keys. The master key can be changed, or rotated in the keyring, as needed. Each tablespace key, when decrypted, remains the same.

The feature requires the keyring plugin.

### SET THE DEFAULT FOR SCHEMAS AND GENERAL TABLESPACE ENCRYPTION

The tables in a general tablespace are either all encrypted or all unencrypted. A tablespace cannot contain a mixture of encrypted tables and unencrypted tables.

The encryption of a schema or a general tablespace is determined by the `default_table_encryption` variable unless you specify the `ENCRYPTION` clause in the `CREATE SCHEMA` or `CREATE TABLESPACE` statement.

You can set the `default_table_encryption` variable in an individual connection.

```
mysql> SET default_table_encryption=ON;
```

### DEFAULT\_TABLE\_ENCRYPTION

Option	Description
Command-line	default-table-encryption
Scope	Session
Dynamic	Yes
Data type	Text
Default	OFF

Defines the default encryption setting for schemas and general tablespaces. The variable allows you to create or alter schemas or tablespaces without specifying the `ENCRYPTION` clause. The default encryption setting applies only to schemas and general tablespaces and is not applied to the MySQL system tablespace.

The variable has the following possible options:

Value	Description
ON	New tables are encrypted. Add <code>ENCRYPTION</code> clause to <code>CREATE TABLE</code> statement to create unencrypted tables.
OFF	By default, new tables are unencrypted. Add <code>ENCRYPTION</code> clause to <code>CREATE TABLE</code> or <code>ALTER TABLE</code> statement to create encrypted tables.

#### Note

The `ALTER TABLE` statement changes the current encryption mode only if you use the `ENCRYPTION` clause.

`INNODB_ENCRYPT_ONLINE_ALTER_LOGS`

Option	Description
Command-line	<code>--innodb-encrypt-online-alter-logs</code>
Scope	Global
Dynamic	Yes
Data type	Boolean
Default	OFF

This variable simultaneously turns on the encryption of files used by InnoDB for full-text search using parallel sorting, building indexes using merge sort, and online DDL logs created by InnoDB for online DDL. Encryption is available for file merges used in queries and backend processes.

## USE ENCRYPTION

If you do not set the default encryption setting, you can create general tablespaces with the `ENCRYPTION` setting.

```
mysql> CREATE TABLESPACE tablespace_name ENCRYPTION='Y';
```

All tables contained in the tablespace are either encrypted or not encrypted. You cannot encrypt only some of the tables in a general tablespace. This feature extends the `CREATE TABLESPACE` statement to accept the `ENCRYPTION='Y/N'` option.

The option is a tablespace attribute and is not allowed with the `CREATE TABLE` or `SHOW CREATE TABLE` statement except with file-per-table tablespaces.

In an encrypted general tablespace, an attempt to create an unencrypted table generates the following error:

```
mysql> CREATE TABLE t3 (a INT, b TEXT) TABLESPACE foo ENCRYPTION='N';
```

 **Expected output**

```
ERROR 1478 (HY000): InnoDB: Tablespace 'foo' can contain only ENCRYPTED tables.
```

The server diagnoses an attempt to create or move tables, including partitioned ones, to a general tablespace with an incompatible encryption setting and aborts the process.

If you must move tables between incompatible tablespaces, create tables with the same structure in another tablespace and run `INSERT INTO SELECT` from each of the source tables into the destination tables.

## EXPORT AN ENCRYPTED GENERAL TABLESPACE

You can only export encrypted file-per-table tablespaces

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-09-01

### Encrypt system tablespace

By default, the system tablespace, which contains the system database and the data dictionary tables, is unencrypted. To change the encryption requires the `CREATE TABLESPACE` privilege on all tables in the instance.

In an `ALTER TABLESPACE` statement, add the `ENCRYPTION` option with the tablespace name to enable encryption.

```
mysql> ALTER TABLESPACE mysql ENCRYPTION='Y';
```

Disable encryption by setting the `ENCRYPTION` option to 'N'.

```
mysql> ALTER TABLESPACE mysql ENCRYPTION='N';
```

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-09-01

### Encrypt doublewrite file pages

InnoDB encrypts doublewrite file pages associated with encrypted tablespaces automatically. Doublewrite files can contain the following page types:

- Unencrypted
- Uncompressed
- Encrypted
- Compressed

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-09-01

## Encrypt temporary files

For InnoDB user-created temporary tables are created in a temporary tablespace file and use the `innodb_temp_tablespace_encrypt` variable.

The `CREATE TEMPORARY TABLE` does not support the `ENCRYPTION` clause. The `TABLESPACE` clause cannot be set to `innodb_temporary`.

The global temporary tablespace datafile `ibtmp1` contains the temporary table undo logs while intrinsic temporary tables and user-created temporary tables are located in the encrypted session temporary tablespace.

To create new temporary tablespaces unencrypted, the following variables must be set to `OFF` at runtime:

- `innodb_temp_tablespace_encrypt`
- `default_table_encryption`

Any existing encrypted user-created temporary files and intrinsic temporary tables remain in an encrypted session.

Temporary tables are only destroyed when the session is disconnected.

The `default_table_encryption` setting in `my.cnf` determines if a temporary table is encrypted.

If the `innodb_temp_tablespace_encrypt` = "OFF" and the `default_table_encryption` = "ON", the user-created temporary tables are encrypted. The temporary tablespace data file `ibtmp1`, which contains undo logs, is not encrypted.

If the `innodb_temp_tablespace_encrypt` is "ON" for the system tablespace, InnoDB generates an encryption key and encrypts the system's temporary tablespace. If you reset the encryption to "OFF", all subsequent pages are written to an unencrypted tablespace. Any generated keys are not erased to allow encrypted tables and undo data to be decrypted.

For each temporary file, an encryption key has the following attributes:

- Generated locally
- Maintained in memory for the lifetime of the temporary file
- Discarded with the temporary file

**ENCRYPT\_TMP\_FILES**

Option	Description
Command-line	<code>--encrypt_tmp_files</code>
Scope	Global
Dynamic	No
Data type	Boolean
Default	OFF

This variable turns “ON” the encryption of temporary files created by the *Percona Server for MySQL*. The default value is `OFF`.

**INNODB\_TEMP\_TABLESPACE\_ENCRYPT**

Option	Description
Command-line	<code>innodb-temp-tablespace-encrypt</code>
Scope	Global
Dynamic	Yes
Data type	Boolean
Default	OFF

When this variable is set to `ON`, the server encrypts the global temporary tablespace and has the `.ibtmp1` file extension and the session temporary tablespace and has the `.ibt` file extension.

The variable does not enforce the encryption of currently open temporary files and does not rebuild the system’s temporary tablespace to encrypt data that has already been written.

## GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-01

**Log encryption**

Describes the redo log encryption and the undo log encryption.

**REDO LOG ENCRYPTION**

Use the `innodb_redo_log_encrypt` option to enable or disable redo log data encryption. By default, the encryption of the redo log is disabled.

InnoDB uses the tablespace encryption key to encrypt the redo log data. If the encryption is enabled, when the server encrypts and writes the redo log data to the disk. When the server reads the redo log data from disk, the data is decrypted.

Changing the encryption does not change existing redo log pages. Setting the option to `ON`, any existing redo log pages remain unencrypted; writing new pages to disk encrypts them. Setting the option to `OFF`, any existing encrypted pages remain encrypted; writing new pages to disk are unencrypted.

The metadata for the redo log encryption includes the tablespace encryption key and is stored in the redo log file header. Removing the encryption metadata disables the redo log encryption.

Without the keyring plugin or component or the encryption key, a normal restart is not possible. InnoDB scans the redo log pages during startup. If the encryption options are not available, InnoDB cannot scan these pages. A forced startup without the redo logs is possible.

#### OPTION

##### `innodb_redo_log_encrypt`

Variable	Description
Command-line	<code>--innodb-redo-log-encrypt[= {ON   OFF}]</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

  

Option	Description
ON	This option is a compatibility alias for <code>OFF</code> ; existing redo log pages remain unencrypted; new pages are encrypted.
OFF	Any existing encrypted pages remain encrypted; writing new pages to disk are unencrypted.

Determines the encryption for the table redo log data. The default option for the variable is `OFF`.

#### UNDO LOG ENCRYPTION

Use the `innodb_undo_log_encrypt` option to enable or disable undo log data encryption. By default, the option to encrypt the undo log data is disabled.

InnoDB uses the tablespace encryption key to encrypt the undo log data. If the encryption is enabled, when the server encrypts and writes the undo log data to the disk. When the server reads the undo log data from disk, the data is decrypted.

Changing the encryption does not change existing undo log pages. Setting the option to `ON`, any existing pages remain unencrypted; writing new pages to disk encrypts them. Setting the option to `OFF`, any existing encrypted pages remain encrypted; writing new pages to disk are unencrypted.

The metadata for the redo log encryption includes the tablespace encryption key and is stored in the undo log file header.

The server requires the keyring plugin or keyring component used to encrypt log data until that data is truncated, even if the current option setting is `OFF`. When the undo tablespace is truncated, the encryption header is removed.

**INNODB\_UNDO\_LOG\_ENCRYPT**

Option	Description
Command-line	<code>--innodb-undo-log-encrypt[= {ON   OFF}]</code>
Scope	Global
Dynamic	Yes
Data type	Boolean
Default	OFF

This system variable defines the encryption status for the undo log data. The default setting is `OFF`, which disables the encryption.

## GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-01

**Rotate the master encryption key**

Rotate the master encryption key periodically and if the key has been compromised.

Rotating the master encryption key changes that key and tablespace keys are re-encrypted and updated in the tablespace headers. The rotation only succeeds if all operations are successful. If the rotation is interrupted, the operation is rolled forward when the server restarts.

The rotation operation does not affect tablespace data. To change a tablespace key, disable and then re-enable encryption for that tablespace.

The `ENCRYPTION_KEY_ADMIN` privilege is required to rotate the master encryption key.

InnoDB reads the encryption data from the tablespace header, if certain tablespace keys have been encrypted with the prior master key, InnoDB retrieves the master key from the keyring to decrypt the tablespace key. InnoDB re-encrypts the tablespace key with the new Master key.

Rotate the master encryption key with following statement:

```
mysql> ALTER INSTANCE ROTATE INNODB MASTER KEY;
```

The rotation operation must complete before any tablespace encryption operation can begin.

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#) [:percona–logo: Get a Percona Expert](#)

Last update: 2023-09-01

## Verify the encryption for tables, tablespaces, and schemas

If a general tablespace contains tables, check the table information to see if the table is encrypted. When the general tablespace contains no tables, you may verify if the tablespace is encrypted or not.

For single tablespaces, verify the ENCRYPTION option using INFORMATION\_SCHEMA.TABLES and the CREATE OPTIONS settings.

```
mysql> SELECT TABLE_SCHEMA, TABLE_NAME, CREATE_OPTIONS FROM
        INFORMATION_SCHEMA.TABLES WHERE CREATE_OPTIONS LIKE '%ENCRYPTION%';
```

### Expected output

TABLE_SCHEMA	TABLE_NAME	CREATE_OPTIONS
sample	t1	ENCRYPTION="Y"

A `flag` field in the `INFORMATION_SCHEMA.INNODB_TABLESPACES` has bit number 13 set if the tablespace is encrypted. This bit can be checked with the `flag & 8192` expression in the following way:

```
SELECT space, name, flag, (flag & 8192) != 0 AS encrypted FROM
        INFORMATION_SCHEMA.INNODB_TABLESPACES WHERE name in ('foo', 'test/t2', 'bar',
        'noencrypt');
```

The encrypted table metadata is contained in the `INFORMATION_SCHEMA.INNODB_TABLESPACES_ENCRYPTION` table. You must have the `Process` privilege to view the table information.

### Note

**This table is in tech preview and may change in future releases.**

```
mysql> DESCRIBE INNODB_TABLESPACES_ENCRYPTION;
```

**Expected output**

Field	Type	Null	Key	Default	Extra
SPACE	int(11) unsigned	NO			
NAME	varchar(655)	YES			
ENCRYPTION_SCHEME	int(11) unsigned	NO			
KEYSERVER_REQUESTS	int(11) unsigned	NO			
MIN_KEY_VERSION	int(11) unsigned	NO			
CURRENT_KEY_VERSION	int(11) unsigned	NO			
KEY_ROTATION_PAGE_NUMBER	bigint(21) unsigned	YES			
KEY_ROTATION_MAX_PAGE_NUMBER	bigint(21) unsigned	YES			
CURRENT_KEY_ID	int(11) unsigned	NO			
ROTATING_OR_FLUSHING	int(1) unsigned	NO			

To identify encryption-enabled schemas, query the INFORMATION\_SCHEMA.SCHEMATA table:

```
mysql> SELECT SCHEMA_NAME, DEFAULT_ENCRYPTION FROM
INFORMATION_SCHEMA.SCHEMATA WHERE DEFAULT_ENCRYPTION='YES';
```

**Expected output**

SCHEMA_NAME	DEFAULT_ENCRYPTION
samples	YES

The `SHOW CREATE SCHEMA` statement returns the `DEFAULT ENCRYPTION` clause.

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-01

# 10. Replicate

## 10.1 Manage group replication flow control

In replication, flow control prevents one member from falling too far behind the cluster and avoids excessive buffering. A cluster is not required to keep members in sync together for replication. The pending transactions in the relay log only increase for the lagging replica. Each member sends statistics to the group.

Flow control sets a threshold on the queue for transactions waiting in the certification queue or the transactions waiting in the applier queue. If the thresholds are exceeded, and during the duration that they are exceeded, flow control adjusts the writer members to the capacity of the delayed member. This action ensures that all members are in sync.

Flow controls work asynchronously and depend on the following:

- Monitoring the throughput and queue sizes of each member
- Throttling members to avoid writing beyond the capacity available

The following system variables set flow control behavior for Group Replication:

- [group\\_replication\\_flow\\_control\\_mode](#)
- [group\\_replication\\_flow\\_control\\_certifier\\_threshold](#)
- [group\\_replication\\_flow\\_control\\_applier\\_threshold](#)

Flow control is enabled and disabled by selecting a value in the `group_replication_flow_control_mode` variable. Flow control can also be enabled on the certifier or applier level or both and sets the threshold level.

### 10.1.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 10.2 Group replication system variables

variable name
<code>group_replication_auto_evict_timeout</code>
<code>group_replication_certification_loop_chunk_size</code>
<code>group_replication_certification_loop_sleep_time</code>

variable name
group_replication_flow_control_mode

## 10.2.1 group\_replication\_auto\_evict\_timeout

The variable is in [tech preview](#) mode. Before using the variable in production, we recommend that you test restoring production from physical backups in your environment, and also use the alternative backup method for redundancy.

Option	Description
Command-line	<code>-group-replication-auto-evict-timeout</code>
Dynamic	Yes
Scope	Global
Type	Integer
Default value	0
Maximum Value	65535
Unit	seconds

The value can be changed while Group Replication is running. The change takes effect immediately. Every node in the group can have a different timeout value, but, to avoid unexpected exits, we recommend that all nodes have the same value.

The variable specifies a period of time in seconds before a node is automatically evicted if the node exceeds the [flow control threshold](#). The default value is 0, which disables the eviction. To set the timeout, change the value with a number higher than zero.

In single-primary mode, the primary server ignores the timeout.

## 10.2.2 group\_replication\_certification\_loop\_chunk\_size

Option	Description
Command-line	<code>-group-replication-certification-loop-chunk-size</code>
Dynamic	Yes
Scope	Global
Data type	ulong
Default value	0

Defines the size of the chunk that must be processed during the certifier garbage collection phase after which the client transactions are allowed to interleave. The default value is 0.

The minimum value is 0. The maximum value is 4294967295.

### 10.2.3 group\_replication\_certification\_loop\_sleep\_time

Option	Description
Command-line	<code>-group-replication-certification-loop</code>
Dynamic	Yes
Scope	Global
Data type	ulong
Default value	0

Defines the sleep time in microseconds that the certifier garbage collection loop allows client transactions to interleave. The default value is 0.

The minimum value is 0. The maximum value is 1000000.

### 10.2.4 group\_replication\_flow\_control\_mode

Option	Description
Command-line	<code>-group_replication_flow_control_m</code>
Dynamic	Yes
Scope	Global
Data type	Enumeration
Default value	Quota
Valid values	DISABLED QUOTA MAJORITY

The “MAJORITY” value is in [tech preview](#) mode. Before using the variable in production, we recommend that you test restoring production from physical backups in your environment, and also use the alternative backup method for redundancy.

The variable specifies the mode use for flow control.

Percona Server for MySQL adds the “MAJORITY” value to the [group\\_replication\\_flow\\_control\\_mode](#) variable. In “MAJORITY” mode, [flow control](#) is activated only if the majority, more than half the number of members, exceed the flow control threshold. The other values are not changed.

### 10.2.5 Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 10.3 Replication Security

### 10.3.1 Encrypt binary log files and relay log files

Encrypt the binary log files and the relay log files to protect them from unauthorized viewing. The encryption uses the Advanced Encryption Standard (AES) symmetric block cipher algorithm. Instead of bits, this algorithm works on bytes.

Any supported encryption keyring plugin or component must be installed and configured to use encryption.

Enabling binary log encryption at system start, the server generates a binary log encryption key before initializing the binary log and relay logs. The key encrypts a file password for each binary log, if binary logging is enabled, and relay log. Keys generated from the file passwords encrypt the data in the files.

When used by the server, the binary log encryption key is called the binary log master key. This key can be rotated as needed and only the file password for each file is re-encrypted.

The binary log index file and relay log index file is never encrypted. Relay log files are encrypted for all channels.

To review if a binary log file is encrypted or not, use the `SHOW BINARY LOGS` statement.

If the server is running, the `BINLOG_ENCRYPTION_ADMIN` privilege is required to enable or disable encryption.

#### `binlog_encryption`

Option	Description
Command-line	<code>--binlog-encryption[= {ON \   OFF}]</code>
Scope	Global
Dynamic	Yes
Data type	Boolean
Default	OFF

This system variable enables binary log file encryption and relay log file encryption on the server. The default value is `OFF`. You can enable encryption for relay log files on a replica without a binary log.

If you set the `binlog_encryption` to `OFF`, the server immediately rotates the binary log file and relay log files and all logging is not encrypted. For any previously encrypted files, the server can still read them and they are not decrypted.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[Community Forum](#)

[Get a Percona Expert](#)

---

Last update: 2023-09-01

# 11. Monitor

## 11.1 Jemalloc memory allocation profiling

Percona Server for MySQL can take advantage of the memory-profiling ability of the jemalloc allocator. This ability provides a method to investigate memory-related issues.

### 11.1.1 Requirements

This memory-profiling requires `jemalloc_detected`. This read-only variable returns `true` if jemalloc with the profiling-enabled option is being used by *Percona Server for MySQL*.

As root, customize jemalloc with the following flags:

Option	Description
<code>--enable-stats</code>	Enables statistics-gathering ability
<code>--enable-prof</code>	Enables heap profiling and the ability

Using `LD_PRELOAD`. Build the library, configure the malloc configuration with the `prof:true` string, and then use `LD_PRELOAD` to preload the `libjemalloc.so` library. The `libprocess` `MemoryProfiler` class detects the library automatically and enables the profiling support.

The following is an example of the required commands:

```
./configure --enable-stats --enable-prof && make && make install
MALLOC_CONF=prof:true
LD_PRELOAD=/usr/lib/libjemalloc.so
```

### 11.1.2 Use Percona Server for MySQL with jemalloc with profiling enabled

To detect if jemalloc is set, run the following command:

```
SELECT @@jemalloc_detected;
```

To enable jemalloc profiling in a MySQL client, run the following command:

```
set global jemalloc_profiling=on;
```

The `malloc_stats_totals` table returns the statistics, in bytes, of the memory usage. The command takes no parameters and returns the results as a table.

The following example commands display this result:

```
use performance_schema;
```

```
SELECT * FROM malloc_stats_totals;
```

**Expected output**

```
+-----+-----+-----+-----+-----+
| id | ALLOCATION | MAPPED   | RESIDENT | RETAINED | METADATA |
+-----+-----+-----+-----+-----+
| 1 | 390977528 | 405291008 | 520167424 | 436813824 | 9933744 |
+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

The `malloc_stats` table returns the cumulative totals, in bytes, of several statistics per type of arena. The command takes no parameters and returns the results as a table.

The following example commands display this result:

```
use performance_schema;
```

```
mysql> SELECT * FROM malloc_stats ORDER BY TYPE DESC LIMIT 3;
```

**Expected output**

```
+-----+-----+-----+-----+-----+
| TYPE   | ALLOCATED | NMALLOC  | NDALLOC  | NREQUESTS |
+-----+-----+-----+-----+-----+
| small  | 23578872  | 586156   | 0        | 2649417   |
| large  | 367382528 | 2218     | 0        | 6355      |
| huge   | 0         | 0        | 0        |           |
+-----+-----+-----+-----+-----+
3 rows in set (0.00 sec)
```

### 11.1.3 Dumping the profile

The profiling samples the `malloc()` calls and stores the sampled stack traces in a separate location in memory. These samples can be dumped into the filesystem. A dump returns a detailed view of the state of the memory.

The process is global; therefore, only a single concurrent run is available and only the most recent runs are stored on disk.

Use the following command to create a profile dump file:

```
flush memory profile;
```

The generated memory profile dumps are written to the `/tmp` directory.

You can analyze the dump files with `jeprof` program, which must be installed on the host system in the appropriate path. This program is a perl script that post-processes the dump files in their raw format. The program has no connection to the `jemalloc` library and the version numbers are not required to match.

To verify the dump, run the following command:

```
ls /tmp/jeprof_mysql*
/tmp/jeprof_mysql.1.0.170013202213
jeprof --show_bytes /tmp/jeprof_mysql.1.0.170013202213 jeprof.*.heap
```

You can also access the memory profile to plot a graph of the memory use. This ability requires that `jeprof` and `dot` are in the `/tmp` path. For the graph to display useful information, the binary file must contain symbol information.

Run the following command:

```
jeprof --dot /usr/sbin/mysqld /tmp/jeprof_mysql.1.0.170013202213 > /tmp/jeprof1.dot
dot --Tpng /tmp/jeprof1.dot > /tmp/jeprof1.png
```

#### Note

An example of [allocation graph](#).

## 11.1.4 PERFORMANCE\_SCHEMA tables

The following tables are implemented to retrieve memory allocation statistics for a running instance or return the cumulative number of allocations requested or allocations returned for a running instance.

More information about the stats that are returned can be found in [jemalloc](#).

### 11.1.5 malloc\_stats\_totals

The current stats for allocations. All measurements are in bytes.

Column Name	Description
ALLOCATED	The total amount the application allo
ACTIVE	The total amount allocated by the ap the page size and this value is greater value. The sum does not include allo stats.arenas.<i>.dirty or stats.arenas
MAPPED	The total amount in chunks that are r extents. This value does not include i large as the stats.active and is a mul
RESIDENT	A maximum number the allocator ha pages. All allocator metadata pages this value. Pages may not be physical demand-zeroed virtual memory that a maximum rather than a precise va The value is greater than the stats.ac
RETAINED	The amount retained by the virtual m system. This value does not include o memory, usually de-committed, untr associated with physical memory an statistics.

Column Name	Description
METADATA	The total amount dedicated to metadata allocations which are used for bootstrap structures. Transparent huge pages

## 11.1.6 malloc\_stats

The cumulative number of allocations requested or allocations returned for a running instance.

Column Name	Description
Type	The type of object: small, large, and huge
ALLOCATED	The number of bytes that are currently allocated
NMALLOC	A cumulative number of times an allocation was made in the small bins. The number includes times when the request was filled from the relevant tcache if <code>opt.tcache</code> is enabled
NDALLOC	A cumulative number of times an allocation was made in the large bins. The number includes times when the request was filled from the relevant tcache if <code>opt.tcache</code> is enabled
NREQUESTS	The cumulative number of allocations requested

## 11.1.7 System variables

The following variables have been added:

### **jemalloc\_detected**

Description: This read-only variable returns `true` if jemalloc with profiling enabled is detected. The following options are required:

- Jemalloc is installed and compiled with profiling enabled
- *Percona Server for MySQL* is configured to use jemalloc by using the environment variable `LD_PRELOAD`.
- The environment variable `MALLOC_CONF` is set to `prof:true`.

The following options are:

- Scope: Global
- Variable Type: Boolean
- Default Value: false

## jemalloc\_profiling

Description: Enables jemalloc profiling. The variable requires `jemalloc_detected`.

- Command Line: `–jemalloc_profiling[=(OFF|ON)]`
- Config File: Yes
- Scope: Global
- Dynamic: Yes
- Variable Type: Boolean
- Default Value: OFF

### 11.1.8 Disable profiling

To disable jemalloc profiling, in a MySQL client, run the following command:

```
set global jemalloc_profiling=off;
```

### 11.1.9 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 11.2 User statistics

This feature adds several `INFORMATION_SCHEMA` tables, several commands, and the `userstat` variable. The tables and commands can be used to understand the server activity better and identify the source of the load.

The functionality is disabled by default and must be enabled by setting `userstat` to `ON`. It works by keeping several hash tables in memory. To avoid contention over global mutexes, each connection has its own local statistics, which are occasionally merged into the global statistics, and the local statistics are then reset to 0.

### 11.2.1 System variables

#### `userstat`

Option	Description
Command-line	Yes
Config file	Yes

Option	Description
Scope	Global
Dynamic	Yes
Data type	BOOLEAN
Default	OFF
Range	ON/OFF

Enables or disables collection of statistics. The default is `OFF`, meaning no statistics are gathered. This is to ensure that the statistics collection doesn't cause any extra load on the server unless desired.

#### `thread_statistics`

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Data type	BOOLEAN
Default	OFF
Range	ON/OFF

Enables or disables collection of thread statistics. The default is `OFF`, meaning no thread statistics are gathered. This is to ensure that the statistics collection doesn't cause any extra load on the server unless desired. The variable `userstat` must be enabled as well in order for thread statistics to be collected.

## 11.2.2 INFORMATION\_SCHEMA Tables

### `INFORMATION_SCHEMA.CLIENT_STATISTICS`

Column Name	Description
'CLIENT'	'The IP address or hostname from which the connection was made.'
'TOTAL_CONNECTIONS'	'The number of connections created from this client.'
'CONCURRENT_CONNECTIONS'	'The number of concurrent connections from this client.'
'CONNECTED_TIME'	'The cumulative number of seconds that the client has been connected to the server from this client.'
'BUSY_TIME'	'The cumulative number of seconds that the client has been busy with the server from this client.'
'CPU_TIME'	'The cumulative CPU time elapsed, in seconds, for all connections from this client.'

Column Name	Description
'BYTES_RECEIVED'	'The number of bytes received from t
'BYTES_SENT'	'The number of bytes sent to this clien
'BINLOG_BYTES_WRITTEN'	'The number of bytes written to the b
'ROWS_FETCHED'	'The number of rows fetched by this c
'ROWS_UPDATED'	'The number of rows updated by this
'TABLE_ROWS_READ'	'The number of rows read from table (different from ROWS_FETCHED.)'
'SELECT_COMMANDS'	'The number of SELECT commands exe
'UPDATE_COMMANDS'	'The number of UPDATE commands e
'OTHER_COMMANDS'	'The number of other commands exe
'COMMIT_TRANSACTIONS'	'The number of COMMIT commands
'ROLLBACK_TRANSACTIONS'	'The number of ROLLBACK command
'DENIED_CONNECTIONS'	'The number of connections denied t
'LOST_CONNECTIONS'	'The number of this client's connecti
'ACCESS_DENIED'	'The number of times this client's con denied.'
'EMPTY_QUERIES'	'The number of times this client's con server.'

This table holds statistics about client connections. The Percona version of the feature restricts this table's visibility to users who have the `SUPER` or `PROCESS` privilege.

For example:

```
mysql>SELECT * FROM INFORMATION_SCHEMA.CLIENT_STATISTICS\G
```

### Expected output

```

***** 1. row *****
      CLIENT: 10.1.12.30
    TOTAL_CONNECTIONS: 20
  CONCURRENT_CONNECTIONS: 0
    CONNECTED_TIME: 0
      BUSY_TIME: 93
      CPU_TIME: 48
    BYTES_RECEIVED: 5031
    BYTES_SENT: 276926
  BINLOG_BYTES_WRITTEN: 217
    ROWS_FETCHED: 81
    ROWS_UPDATED: 0
  TABLE_ROWS_READ: 52836023
  SELECT_COMMANDS: 26
  UPDATE_COMMANDS: 1
    OTHER_COMMANDS: 145
  COMMIT_TRANSACTIONS: 1
  ROLLBACK_TRANSACTIONS: 0
  DENIED_CONNECTIONS: 0
  LOST_CONNECTIONS: 0
    ACCESS_DENIED: 0
    EMPTY_QUERIES: 0

```

## 11.2.3 INFORMATION\_SCHEMA tables

### INFORMATION\_SCHEMA.INDEX\_STATISTICS

Column Name	Description
'TABLE_SCHEMA'	'The schema (database) name.'
'TABLE_NAME'	'The table name.'
'INDEX_NAME'	'The index name (as visible in SHOW INDEX)'
'ROWS_READ'	'The number of rows read from this index'

This table shows statistics on index usage. An older version of the feature contained a single column that had the `TABLE_SCHEMA`, `TABLE_NAME`, and `INDEX_NAME` columns concatenated together. The *Percona* version of the feature separates these into three columns. Users can see entries only for tables to which they have `SELECT` access.

This table makes it possible to do many things that were difficult or impossible previously. For example, you can use it to find unused indexes and generate `DROP` commands to remove them.

Example:

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INDEX_STATISTICS WHERE
TABLE_NAME='tables_priv';
```

**Expected output**

```

+-----+-----+-----+-----+
| TABLE_SCHEMA | TABLE_NAME          | INDEX_NAME          | ROWS_READ |
+-----+-----+-----+-----+
| mysql         | tables_priv          | PRIMARY             |          2 |
+-----+-----+-----+-----+

```

**Note**

The current implementation of index statistics doesn't support partitioned tables.

**INFORMATION\_SCHEMA.TABLE\_STATISTICS**

Column Name	Description
'TABLE_SCHEMA'	'The schema (database) name.'
'TABLE_NAME'	'The table name.'
'ROWS_READ'	'The number of rows read from the table.'
'ROWS_CHANGED'	'The number of rows changed in the table.'
'ROWS_CHANGED_X_INDEXES'	'The number of rows changed in the table and indexes changed.'

This table is similar in function to the `INDEX_STATISTICS` table.

For example:

```
mysql> SELECT * FROM INFORMATION_SCHEMA.TABLE_STATISTICS WHERE
TABLE_NAME='`tables_priv`';
```

**Expected output**

```

+-----+-----+-----+-----+
| TABLE_SCHEMA | TABLE_NAME          | ROWS_READ | ROWS_CHANGED |
| ROWS_CHANGED_X_INDEXES |
+-----+-----+-----+-----+
| mysql         | tables_priv          |          2 |          0 |
| 0 |
+-----+-----+-----+-----+

```

**Note**

The current implementation of table statistics doesn't support partitioned tables.

**INFORMATION\_SCHEMA.THREAD\_STATISTICS**

Column Name	Description
'THREAD_ID'	'Thread ID'
'TOTAL_CONNECTIONS'	'The number of connections created'
'CONNECTED_TIME'	'The cumulative number of seconds from this thread.'
'BUSY_TIME'	'The cumulative number of seconds'
'CPU_TIME'	'The cumulative CPU time elapsed w'
'BYTES_RECEIVED'	'The number of bytes received from t'
'BYTES_SENT'	'The number of bytes sent to this thre'
'BINLOG_BYTES_WRITTEN'	'The number of bytes written to the b'
'ROWS_FETCHED'	'The number of rows fetched by this t'
'ROWS_UPDATED'	'The number of rows updated by this'
'TABLE_ROWS_READ'	'The number of rows read from table'
'SELECT_COMMANDS'	'The number of SELECT commands exe'
'UPDATE_COMMANDS'	'The number of UPDATE commands e'
'OTHER_COMMANDS'	'The number of other commands exe'
'COMMIT_TRANSACTIONS'	'The number of COMMIT commands s'
'ROLLBACK_TRANSACTIONS'	'The number of ROLLBACK command'
'DENIED_CONNECTIONS'	'The number of connections denied t'
'LOST_CONNECTIONS'	'The number of thread connections t'
'ACCESS_DENIED'	'The number of times this thread issu'
'EMPTY_QUERIES'	'The number of times this thread sen'
'TOTAL_SSL_CONNECTIONS'	'The number of thread connections t'

In order for this table to be populated with statistics, the additional variable `thread_statistics` should be set to `ON`.

**INFORMATION\_SCHEMA.USER\_STATISTICS**

Column Name	Description
'USER'	'The username. The value <code>#mysql_sys</code> username (such as for the replica SC
'TOTAL_CONNECTIONS'	'The number of connections created'
'CONCURRENT_CONNECTIONS'	'The number of concurrent connectio'
'CONNECTED_TIME'	'The cumulative number of seconds from this user.'
'BUSY_TIME'	

Column Name	Description
	'The cumulative number of seconds this user.'
'CPU_TIME'	'The cumulative CPU time elapsed, in connections.'
'BYTES_RECEIVED'	'The number of bytes received from t
'BYTES_SENT'	'The number of bytes sent to this use
'BINLOG_BYTES_WRITTEN'	'The number of bytes written to the b
'ROWS_FETCHED'	'The number of rows fetched by this u
'ROWS_UPDATED'	'The number of rows updated by this
'TABLE_ROWS_READ'	'The number of rows read from table (different from ROWS_FETCHED.)'
'SELECT_COMMANDS'	'The number of SELECT commands exe
'UPDATE_COMMANDS'	'The number of UPDATE commands e
'OTHER_COMMANDS'	'The number of other commands exe
'COMMIT_TRANSACTIONS'	'The number of COMMIT commands
'ROLLBACK_TRANSACTIONS'	'The number of ROLLBACK command
'DENIED_CONNECTIONS'	'The number of connections denied t
'LOST_CONNECTIONS'	'The number of this user's connectio
'ACCESS_DENIED'	'The number of times this user's conr denied.'
'EMPTY_QUERIES'	'The number of times this user's conr server.'

This table contains information about user activity. The *Percona* version of the patch restricts this table's visibility to users who have the `SUPER` or `PROCESS` privilege.

The table gives answers to questions such as which users cause the most load, and whether any users are being abusive. It also lets you measure how close to capacity the server may be. For example, you can use it to find out whether replication is likely to start falling behind.

Example:

```
mysql>SELECT * FROM INFORMATION_SCHEMA.USER_STATISTICS\G
```

### Expected output

```

***** 1. row *****
      USER: root
    TOTAL_CONNECTIONS: 5592
  CONCURRENT_CONNECTIONS: 0
      CONNECTED_TIME: 6844
        BUSY_TIME: 179
          CPU_TIME: 72
    BYTES_RECEIVED: 603344
      BYTES_SENT: 15663832
  BINLOG_BYTES_WRITTEN: 217
      ROWS_FETCHED: 9793
      ROWS_UPDATED: 0
    TABLE_ROWS_READ: 52836023
    SELECT_COMMANDS: 9701
    UPDATE_COMMANDS: 1
      OTHER_COMMANDS: 2614
    COMMIT_TRANSACTIONS: 1
  ROLLBACK_TRANSACTIONS: 0
    DENIED_CONNECTIONS: 0
      LOST_CONNECTIONS: 0
      ACCESS_DENIED: 0
      EMPTY_QUERIES: 0

```

## 11.2.4 Commands Provided

- `FLUSH CLIENT_STATISTICS`
- `FLUSH INDEX_STATISTICS`
- `FLUSH TABLE_STATISTICS`
- `FLUSH THREAD_STATISTICS`
- `FLUSH USER_STATISTICS`

These commands discard the specified type of stored statistical information.

- `SHOW CLIENT_STATISTICS`
- `SHOW INDEX_STATISTICS`
- `SHOW TABLE_STATISTICS`
- `SHOW THREAD_STATISTICS`
- `SHOW USER_STATISTICS`

These commands are another way to display the information you can get from the `INFORMATION_SCHEMA` tables. The commands accept `WHERE` clauses. They also accept but ignore `LIKE` clauses.

## 11.2.5 Status Variables

### `Com_show_client_statistics`

Option	Description
Scope	Global/Session

Option	Description
Data type	numeric

The `Com_show_client_statistics` statement counter variable indicates the number of times the statement `SHOW CLIENT_STATISTICS` has been executed.

#### `Com_show_index_statistics`

Option	Description
Scope	Global/Session
Data type	numeric

The `Com_show_index_statistics` statement counter variable indicates the number of times the statement `SHOW INDEX_STATISTICS` has been executed.

#### `Com_show_table_statistics`

Option	Description
Scope	Global/Session
Data type	numeric

The `Com_show_table_statistics` statement counter variable indicates the number of times the statement `SHOW TABLE_STATISTICS` has been executed.

#### `Com_show_thread_statistics`

Option	Description
Scope	Global/Session
Data type	numeric

The `Com_show_thread_statistics` statement counter variable indicates the number of times the statement `SHOW THREAD_STATISTICS` has been executed.

#### `Com_show_user_statistics`

Option	Description
Scope	Global/Session
Data type	numeric

The `Com_show_user_statistics` statement counter variable indicates the number of times the statement `SHOW USER_STATISTICS` has been executed.

## 11.2.6 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-25

## 11.3 Slow query log

This feature adds microsecond time resolution and additional statistics to the slow query log output. It lets you enable or disable the slow query log at runtime, adds logging for the replica SQL thread, and adds fine-grained control over what and how much to log into the slow query log.

You can use Percona-Toolkit's [pt-query-digest](#) tool to aggregate similar queries together and report on those that consume the most execution time.

### 11.3.1 System Variables

#### Log\_slow\_filter

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global, Session
Dynamic	Yes

Filters the slow log by the query's execution plan. The value is a comma-delimited string, and can contain any combination of the following values:

- `full_scan`: The query performed a full table scan.
- `full_join`: The query performed a full join (a join without indexes).
- `tmp_table`: The query created an implicit internal temporary table.
- `tmp_table_on_disk`: The query's temporary table was stored on disk.
- `filesort`: The query used a filesort.
- `filesort_on_disk`: The filesort was performed on disk.

Values are OR'ed together. If the string is empty, then the filter is disabled. If it is not empty, then queries will only be logged to the slow log if their execution plan matches one of the types of plans present in the filter.

For example, to log only queries that perform a full table scan, set the value to `full_scan`. To log only queries that use on-disk temporary storage for intermediate results, set the value to `tmp_table_on_disk,filesort_on_disk`.

**log\_slow\_rate\_type**

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	yes
Data type	Enumerated
Default	session, query

Specifies semantic of `log_slow_rate_limit` - `session` or `query`.

**log\_slow\_rate\_limit**

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global, session
Dynamic	yes
Default	1
Range	1-1000

Behavior of this variable depends on the selected `log_slow_rate_type`.

Specifies that only a fraction of `session/query` should be logged. Logging is enabled for every `n`th `session/query`. By default, `n` is 1, so logging is enabled for every `session/query`. Please note: when `log_slow_rate_type` is `session` rate limiting is disabled for the replication thread.

Logging all queries might consume I/O bandwidth and cause the log file to grow large.

\* When `log_slow_rate_type` is `session`, this option lets you log full sessions, so you have complete records of sessions for later analysis; but you can rate-limit the number of sessions that are logged. Note that this feature will not work well if your application uses any type of connection pooling or persistent connections. Note that you change `log_slow_rate_limit` in `session` mode, you should reconnect for get effect.

\* When `log_slow_rate_type` is `query`, this option lets you log just some queries for later analysis. For example, if you set the value to 100, then one percent of queries will be logged.

Note that every query has global unique `query_id` and every connection can has it own (session) `log_slow_rate_limit`. Decision “log or no” calculated in following manner:

- if `log_slow_rate_limit` is 1 - log every query
- If `log_slow_rate_limit` 1 - randomly log every  $1 / \text{log\_slow\_rate\_limit}$  query.

This allows flexible setup logging behavior.

For example, if you set the value to 100, then one percent of `sessions/queries` will be logged. In *Percona Server for MySQL* information about the `log_slow_rate_limit` has been added to the slow query log. This means that if the `log_slow_rate_limit` is effective it will be reflected in the slow query log for each written query.

#### Expected output

```
Log_slow_rate_type: query Log_slow_rate_limit: 10
```

#### Log\_slow\_sp\_statements

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Data type	Enumerated
Default	session
Range	session, query

If `TRUE`, statements executed by stored procedures are logged to the slow if it is open.

*Percona Server for MySQL* implemented improvements for logging of stored procedures to the slow query log:

- \* Each query from a stored procedure is now logged to the slow query log individually
- \* ``CALL`` itself isn't logged to the slow query log anymore as this would be counting twice for the same query which would lead to incorrect results
- \* Queries that were called inside of stored procedures are annotated in the slow query log with the stored procedure name in which they run.

Example of the improved stored procedure slow query log entry:

```
mysqlDELIMITER //
mysqlCREATE PROCEDURE improved_sp_log()
  BEGIN
    SELECT * FROM City;
```

```

SELECT * FROM Country;
END//
mysqlDELIMITER ;
mysqlCALL improved_sp_log();

```

When we check the slow query log after running the stored procedure, with `log_slow_sp_statements` set to `TRUE`, it should look like this:

#### Expected output

```

# Time: 150109 11:38:55
# User@Host: root[root] @ localhost []
# Thread_id: 40 Schema: world Last_errno: 0 Killed: 0
# Query_time: 0.012989 Lock_time: 0.000033 Rows_sent: 4079 Rows_examined: 4079
Rows_affected: 0 Rows_read: 4079
# Bytes_sent: 161085
# Stored routine: world.improved_sp_log
SET timestamp=1420803535;
SELECT * FROM City;
# User@Host: root[root] @ localhost []
# Thread_id: 40 Schema: world Last_errno: 0 Killed: 0
# Query_time: 0.001413 Lock_time: 0.000017 Rows_sent: 4318 Rows_examined: 4318
Rows_affected: 0 Rows_read: 4318
# Bytes_sent: 194601
# Stored routine: world.improved_sp_log
SET timestamp=1420803535;

```

If variable `log_slow_sp_statements` is set to `FALSE`:

- Entry is added to a slow-log for a `CALL` statement only and not for any of the individual statements run in that stored procedure
- Execution time is reported for the `CALL` statement as the total execution time of the `CALL` including all its statements

If we run the same stored procedure with the `log_slow_sp_statements` is set to `FALSE` slow query log should look like this:

#### Expected output

```

# Time: 150109 11:51:42
# User@Host: root[root] @ localhost []
# Thread_id: 40 Schema: world Last_errno: 0 Killed: 0
# Query_time: 0.013947 Lock_time: 0.000000 Rows_sent: 4318 Rows_examined: 4318
Rows_affected: 0 Rows_read: 4318
# Bytes_sent: 194612
SET timestamp=1420804302;
CALL improved_sp_log();

```

#### Note

Support for logging stored procedures doesn't involve triggers, so they won't be logged even if this feature is enabled.

**log\_slow\_verbosity**

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global, session
Dynamic	Yes

Specifies how much information to include in your slow log. The value is a comma-delimited string, and can contain any combination of the following values:

- `microtime`: Log queries with microsecond precision.
- `query_plan`: Log information about the query's execution plan.
- `innodb`: Log *InnoDB* statistics.
- `minimal`: Equivalent to enabling just `microtime`.
- `standard`: Equivalent to enabling `microtime,query_plan`.
- `full`: Equivalent to all other values OR'ed together without the `profiling` and `profiling_use_getrusage` options.
- `profiling`: Enables profiling of all queries in all connections.
- `profiling_use_getrusage`: Enables usage of the `getrusage` function.
- `query_info`: Enables printing `Query_tables` and `Query_digest` into the slow query log. These fields are disabled by default.

Values are OR'ed together.

For example, to enable microsecond query timing and *InnoDB* statistics, set this option to `microtime,innodb` or `standard`. To turn all options on, set the option to `full`.

**slow\_query\_log\_use\_global\_control**

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Default	None

Specifies which variables have global scope instead of local. For such variables, the global variable value is used in the current session, but without copying this value to the session value. Value is a “flag” variable - you can specify multiple values separated by commas

- `none` : All variables use local scope
- `log_slow_filter` : Global variable `log_slow_filter` has effect (instead of local)
- `log_slow_rate_limit` : Global variable `log_slow_rate_limit` has effect (instead of local)
- `log_slow_verbosity` : Global variable `log_slow_verbosity` has effect (instead of local)
- `long_query_time` : Global variable `long_query_time` has effect (instead of local)
- `min_examined_row_limit` : Global variable `min_examined_row_limit` has effect (instead of local)
- `all` Global variables has effect (instead of local)

#### `slow_query_log_always_write_time`

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Default	10

This variable can be used to specify the query execution time after which the query will be written to the slow query log. It can be used to specify an additional execution time threshold for the slow query log, that, when exceeded, will cause a query to be logged unconditionally, that is, `log_slow_rate_limit` will not apply to it.

## 11.3.2 Other information

### Changes to the log format

The feature adds more information to the slow log output.

```

# Time: 130601 8:01:06.058915
# User@Host: root[root] @ localhost [] Id: 42
# Schema: imdb Last_errno: 0 Killed: 0
# Query_time: 7.725616 Lock_time: 0.000328 Rows_sent: 4 Rows_examined: 1543720
Rows_affected: 0
# Bytes_sent: 272 Tmp_tables: 0 Tmp_disk_tables: 0 Tmp_table_sizes: 0
# Full_scan: Yes Full_join: No Tmp_table: No Tmp_table_on_disk: No
# Filesort: No Filesort_on_disk: No Merge_passes: 0
SET timestamp=1370073666;
SELECT id,title,production_year FROM title WHERE title = 'Bambi';

```

Another example (`log_slow_verbosity =profiling`):

### Expected output

```
# Time: 130601 8:03:20.700441
# User@Host: root[root] @ localhost [] Id: 43
# Schema: imdb Last_errno: 0 Killed: 0
# Query_time: 7.815071 Lock_time: 0.000261 Rows_sent: 4 Rows_examined: 1543720
Rows_affected: 0
# Bytes_sent: 272
# Profile_starting: 0.000125 Profile_starting_cpu: 0.000120
Profile_checking_permissions: 0.000021 Profile_checking_permissions_cpu: 0.000021
Profile_Opening_tables: 0.000049 Profile_Opening_tables_cpu: 0.000048 Profile_init: 0.000048
Profile_init_cpu: 0.000049 Profile_System_lock: 0.000049 Profile_System_lock_cpu: 0.000048
Profile_optimizing: 0.000024 Profile_optimizing_cpu: 0.000024 Profile_statistics: 0.000036
Profile_statistics_cpu: 0.000037 Profile_preparing: 0.000029 Profile_preparing_cpu: 0.000029
Profile_executing: 0.000012 Profile_executing_cpu: 0.000012 Profile_Sending_data: 7.814583
Profile_Sending_data_cpu: 7.811634 Profile_end: 0.000013 Profile_end_cpu: 0.000012
Profile_query_end: 0.000014 Profile_query_end_cpu: 0.000014 Profile_closing_tables: 0.000023
Profile_closing_tables_cpu: 0.000023 Profile_freeing_items: 0.000051
Profile_freeing_items_cpu: 0.000050 Profile_logging_slow_query: 0.000006
Profile_logging_slow_query_cpu: 0.000006
# Profile_total: 7.815085 Profile_total_cpu: 7.812127
SET timestamp=1370073800;
SELECT id,title,production_year FROM title WHERE title = 'Bambi';
```

Notice that the `Killed: \`` keyword is followed by zero when the query successfully completes. If the query was killed, the ```Killed:`` keyword is followed by a number other than zero:

Killed Numeric Code	Exception
0	NOT_KILLED
1	KILL_BAD_DATA
1053	ER_SERVER_SHUTDOWN (see MySQL D
1317	ER_QUERY_INTERRUPTED (see MySQL
3024	ER_QUERY_TIMEOUT (see MySQL Doc
Any other number	KILLED_NO_VALUE (Catches all other

### Connection and Schema Identifier

Each slow log entry now contains a connection identifier, so you can trace all the queries coming from a single connection. This is the same value that is shown in the `Id` column in `SHOW FULL PROCESSLIST` or returned from the `CONNECTION_ID()` function.

Each entry also contains a schema name, so you can trace all the queries whose default database was set to a particular schema.

### Expected output

```
# Id: 43 Schema: imdb
```

## Microsecond time resolution and extra row information

This is the original functionality offered by the `microslow` feature. `Query_time` and `Lock_time` are logged with microsecond resolution.

The feature also adds information about how many rows were examined for `SELECT` queries, and how many were analyzed and affected for `UPDATE`, `DELETE`, and `INSERT` queries,

```

Expected output

# Query_time: 0.962742 Lock_time: 0.000202 Rows_sent: 4 Rows_examined: 1543719
Rows_affected: 0
  
```

Values and context:

- `Rows_examined`: Number of rows scanned - `SELECT`
- `Rows_affected`: Number of rows changed - `UPDATE`, `DELETE`, `INSERT`

## Memory footprint

The feature provides information about the amount of bytes sent for the result of the query and the number of temporary tables created for its execution - differentiated by whether they were created on memory or on disk - with the total number of bytes used by them.

```

Expected output

# Bytes_sent: 8053 Tmp_tables: 1 Tmp_disk_tables: 0 Tmp_table_sizes: 950528
  
```

Values and context:

- `Bytes_sent`: The amount of bytes sent for the result of the query
- `Tmp_tables`: Number of temporary tables created on memory for the query
- `Tmp_disk_tables`: Number of temporary tables created on disk for the query
- `Tmp_table_sizes`: Total Size in bytes for all temporary tables used in the query

## Query plan information

Each query can be executed in various ways. For example, it may use indexes or do a full table scan, or a temporary table may be needed. These are the things that you can usually see by running `EXPLAIN` on the query. The feature will now allow you to see the most important facts about the execution in the log file.

```

Expected output

# Full_scan: Yes Full_join: No Tmp_table: No Tmp_table_on_disk: No
# Filesort: No Filesort_on_disk: No Merge_passes: 0
  
```

The values and their meanings are documented with the `log_slow_filter` option.

### InnoDB usage information

The final part of the output is the *InnoDB* usage statistics. *MySQL* currently shows many per-session statistics for operations with `SHOW SESSION STATUS`, but that does not include those of *InnoDB*, which are always global and shared by all threads. This feature lets you see those values for a given query.

#### Expected output

```
# InnoDB_IO_r_ops: 6415 InnoDB_IO_r_bytes: 105103360 InnoDB_IO_r_wait: 0.001279
# InnoDB_rec_lock_wait: 0.000000 InnoDB_queue_wait: 0.000000
# InnoDB_pages_distinct: 6430
```

#### Values:

- `innodb_IO_r_ops`: Counts the number of page read operations scheduled. The actual number of read operations may be different, but since this can be done asynchronously, there is no good way to measure it.
- `innodb_IO_r_bytes`: Similar to `innodb_IO_r_ops`, but the unit is bytes.
- `innodb_IO_r_wait`: Shows how long (in seconds) it took *InnoDB* to actually read the data from storage.
- `innodb_rec_lock_wait`: Shows how long (in seconds) the query waited for row locks.
- `innodb_queue_wait`: Shows how long (in seconds) the query spent either waiting to enter the *InnoDB* queue or inside that queue waiting for execution.
- `innodb_pages_distinct`: Counts approximately the number of unique pages the query accessed. The approximation is based on a small hash array representing the entire buffer pool, because it could take a lot of memory to map all the pages. The inaccuracy grows with the number of pages accessed by a query, because there is a higher probability of hash collisions.

If the query did not use *InnoDB* tables, that information is written into the log instead of the above statistics.

## 11.3.3 Related reading

- [Impact of logging on MySQL's performance](#)
- [log\\_slow\\_filter Usage](#)
- [Added microseconds to the slow query log event time](#)

### 11.3.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-23

## 11.4 Process list

This page describes Percona changes to both the standard MySQL `SHOW PROCESSLIST` command and the standard MySQL `INFORMATION_SCHEMA` table `PROCESSLIST`.

### 11.4.1 INFORMATION\_SCHEMA Tables

`INFORMATION_SCHEMA.PROCESSLIST`

This table implements modifications to the standard MySQL `INFORMATION_SCHEMA` table `PROCESSLIST`.

Column Name	Description
'ID'	'The connection identifier.'
'USER'	'The MySQL user who issued the statement.'
'HOST'	'The host name of the client issuing the statement.'
'DB'	'The default database, if one is selected.'
'COMMAND'	'The type of command the thread is executing.'
'TIME'	'The time in seconds that the thread has been running.'
'STATE'	'An action, event, or state that indicates the thread's current state.'
'INFO'	'The statement that the thread is executing. For statements that are not fully executed, the value is truncated to the length of the column.'
'TIME_MS'	'The time in milliseconds that the thread has been running.'
'ROWS_EXAMINED'	'The number of rows examined by the thread. For statements that are not fully executed, the value is truncated to the length of the column. For statements that are fully executed, the value is correct after the statement has completed. For statements that are not fully executed, the value is not updated for each examined row. For statements that are fully executed, the value is correct after the statement has completed.'
'ROWS_SENT'	'The number of rows sent by the thread. For statements that are not fully executed, the value is truncated to the length of the column. For statements that are fully executed, the value is correct after the statement has completed.'
'TID'	'The Linux Thread ID. For Linux, this column is the same as the LWP ID and can be seen in the <code>ps -l</code> command. If the <code>show_thread_info</code> system variable is enabled, "TID" is not null for only current statements received via "extra" connections.'

### 11.4.2 Example output

Table `PROCESSLIST`:

```
mysql> SELECT * FROM INFORMATION_SCHEMA.PROCESSLIST;
```

### Expected output

```

+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| ID | USER | HOST      | DB              | COMMAND | TIME | STATE   |
INFO |      |           |                 |         |     |        |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| 12 | root | localhost | information_schema | Query   | 0   | executing |
processlist |      |           |                 |         |     |        |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+

```

## 11.4.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-23

## 11.5 Misc. INFORMATION\_SCHEMA tables

This page lists the `INFORMATION_SCHEMA` tables added to standard *MySQL* by *Percona Server for MySQL* that don't exist elsewhere in the documentation.

### 11.5.1 Temporary tables

#### Note

**This feature implementation is considered tech preview quality.**

Only the temporary tables that were explicitly created with `CREATE TEMPORARY TABLE` or `ALTER TABLE` are shown, and not the ones created to process complex queries.

#### INFORMATION\_SCHEMA.GLOBAL\_TEMPORARY\_TABLES

Column Name	Description
'SESSION_ID'	'MySQL connection id'
'TABLE_SCHEMA'	'Schema in which the temporary table'
'TABLE_NAME'	'Name of the temporary table'
'ENGINE'	'Engine of the temporary table'
'NAME'	'Internal name of the temporary table'

Column Name	Description
'TABLE_ROWS'	'Number of rows of the temporary table'
'AVG_ROW_LENGTH'	'Average row length of the temporary table'
'DATA_LENGTH'	'Size of the data (Bytes)'
'INDEX_LENGTH'	'Size of the indexes (Bytes)'
'CREATE_TIME'	'Date and time of creation of the temporary table'
'UPDATE_TIME'	'Date and time of the latest update of the temporary table'

This table holds information on the temporary tables that exist for all connections. You don't need the `SUPER` privilege to query this table.

#### INFORMATION\_SCHEMA.TEMPORARY\_TABLES

Column Name	Description
'SESSION_ID'	'MySQL connection id'
'TABLE_SCHEMA'	'Schema in which the temporary table exists'
'TABLE_NAME'	'Name of the temporary table'
'ENGINE'	'Engine of the temporary table'
'NAME'	'Internal name of the temporary table'
'TABLE_ROWS'	'Number of rows of the temporary table'
'AVG_ROW_LENGTH'	'Average row length of the temporary table'
'DATA_LENGTH'	'Size of the data (Bytes)'
'INDEX_LENGTH'	'Size of the indexes (Bytes)'
'CREATE_TIME'	'Date and time of creation of the temporary table'
'UPDATE_TIME'	'Date and time of the latest update of the temporary table'

This table holds information on the temporary tables existing for the running connection.

## 11.5.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 12. Troubleshoot

### 12.1 Use Percona Monitoring and Management (PMM) Advisors

Percona Monitoring and Management (PMM) provides several categories of Advisors. Each Advisor contains a set of automated checks. These checks investigate your database settings for a specific range of possible issues.

The Percona Platform hosts the Advisors.

The PMM Server automatically downloads the Advisors if the `Advisors` and `Telemetry` options are enabled in `Configuration > Settings > Advanced Settings`. Both options are enabled by default.

#### See also

[PMM documentation - Advisor checks](#)

#### 12.1.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-04-19

### 12.2 Too many connections warning

If the `log_error_verbosity` system variable is set to `2` or higher, this feature generates the `Too many connections` warning in the log.

#### 12.2.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 12.3 Handle corrupted tables

When a server subsystem tries to access a corrupted table, the server may crash. If this outcome is not desirable when a corrupted table is encountered, set the new system `innodb_corrupt_table_action` variable to a value which allows the ongoing operation to continue without crashing the server.

The server error log registers attempts to access corrupted table pages.

### Interacting with the `innodb_force_recovery` variable

The `innodb_corrupt_table_action` variable may work in conjunction with the `innodb_force_recovery` variable which considerably reduces the effect of *InnoDB* subsystems running in the background.

If the `innodb_force_recovery` option is <4, corrupted pages are lost and the server may continue to run due to the `innodb_corrupt_table_action` variable having a non-default value.

For more information about the `innodb_force_recovery` variable, see [Forcing InnoDB Recovery](#) from the MySQL Reference Manual.

This feature adds a system variable.

### 12.3.1 System variables

#### `innodb_corrupt_table_action`

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Data type	ULONG
Default	assert
Range	assert, warn, salvage

- Enabling `innodb_file_per_table` and using the `assert` value creates an assertion failure which causes *XtraDB* to intentionally crash the server. This action is expected when detecting corrupted data in a single-table tablespace.
- Enabling `innodb_file_per_table` and using the `warn` value causes *XtraDB* to pass the table corruption as `corrupt table` instead of crashing the server. Detecting the file as corrupt also disables the file I/O for that data file, except for the deletion operation.
- Enabling `innodb_file_per_table` and using the `salvage` value causes *XtraDB* to allow read access to the corrupted tablespace but ignores any corrupted pages.

## 12.3.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 12.4 Thread based profiling

Percona Server for MySQL now uses thread based profiling by default, instead of process based profiling. This was implemented because with process based profiling, threads on the server, other than the one being profiled, can affect the profiling information.

Thread based profiling is using the information provided by the kernel `getrusage` function. Since the 2.6.26 kernel version, thread based resource usage is available with the `RUSAGE_THREAD`. This means that the thread based profiling will be used if you're running the 2.6.26 kernel or newer, or if the `RUSAGE_THREAD` has been ported back.

This feature is enabled by default if your system supports it, in other cases it uses process based profiling.

### 12.4.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-09-25

## 12.5 Stack trace

Developers use the stack trace in the debug process, either an interactive investigation or during the post-mortem. No configuration is required to generate a stack trace.

Stack trace adds the following:

Name	Description
Prints binary BuildID	The Strip utility removes unneeded symbols to reduce the size. This method is standard. The image is essential. The BuildID lets you identify the image. The strip utility removes the binary symbols to

Name	Description
Print the server version information	The version information establishes the version of the operating system and applications, such as MySQL, only prior to the crash. When the crash occurs, the size of the dump is truncated.

## 12.5.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-25

## 12.6 Using libcoredumper

A core dump file is the documented moment of a computer when either the computer or an application exits. Developers examine the dump as one of the tasks when searching for the cause of a failure.

The `libcoredumper` is a free and Open Source fork of `google-coreddumper`, enhanced to work on newer Linux versions, and GCC and CLANG.

### Enabling the `libcoredumper`

Enable core dumps for troubleshooting purposes.

To enable the `libcoredumper`, add the `coreddumper` variable to the `mysqld` section of `my.cnf`. This variable is independent of the older `core-file` variable.

The variable can have the following possible values:

Value	Description
Blank	The core dump is saved under MySQL
A path ending with /	The core dump is saved under the sp
Full path with a filename	The core dump is saved under the sp

Restart the server.

### Verifying the `libcoredumper` is active

MySQL writes to the log when generating a core file and delegates the core dump operation to the Linux kernel.

```
Writing a core file
```

MySQL using the `libcoredumper` to generate the file creates the following message in the log:

```
Writing a core file using lib coredumper
```

Every core file adds a crash timestamp instead of a PID for the following reasons:

- Correlates the core file with the crash. MySQL prints a UTC timestamp on the crash log.

```
10:02:09 UTC - mysqld got signal 11;
```

- Stores multiple core files.

#### Note

For example, operators and containers run as the process id of PID 1. If the process ID is used to identify the core file, each container crash generates a core dump that overwrites the previous core file.

### Disabling the libcoredumper

You can disable the libcoredumper. A core file may contain sensitive data and takes disk space.

To disable the `libcoredumper` you must do the following:

1. In the `mysqld` section of `my.cnf`, remove the `libcoredumper` variable.
1. Restart the server.

## 12.6.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-04

# 13. Storage engines

## 13.1 Percona XtraDB storage engine

### 13.1.1 The Percona XtraDB storage engine

Percona XtraDB is an enhanced version of the *InnoDB* storage engine, designed to better scale on modern hardware. It also includes a variety of other features useful in high-performance environments. It is fully backward compatible, and so can be used as a drop-in replacement for standard *InnoDB*.

Percona XtraDB includes all of *InnoDB*'s robust, reliable ACID-compliant design and advanced MVCC architecture, and builds on that solid foundation with more features, more tunability, more metrics, and more scalability. In particular, it is designed to scale better on many cores, use memory more efficiently, and be more convenient and useful. The new features are specially designed to alleviate some of *InnoDB*'s limitations. We choose features and fixes based on customer requests and on our best judgment of real-world needs as a high-performance consulting company.

Percona XtraDB engine will not have further binary releases, it is distributed as part of the *Percona Server for MySQL*.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-07-27

### 13.1.2 Improved MEMORY storage engine

A *Fixed Row Format* (FRF) is used in the *MEMORY* storage engine. The fixed row format imposes restrictions on the type of columns as it assigns in advance a limited amount of memory per row. This renders a *VARCHAR* field in a *CHAR* field in practice and makes it impossible to have a *TEXT* or *BLOB* field with that engine implementation.

To overcome this limitation, the *Improved MEMORY Storage Engine* is introduced in this release for supporting **true** *VARCHAR*, *VARBINARY*, *TEXT*, and *BLOB* fields in the *MEMORY* tables.

This implementation is based on the *Dynamic Row Format* (DFR) introduced by the [mysql-heap-dynamic-rows](#) patch.

DFR is used to store column values in a variable-length form, thus helping to decrease the memory footprint of those columns and making possible *BLOB* and *TEXT* fields and real *VARCHAR* and *VARBINARY*.

Unlike the fixed implementation, each column value in `DRF` uses only as much space as required. Variable-length values can use up to 4 bytes to store the actual value length, and only the necessary number of blocks is used to store the value.

Rows in `DRF` are represented internally by multiple memory blocks, which means that a single row can consist of multiple blocks organized into one set. Each row occupies at least one block, there can not be multiple rows within a single block. Block size can be configured when creating a table (see below).

This `DRF` implementation has two caveats regarding ordering and indexes.

## Caveats

### ORDERING OF ROWS

In the absence of `ORDER BY`, records may be returned in a different order than the previous `MEMORY` implementation.

This is not a bug. Any application relying on a specific order without an `ORDER BY` clause may deliver unexpected results. A specific order without `ORDER BY` is a side effect of a storage engine and query optimizer implementation which may and will change between minor *MySQL* releases.

### INDEXING

It is currently impossible to use indexes on `BLOB` columns due to some limitations of the *Dynamic Row Format*. Trying to create such an index will fail with the following error:

 **Expected output**

```
BLOB column '<name>' can't be used in key specification with the used table type.
```

## Restrictions

For performance reasons, a mixed solution is implemented: the fixed format is used at the beginning of the row, while the dynamic one is used for the rest of it.

The size of the fixed-format portion of the record is chosen automatically on `CREATE TABLE` and cannot be changed later. This, in particular, means that no indexes can be created later with `CREATE INDEX` or `ALTER TABLE` when the dynamic row format is used.

All values for columns used in indexes are stored in fixed format at the first block of the row, then the following columns are handled with `DRF`.

This sets two restrictions to tables:

- the order of the fields and therefore,
- the minimum size of the block used in the table.

### ORDERING OF COLUMNS

The columns used in fixed format must be defined before the dynamic ones in the `CREATE TABLE` statement. If this requirement is not met, the engine will not be able to add blocks to the set for these fields and they will be treated as fixed.

## MINIMUM BLOCK SIZE

The block size has to be big enough to store all fixed-length information in the first block. If not, the `CREATE TABLE` or `ALTER TABLE` statements will fail (see below).

## Limitations

*MyISAM* tables are still used for query optimizer internal temporary tables where the `MEMORY` tables could be used now instead: for temporary tables containing large `VARCHAR`, `BLOB`, and `TEXT` columns.

## Setting row format

Taking the restrictions into account, the *Improved MEMORY Storage Engine* will choose `DRF` over `FRF` at the moment of creating the table according to following criteria:

- There is an implicit request of the user in the column types **OR**
- There is an explicit request of the user **AND** the overhead incurred by `DRF` is beneficial.

## IMPLICIT REQUEST

The implicit request by the user is taken when there is at least one `BLOB` or `TEXT` column in the table definition. If there are none of these columns and no relevant option is given, the engine will choose `FRF`.

For example, this will yield the use of the dynamic format:

```
mysql> CREATE TABLE t1 (f1 VARCHAR(32), f2 TEXT, PRIMARY KEY (f1)) ENGINE=HEAP;
```

While this will not:

```
mysql> CREATE TABLE t1 (f1 VARCHAR(16), f2 VARCHAR(16), PRIMARY KEY (f1))
ENGINE=HEAP;
```

## EXPLICIT REQUEST

The explicit request is set with one of the following options in the `CREATE TABLE` statement:

- `KEY_BLOCK_SIZE = <value>`
- Requests the `DRF` with the specified block size (in bytes)

Despite its name, the `KEY_BLOCK_SIZE` option refers to a block size used to store data rather than indexes. The reason for this is that an existing `CREATE TABLE` option is reused to avoid introducing new ones.

*The Improved MEMORY Engine* checks whether the specified block size is large enough to keep all key column values. If it is too small, table creation will abort with an error.

After `DRF` is requested explicitly and there are no `BLOB` or `TEXT` columns in the table definition, the *Improved MEMORY Engine* will check if using the dynamic format provides any space saving benefits as compared to the fixed one:

- if the fixed row length is less than the dynamic block size (plus the dynamic row overhead - platform dependent) **OR**
- there isn't any variable-length columns in the table or `VARCHAR` fields are declared with length 31 or less,

the engine will revert to the fixed format as it is more space efficient in such case. The row format being used by the engine can be checked using `SHOW TABLE STATUS`.

## Examples

On a 32-bit platform:

```
mysql> CREATE TABLE t1 (f1 VARCHAR(32), f2 VARCHAR(32), f3 VARCHAR(32), f4
VARCHAR(32), PRIMARY KEY (f1)) KEY_BLOCK_SIZE=124 ENGINE=HEAP;

mysql> SHOW TABLE STATUS LIKE 't1';
```

### Expected output

Name	Engine	Version	Rows	Avg_row_length	Data_length	Max_data_length	Index_length	Data_free	Auto_increment	Create_time	Update_time	Check_time	Collation	Checksum	Create_options	Comment
t1	MEMORY	10	X	0	0	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	latin1_swedish_ci	NULL	row_format=DYNAMIC KEY_BLOCK_SIZE=124

On a 64-bit platform:

```
mysql> CREATE TABLE t1 (f1 VARCHAR(32), f2 VARCHAR(32), f3 VARCHAR(32), f4
VARCHAR(32), PRIMARY KEY (f1)) KEY_BLOCK_SIZE=124 ENGINE=HEAP;
```

```
mysql> SHOW TABLE STATUS LIKE 't1';
```

### Expected output

Name	Engine	Version	Rows	Avg_row_length	Data_length	Max_data_length	Index_length	Data_free	Auto_increment	Create_time	Update_time	Check_time	Collation	Checksum	Create_options	Comment
t1	MEMORY	10	X	0	0	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	latin1_swedish_ci	NULL	KEY_BLOCK_SIZE=124

## Implementation details

MySQL MEMORY tables keep data in arrays of fixed-size chunks. These chunks are organized into two groups of HP\_BLOCK structures:

- group1 contains indexes, with one HP\_BLOCK per key (part of HP\_KEYDEF),
- group2 contains record data, with a single HP\_BLOCK for all records.

While columns used in indexes are usually small, other columns in the table may need to accommodate larger data. Typically, larger data is placed into VARCHAR or BLOB columns.

The Improved MEMORY Engine implements the concept of dataspace, HP\_DATASPACE, which incorporates the HP\_BLOCK structures for the record data, adding more information for managing variable-sized records.

Variable-size records are stored in multiple “chunks”, which means that a single record of data (a database “row”) can consist of multiple chunks organized into one “set”, contained in HP\_BLOCK structures.

In variable-size format, one record is represented as one or many chunks depending on the actual data, while in fixed-size mode, one record is always represented as one chunk. The index structures would always point to the first chunk in the chunkset.

Variable-size records are necessary only in the presence of variable-size columns. The Improved Memory Engine will be looking for BLOB or VARCHAR columns with a declared length of 32 or more. If no such columns are found, the table will be switched to the fixed-size format. You should always put such columns at the end of the table definition in order to use the variable-size format.

Whenever data is being inserted or updated in the table, the Improved Memory Engine will calculate how many chunks are necessary.

For INSERT operations, the engine only allocates new chunksets in the recordspace. For UPDATE operations it will modify the length of the existing chunkset if necessary, unlinking unnecessary chunks at the end, or allocating and adding more if a larger length is needed.

When writing data to chunks or copying data back to a record, fixed-size columns are copied in their full format, while VARCHAR and BLOB columns are copied based on their actual length, skipping any NULL values.

When allocating a new chunkset of N chunks, the engine will try to allocate chunks one-by-one, linking them as they become allocated. For allocating a single chunk, it will attempt to reuse a deleted (freed) chunk. If no free chunks are available, it will try to allocate a new area inside a HP\_BLOCK.

When freeing chunks, the engine will place them at the front of a free list in the dataspace, each one containing a reference to the previously freed chunk.

The allocation and contents of the actual chunks varies between fixed and variable-size modes:

- Format of a fixed-size chunk:
  - `uchar[]` \* With `sizeof=chunk_dataspace_length`, but at least `sizeof(uchar*)` bytes. It keeps actual data or pointer to the next deleted chunk, where `chunk_dataspace_length` equals to full record length
  - `uchar` \* Status field (1 means “in use”, 0 means “deleted”)
- Format of a variable-size chunk:
  - `uchar[]` \* With `sizeof=chunk_dataspace_length`, but at least `sizeof(uchar*)` bytes. It keeps actual data or pointer to the next deleted chunk, where `chunk_dataspace_length` is set according to table’s `key_block_size`
  - `uchar*` \* Pointer to the next chunk in this chunkset, or NULL for the last chunk
  - `uchar` \* Status field (1 means “first”, 0 means “deleted”, 2 means “linked”)

Total chunk length is always aligned to the next `sizeof(uchar*)`.

#### See also

[Dynamic row format for MEMORY tables](#)

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 13.1.3 Improved InnoDB I/O scalability

Because *InnoDB* is a complex storage engine it must be configured properly in order to perform at its best. Some points are not configurable in standard *InnoDB*. The goal of this feature is to provide a more exhaustive set of options for *XtraDB*.

### System variables

`INNODB_FLUSH_METHOD`

Option	Description
Command-line	Yes
Config file	Yes

Option	Description
Scope	Global
Dynamic	No
Data type	Enumeration
Default	NULL
Allowed values	fsync, O_DSYNC, O_DIRECT, O_DIRECT

The following values are allowed:

- `fdatasync`: use `fsync()` to flush data, log, and parallel doublewrite files.
- `O_SYNC`: use `O_SYNC` to open and flush the log and parallel doublewrite files; use `fsync()` to flush the data files. Do not use `fsync()` to flush the parallel doublewrite file.
- `O_DIRECT`: use `O_DIRECT` to open the data files and `fsync()` system call to flush data, log, and parallel doublewrite files.
- `O_DIRECT_NO_FSYNC`: use `O_DIRECT` to open the data files and parallel doublewrite files, but does not use the `fsync()` system call to flush the data files, log files, and parallel doublewrite files. Do not use this option for the XFS file system.

#### Note

On an ext4 filesystem, set `innodb_log_write_ahead_size` to match the filesystem's write-ahead block size. This variable avoids `unaligned AIO/DIO` warnings.

### Status variables

The following information has been added to `SHOW ENGINE INNODB STATUS` to confirm the checkpointing activity:

```

The current checkpoint age target
The current age of the oldest page modification which has not been flushed to disk
yet.
The current age of the last checkpoint
...
---
LOG
---
Log sequence number 0 1059494372
Log flushed up to   0 1059494372
Last checkpoint at  0 1055251010
Max checkpoint age  162361775
Checkpoint age target 104630090
Modified age        4092465
Checkpoint age      4243362
0 pending log writes, 0 pending chkp writes
...

```

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 13.1.4 Enforcing storage engine

Percona Server for MySQL implements a variable to enforce the use of a specific storage engine.

When this variable is specified and a user tries to create a table using an explicit storage engine that is not the specified enforced engine, the user will get either an error if the `NO_ENGINE_SUBSTITUTION` SQL mode is enabled or a warning if `NO_ENGINE_SUBSTITUTION` is disabled and the table will be created anyway using the enforced engine (this is consistent with the default *MySQL* way of creating the default storage engine if other engines are not available unless `NO_ENGINE_SUBSTITUTION` is set).

In case user tries to enable `enforce_storage_engine` with engine that isn't available, system will not start.

### Note

If you're using `enforce_storage_engine`, you must either disable it before doing `mysql_upgrade` or perform `mysql_upgrade` with server started with `--skip-grants-tables`.

## System variables

`ENFORCE_STORAGE_ENGINE`

Option	Description
Command Line:	Yes
Config file	Yes
Scope:	Global
Dynamic:	No
Data type	String
Default value	NULL

This variable is not case sensitive.

## Example

Adding following option to `my.cnf` will start the server with InnoDB as enforced storage engine.

```
enforce_storage_engine=InnoDB
```

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-25

## 13.1.5 Extended show engine InnoDB status

This feature reorganizes the output of `SHOW ENGINE INNODB STATUS` to improve readability and to provide additional information. The variable `innodb_show_locks_held` controls the number of locks held to print for each *InnoDB* transaction.

This feature modified the `SHOW ENGINE INNODB STATUS` command as follows:

- Added extended information about *InnoDB* internal hash table sizes (in bytes) in the `BUFFER POOL AND MEMORY` section; also added buffer pool size in bytes.
- Added additional LOG section information.

### Other information

- Author / Origin: Baron Schwartz, <https://lists.mysql.com/internals/35174>

### System variables

`INNODB_SHOW_LOCKS_HELD`

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Data type	ULONG
Default	10
Range	0 - 1000

Specifies the number of locks held to print for each *InnoDB* transaction in `SHOW ENGINE INNODB STATUS`.

**INNODB\_PRINT\_LOCK\_WAIT\_TIMEOUT\_INFO**

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Data type	Boolean
Default	OFF

Makes *InnoDB* to write information about all lock wait timeout errors into the log file.

This allows to find out details about the failed transaction, and, most importantly, the blocking transaction. Query string can be obtained from `EVENTS_STATEMENTS_CURRENT` table, based on the `PROCESSLIST_ID` field, which corresponds to `thread_id` from the log output.

Taking into account that blocking transaction is often a multiple statement one, following query can be used to obtain blocking thread statements history:

```
SELECT s.SQL_TEXT FROM performance_schema.events_statements_history s
INNER JOIN performance_schema.threads t ON t.THREAD_ID = s.THREAD_ID
WHERE t.PROCESSLIST_ID = %d
UNION
SELECT s.SQL_TEXT FROM performance_schema.events_statements_current s
INNER JOIN performance_schema.threads t ON t.THREAD_ID = s.THREAD_ID
WHERE t.PROCESSLIST_ID = %d;
```

The `PROCESSLIST_ID` in this example is exactly the thread id from error log output.

## Status variables

The status variables here contain information available in the output of `SHOW ENGINE INNODB STATUS`, organized by the sections `SHOW ENGINE INNODB STATUS` displays. If you are familiar with the output of `SHOW ENGINE INNODB STATUS`, you will probably already recognize the information these variables contain.

### BACKGROUND THREAD

The following variables contain information in the `BACKGROUND THREAD` section of the output from `SHOW ENGINE INNODB STATUS`.

#### Expected output

```
-----
BACKGROUND THREAD
-----
srv_master_thread loops: 1 srv_active, 0 srv_shutdown, 11844 srv_idle
srv_master_thread log flush and writes: 11844
```

*InnoDB* has a source thread which performs background tasks depending on the server state, once per second. If the server is under workload, the source thread runs the following: performs background table drops; performs change buffer merge, adaptively; flushes the redo log to disk; evicts tables from the dictionary cache if needed to satisfy its size limit; makes a checkpoint. If the server is idle: performs background table drops, flushes and/or checkpoints the redo log if needed due to the checkpoint age; performs change buffer merge at full I/O capacity; evicts tables from the dictionary cache if needed; and makes a checkpoint.

**INNODB\_MASTER\_THREAD\_ACTIVE\_LOOPS**

Option	Description
Scope	Global
Data type	Numeric

This variable shows the number of times the above one-second loop was executed for active server states.

**INNODB\_MASTER\_THREAD\_IDLE\_LOOPS**

Option	Description
Scope	Global
Data type	Numeric

This variable shows the number of times the above one-second loop was executed for idle server states.

**INNODB\_BACKGROUND\_LOG\_SYNC**

Option	Description
Scope	Global
Data type	Numeric

This variable shows the number of times the *InnoDB* source thread has written and flushed the redo log.

**SEMAPHORES**

The following variables contain information in the **SEMAPHORES** section of the output from **SHOW ENGINE INNODB STATUS**. An example of that output is:

```

Expected output

-----
SEMAPHORES
-----
OS WAIT ARRAY INFO: reservation count 9664, signal count 11182
Mutex spin waits 20599, rounds 223821, OS waits 4479
RW-shared spins 5155, OS waits 1678; RW-excl spins 5632, OS waits 2592
Spin rounds per wait: 10.87 mutex, 15.01 RW-shared, 27.19 RW-excl

```

## INSERT BUFFER AND ADAPTIVE HASH INDEX

The following variables contain information in the `INSERT BUFFER AND ADAPTIVE HASH INDEX` section of the output from `SHOW ENGINE INNODB STATUS`. An example of that output is:

 **Expected output**

```
-----
INSERT BUFFER AND ADAPTIVE HASH INDEX
-----
Ibuf: size 1, free list len 6089, seg size 6091,
44497 inserts, 44497 merged recs, 8734 merges
0.00 hash searches/s, 0.00 non-hash searches/s
```

`INNODB_IBUF_FREE_LIST`

Option	Description
Scope	Global
Data type	Numeric

`INNODB_IBUF_SEGMENT_SIZE`

Option	Description
Scope	Global
Data type	Numeric

## LOG

The following variables contain information in the `LOG` section of the output from `SHOW ENGINE INNODB STATUS`. An example of that output is:

 **Expected output**

```
LOG
---
Log sequence number 10145937666
Log flushed up to   10145937666
Pages flushed up to 10145937666
Last checkpoint at  10145937666
Max checkpoint age  80826164
Checkpoint age target 78300347
Modified age        0
Checkpoint age      0
0 pending log writes, 0 pending chkp writes
9 log i/o's done, 0.00 log i/o's/second
Log tracking enabled
Log tracked up to   10145937666
Max tracked LSN age 80826164
```

`INNODB_LSN_CURRENT`

Option	Description
Scope	Global

Option	Description
Data type	Numeric

This variable shows the current log sequence number.

#### INNODB\_LSN\_FLUSHED

Option	Description
Scope	Global
Data type	Numeric

This variable shows the current maximum LSN that has been written and flushed to disk.

#### INNODB\_LSN\_LAST\_CHECKPOINT

Option	Description
Scope	Global
Data type	Numeric

This variable shows the LSN of the latest completed checkpoint.

#### INNODB\_CHECKPOINT\_AGE

Option	Description
Scope	Global
Data type	Numeric

This variable shows the current *InnoDB* checkpoint age, i.e., the difference between the current LSN and the LSN of the last completed checkpoint.

#### BUFFER POOL AND MEMORY

The following variables contain information in the `BUFFER POOL AND MEMORY` section of the output from `SHOW ENGINE INNODB STATUS`. An example of that output is:

### Expected output

```

-----
BUFFER POOL AND MEMORY
-----
Total memory allocated 137363456; in additional pool allocated 0
Total memory allocated by read views 88
Internal hash tables (constant factor + variable factor)
  Adaptive hash index 2266736      (2213368 + 53368)
  Page hash           139112 (buffer pool 0 only)
  Dictionary cache    729463 (554768 + 174695)
  File system         824800 (812272 + 12528)
  Lock system         333248 (332872 + 376)
  Recovery system     0        (0 + 0)
Dictionary memory allocated 174695
Buffer pool size      8191
Buffer pool size, bytes 134201344
Free buffers          7481
Database pages        707
Old database pages    280
Modified db pages     0
Pending reads         0
Pending writes: LRU 0, flush list 0 single page 0
Pages made young 0, not young 0
0.00 youngs/s, 0.00 non-youngs/s
Pages read 707, created 0, written 1
0.00 reads/s, 0.00 creates/s, 0.00 writes/s
No buffer pool page gets since the last printout
Pages read ahead 0.00/s, evicted without access 0.00/s, Random read ahead 0.00/s
LRU len: 707, unzip_LRU len: 0

```

#### INNODB\_MEM\_ADAPTIVE\_HASH

Option	Description
Scope	Global
Data type	Numeric

This variable shows the current size, in bytes, of the adaptive hash index.

#### INNODB\_MEM\_DICTIONARY

Option	Description
Scope	Global
Data type	Numeric

This variable shows the current size, in bytes, of the *InnoDB* in-memory data dictionary info.

#### INNODB\_MEM\_TOTAL

Option	Description
Scope	Global
Data type	Numeric

This variable shows the total amount of memory, in bytes, *InnoDB* has allocated in the process heap memory.

**INNODB\_BUFFER\_POOL\_PAGES\_LRU\_FLUSHED**

Option	Description
Scope	Global
Data type	Numeric

This variable shows the total number of buffer pool pages which have been flushed from the LRU list, i.e., too old pages which had to be flushed in order to make buffer pool room to read in new data pages.

**INNODB\_BUFFER\_POOL\_PAGES\_MADE\_NOT\_YOUNG**

Option	Description
Scope	Global
Data type	Numeric

This variable shows the number of times a buffer pool page was not marked as accessed recently in the LRU list because of `innodb_old_blocks_time` variable setting.

**INNODB\_BUFFER\_POOL\_PAGES\_MADE\_YOUNG**

Option	Description
Scope	Global
Data type	Numeric

This variable shows the number of times a buffer pool page was moved to the young end of the LRU list due to its access, to prevent its eviction from the buffer pool.

**INNODB\_BUFFER\_POOL\_PAGES\_OLD**

Option	Description
Scope	Global
Data type	Numeric

This variable shows the total number of buffer pool pages which are considered to be old according to the [Making the Buffer Pool Scan Resistant manual page](#).

**TRANSACTIONS**

The following variables contain information in the `TRANSACTIONS` section of the output from `SHOW INNODB STATUS`. An example of that output is:

### Expected output

```

-----
TRANSACTIONS
-----
Trx id counter F561FD
Purge done for trx's n:o < F561EB undo n:o < 0
History list length 19
LIST OF TRANSACTIONS FOR EACH SESSION:
---TRANSACTION 0, not started, process no 993, OS thread id 140213152634640
mysql thread id 15933, query id 32109 localhost root
show innodb status
---TRANSACTION F561FC, ACTIVE 29 sec, process no 993, OS thread id 140213152769808 updating or
deleting
mysql tables in use 1, locked 1

```

#### INNODB\_MAX\_TRX\_ID

Option	Description
Scope	Global
Data type	Numeric

This variable shows the next free transaction id number.

#### INNODB\_OLDEST\_VIEW\_LOW\_LIMIT\_TRX\_ID

Option	Description
Scope	Global
Data type	Numeric

This variable shows the highest transaction id, above which the current oldest open read view does not see any transaction changes. Zero if there is no open view.

#### INNODB\_PURGE\_TRX\_ID

Option	Description
Scope	Global
Data type	Numeric

This variable shows the oldest transaction id whose records have not been purged yet.

#### INNODB\_PURGE\_UNDO\_NO

Option	Description
Scope	Global
Data type	Numeric

### INFORMATION\_SCHEMA Tables

The following table contains information about the oldest active transaction in the system.

`INFORMATION_SCHEMA.XTRADB_READ_VIEW`

The data type for the following columns is `BIGINT UNSIGNED`. The columns contain 64-bit integers.

Column Name	Description
<code>'READ_VIEW_LOW_LIMIT_TRX_NUMBER'</code>	This is the highest transactions number.
<code>'READ_VIEW_UPPER_LIMIT_TRX_ID'</code>	This is the highest transactions ID at the time of the snapshot. This means that it should not see newer transactions than that value.
<code>'READ_VIEW_LOW_LIMIT_TRX_ID'</code>	This is the latest committed transaction ID at the time of the snapshot. This means that it should see transactions with IDs greater than or equal to that value.

The following table contains information about the memory usage for InnoDB/XtraDB hash tables.

`INFORMATION_SCHEMA.XTRADB_INTERNAL_HASH_TABLES`

Column Name	Description
<code>'INTERNAL_HASH_TABLE_NAME'</code>	Hash table name
<code>'TOTAL_MEMORY'</code>	Total amount of memory
<code>'CONSTANT_MEMORY'</code>	Constant memory
<code>'VARIABLE_MEMORY'</code>	Variable memory

### Other reading

- [SHOW INNODB STATUS walk through](#)
- [Table locks in SHOW INNODB STATUS](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 13.1.6 Show storage engines

This feature changes the comment field displayed when the `SHOW STORAGE ENGINES` command is executed and *XtraDB* is the storage engine.

Before the Change:

```
mysql> show storage engines;
```

**Expected output**

```

+-----+-----+-----+
+-----+-----+-----+
| Engine   | Support | Comment                                     |
Transactions | XA   | Savepoints |
+-----+-----+-----+
+-----+-----+-----+
| InnoDB   | YES     | Supports transactions, row-level locking, and foreign keys |
YES       | YES    | YES       |
...
+-----+-----+-----+
+-----+-----+-----+

```

After the Change:

```
mysql> show storage engines;
```

**Expected output**

```

+-----+-----+-----+-----+
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| Engine   | Support | Comment                                     | Transactions | XA
| Savepoints |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| InnoDB   | YES     | Percona-XtraDB, Supports transactions, row-level locking, and foreign
keys |         YES | YES | YES       |
...
+-----+-----+-----+-----+
+-----+-----+-----+-----+

```

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-25

## 13.1.7 Compressed columns with dictionaries

The `per-column compression` feature is a data type modifier, independent from user-level SQL and *InnoDB* data compression, that causes the data stored in the column to be compressed on writing to storage and decompressed on reading. For all other purposes, the data type is identical to the one

without the modifier, i.e. no new data types are created. Compression is done by using the `zlib` library.

Additionally, it is possible to pre-define a set of strings for each compressed column to achieve a better compression ratio on relatively small individual data items.

This feature provides:

- a better compression ratio for text data which consists of a large number of predefined words (e.g. JSON or XML) using compression methods with static dictionaries
- a way to select columns in the table to compress (in contrast to the *InnoDB* row compression method). This feature is based on a patch provided by Weixiang Zhai.

## Specifications

The feature is limited to InnoDB/XtraDB storage engine and to columns of the following data types:

- BLOB (including `TINYBLOB`, `MEDIUMBLOB`, `LOB`)
- TEXT (including `TINYTEXT`, `MEDIUMTEXT`, `LONGTEXT`)
- VARCHAR (including `NATIONAL VARCHAR`)
- VARBINARY
- JSON

A compressed column is declared by using the syntax that extends the existing `COLUMN_FORMAT` modifier: `COLUMN_FORMAT COMPRESSED`. If this modifier is applied to an unsupported column type or storage engine, an error is returned.

The compression can be specified:

- when creating a table: `CREATE TABLE ... (... , foo BLOB COLUMN_FORMAT COMPRESSED, ...);`
- when altering a table and modifying a column to the compressed format: `ALTER TABLE ... MODIFY [COLUMN] ... COLUMN_FORMAT COMPRESSED`, or `ALTER TABLE ... CHANGE [COLUMN] ... COLUMN_FORMAT COMPRESSED`.

Unlike Oracle MySQL, compression is applicable to generated stored columns. Use this syntax extension as follows:

```
mysql> CREATE TABLE t1(
  id INT,
  a BLOB,
  b JSON COLUMN_FORMAT COMPRESSED,
  g BLOB GENERATED ALWAYS AS (a) STORED COLUMN_FORMAT COMPRESSED WITH
  COMPRESSION_DICTIONARY numbers
) ENGINE=InnoDB;
```

To decompress a column, specify a value other than `COMPRESSED` to `COLUMN_FORMAT`: `FIXED`, `DYNAMIC`, or `DEFAULT`. If there is a column compression/decompression request in an `ALTER TABLE`, it is forced to the `COPY` algorithm.

Two new variables: `innodb_compressed_columns_zip_level` and `innodb_compressed_columns_threshold` have been implemented.

## Compression dictionary support

To achieve a better compression ratio on relatively small individual data items, it is possible to predefine a compression dictionary, which is a set of strings for each compressed column.

Compression dictionaries can be represented as a list of words in the form of a string (a comma or any other character can be used as a delimiter although not required). In other words, `a, bb, ccc`, `a bb ccc`, and `abbccc` will have the same effect. However, the latter is more compact. The Quote symbol quoting is handled by regular SQL quoting. The maximum supported dictionary length is 32506 bytes (`zlib` limitation).

The compression dictionary is stored in a new system *InnoDB* table. As this table is of the data dictionary kind, concurrent reads are allowed, but writes are serialized, and reads are blocked by writes. Table read through old read views are not supported, similar to *InnoDB* internal DDL transactions.

### INTERACTION WITH `INNODB_FORCE_RECOVERY` VARIABLE

Compression dictionary operations are treated like DDL operations with the exception when `innodb_force_value` is set to `3`: with values less than `3`, compression dictionary operations are allowed, and with values `>= 3`, they are forbidden.

### EXAMPLE

In order to use the compression dictionary, you need to create it. This can be done by running:

```
mysql> SET @dictionary_data = 'one' 'two' 'three' 'four';
```

#### Expected output

```
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> CREATE COMPRESSION_DICTIONARY numbers (@dictionary_data);
```

#### Expected output

```
Query OK, 0 rows affected (0.00 sec)
```

To create a table that has both compression and compressed dictionary support you should run:

```
mysql> CREATE TABLE t1(
  id INT,
  a BLOB COLUMN_FORMAT COMPRESSED,
  b BLOB COLUMN_FORMAT COMPRESSED WITH COMPRESSION_DICTIONARY numbers
) ENGINE=InnoDB;
```

The following example shows how to insert a sample of JSON data into the table:

```
SET @json_value =
'[\n'
```

```

' {\n'
' "one" = 0,\n'
' "two" = 0,\n'
' "three" = 0,\n'
' "four" = 0\n'
' },\n'
' {\n'
' "one" = 0,\n'
' "two" = 0,\n'
' "three" = 0,\n'
' "four" = 0\n'
' },\n'
' {\n'
' "one" = 0,\n'
' "two" = 0,\n'
' "three" = 0,\n'
' "four" = 0\n'
' },\n'
' {\n'
' "one" = 0,\n'
' "two" = 0,\n'
' "three" = 0,\n'
' "four" = 0\n'
' }\n'
']\n'
;

```

```

mysql> INSERT INTO t1 VALUES(0, @json_value, @json_value);
Query OK, 1 row affected (0.01 sec)

```

## INFORMATION\_SCHEMA Tables

This feature implements two new `INFORMATION_SCHEMA` tables.

### `INFORMATION_SCHEMA.COMPRESSION_DICTIONARY`

Column Name	Description
'BIGINT(21)_UNSIGNED dict_version'	'dictionary version'
'VARCHAR(64) dict_name'	'dictionary name'
'BLOB dict_data'	'compression dictionary string'

This table provides a view of the internal compression dictionary. The `SUPER` privilege is required to query it.

### `INFORMATION_SCHEMA.COMPRESSION_DICTIONARY_TABLES`

Column Name	Description
'BIGINT(21)_UNSIGNED table_schema'	'table schema'
'BIGINT(21)_UNSIGNED table_name'	'table ID from INFORMATION_SCHEMA'
'BIGINT(21)_UNSIGNED column_name'	'column position (starts from 0 as in INFORMATION_SCHEMA.INNODB_SYS_

Column Name	Description
'BIGINT(21)_UNSIGNED dict_name'	'dictionary ID'

This table provides a view over the internal table that stores the mapping between the compression dictionaries and the columns using them. The `SUPER` privilege is required to query it.

### Limitations

Compressed columns cannot be used in indices (neither on their own nor as parts of composite keys).

#### Note

`CREATE TABLE t2 AS SELECT * FROM t1` will create a new table with a compressed column, whereas `CREATE TABLE t2 AS SELECT CONCAT(a,') AS a FROM t1` will not create compressed columns.

At the same time, after executing the `CREATE TABLE t2 LIKE t1` statement, `t2.a` will have the `COMPRESSED` attribute.

`ALTER TABLE ... DISCARD/IMPORT TABLESPACE` is not supported for tables with compressed columns. To export and import tablespaces with compressed columns, you uncompress them first with: `ALTER TABLE ... MODIFY ... COLUMN_FORMAT DEFAULT.`

### mysqldump command line parameters

By default, with no additional options, `mysqldump` will generate a MySQL compatible SQL output.

All `/*!50633 COLUMN_FORMAT COMPRESSED */` and `/*!50633 COLUMN_FORMAT COMPRESSED WITH COMPRESSION_DICTIONARY <dictionary> */` won't be in the dump.

When a new option `enable-compressed-columns` is specified, all `/*!50633 COLUMN_FORMAT COMPRESSED */` will be left intact and all `/*!50633 COLUMN_FORMAT COMPRESSED WITH COMPRESSION_DICTIONARY <dictionary> */` will be transformed into `/*!50633 COLUMN_FORMAT COMPRESSED */`. In this mode, the dump will contain the necessary SQL statements to create compressed columns, but without dictionaries.

When a new `enable-compressed-columns-with-dictionaries` option is specified, dump will contain all compressed column attributes and compression dictionary.

Moreover, the following dictionary creation fragments will be added before `CREATE TABLE` statements which are going to use these dictionaries for the first time.

```
/*!50633 DROP COMPRESSION_DICTIONARY IF EXISTS <dictionary>; */
/*!50633 CREATE COMPRESSION_DICTIONARY <dictionary>(...); */
```

Two new options `add-drop-compression-dictionary` and `skip-add-drop-compression-dictionary` will control if `/*!50633 DROP COMPRESSION_DICTIONARY IF EXISTS <dictionary> */` part from previous paragraph will be skipped or not. By default, `add-drop-compression-dictionary` the mode will be used.

When both `enable-compressed-columns-with-dictionaries` and `--tab=<dir>` (separate file for each table) options are specified, necessary compression dictionaries will be created in each output file

using the following fragment (regardless of the values of `add-drop-compression-dictionary` and `skip-add-drop-compression-dictionary` options).

```
/*!50633 CREATE COMPRESSION_DICTIONARY IF NOT EXISTS <dictionary>(...); */
```

## System variables

### INNODB\_COMPRESSED\_COLUMNS\_ZIP\_LEVEL

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Data type	Numeric
Default	6
Range	0-9

This variable is used to specify the compression level used for compressed columns. Specifying `0` will use no compression, `1` the fastest, and `9` the best compression. The default value is `6`.

### INNODB\_COMPRESSED\_COLUMNS\_THRESHOLD

Option	Description
Command-line	Yes
Config file	Yes
Scope	Global
Dynamic	Yes
Data type	Numeric
Default	96
Range	1 - 2 <sup>64</sup> -1 (or 2 <sup>32</sup> -1 for 32-bit releases)

By default, a value being inserted will be compressed if its length exceeds `innodb_compressed_columns_threshold` bytes. Otherwise, it will be stored in the raw (uncompressed) form.

Please also note that because of the nature of some data, the compressed representation can be longer than the original value. In this case, it does not make sense to store such values in compressed form as *Percona Server for MySQL* would have to waste both memory space and CPU resources for unnecessary decompression. Therefore, even if the length of such non-compressible values exceeds `innodb_compressed_columns_threshold`, they will be stored in an uncompressed form (however, an attempt to compress them will still be made).

This parameter can be tuned to skip unnecessary attempts of data compression for values that are known in advance by the user to have a bad compression ratio of their first N bytes.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 13.1.8 InnoDB full-text search improvements

### 13.1.9 Ignoring stopword list

By default, all Full-Text Search indexes check the [stopwords list](#), to see if any indexed elements contain words on that list.

Using this list for n-gram indexes isn't always suitable, for example, any item that contains `a` or `i` will be ignored. Another word that can't be searched is `east`, this one will find no matches because `a` is on the FTS stopword list.

To resolve this issue, *Percona Server for MySQL* has the `innodb_ft_ignore_stopwords` variable to control whether *InnoDB* Full-Text Search should ignore the stopword list.

Although this variable is introduced to resolve n-gram issues, it affects all Full-Text Search indexes as well.

Being a stopword doesn't just mean being one of the predefined words from the list. Tokens shorter than `innodb_ft_min_token_size` or longer than `innodb_ft_max_token_size` are also considered stopwords. Therefore, when `innodb_ft_ignore_stopwords` is set to `ON` even for non-ngram FTS, `innodb_ft_min_token_size / innodb_ft_max_token_size` will be ignored meaning that in this case very short and very long words will also be indexed.

#### System variables

`INNODB_FT_IGNORE_STOPWORDS`

Option	Description
Command-line	Yes
Config file	Yes
Scope	Session, Global
Dynamic	Yes
Data type	Boolean
Default	OFF

When enabled, this variable will instruct *InnoDB* Full Text Search parser to ignore the stopword list when building/updating an FTS index.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 13.1.10 XtraDB performance improvements for I/O-bound highly-concurrent workloads

### Priority refill for the buffer pool free list

In highly-concurrent I/O-bound workloads the following situation may happen:

- Buffer pool free lists are used faster than they are refilled by the LRU cleaner thread.
- Buffer pool free lists become empty and more and more query and utility (i.e., purge) threads stall, checking whether a buffer pool free list has become non-empty, sleeping, performing single-page LRU flushes.
- The number of buffer pool free list mutex waiters increases.
- When the LRU manager thread (or a single page LRU flush by a query thread) finally produces a free page, it is starved from putting it on the buffer pool free list as it must acquire the buffer pool free list mutex too. However, being one thread in up to hundreds, the chances of a prompt acquisition are low.

This is addressed by delegating all the LRU flushes to the to the LRU manager thread, never attempting to evict a page or perform a LRU single page flush by a query thread, and introducing a backoff algorithm to reduce buffer pool free list mutex pressure on empty buffer pool free lists. This is controlled through a new system variable `innodb_empty_free_list_algorithm`.

#### INNODB\_EMPTY\_FREE\_LIST\_ALGORITHM

Option	Description
Command-line:	Yes
Config file:	Yes
Scope:	Global
Dynamic:	Yes
Data type:	legacy, backoff
Default	legacy

When `legacy` option is set, server will use the upstream algorithm and when the `backoff` is selected, Percona implementation will be used.

## Multi-threaded LRU flusher

Percona Server for MySQL features a true multi-threaded LRU flushing. In this scheme, each buffer pool instance has its own dedicated LRU manager thread that is tasked with performing LRU flushes and evictions to refill the free list of that buffer pool instance. Existing multi-threaded flusher no longer does any LRU flushing and is tasked with flush list flushing only.

- All threads still synchronize on each coordinator thread iteration. If a particular flushing job is stuck on one of the worker threads, the rest will idle until the stuck one completes.
- The coordinator thread heuristics focus on flush list adaptive flushing without considering the state of free lists, which might be in need of urgent refill for a subset of buffer pool instances on a loaded server.
- LRU flushing is serialized with flush list flushing for each buffer pool instance, introducing the risk that the right flushing mode will not happen for a particular instance because it is being flushed in the other mode.

The following InnoDB metrics are no longer accounted, as their semantics do not make sense under the current LRU flushing design: `buffer_LRU_batch_flush_avg_time_slot`, `buffer_LRU_batch_flush_avg_pass`, `buffer_LRU_batch_flush_avg_time_thread`, `buffer_LRU_batch_flush_avg_time_est`.

The need for InnoDB recovery thread writer threads is also removed, consequently all associated code is deleted.

### INNODB\_SCHED\_PRIORITY\_MASTER

Option	Description
Command-line:	Yes
Config file:	Yes
Scope:	Global
Dynamic:	Yes
Data type:	Boolean

This variable can be added to the configuration file.

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-09-25

### 13.1.11 Multiple page asynchronous I/O requests

The I/O unit size in *InnoDB* is only one page, even if the server doing read ahead. A 16KB I/O unit size is too small for sequential reads, and less efficient than a larger I/O unit size. *InnoDB* uses Linux asynchronous I/O (`aio`) by default. By submitting multiple, consecutive 16KB read requests at the same time, Linux internally merges the requests and reads more efficiently.

This feature is able to submit multiple page I/O requests and works in the background. You can manage the feature with the [linear read-ahead technique](#). This technique adds pages to the buffer pool based on the buffer pool pages being accessed sequentially. The `[innodb_read_ahead_threshold]` configuration parameter controls this operation.

On a HDD RAID 1+0 environment, more than 1000MB/s disk reads can be achieved by submitting 64 consecutive pages requests at once, while only 160MB/s disk reads is shown by submitting single page request.

#### Status variables

`INNODB_BUFFERED_AIO_SUBMITTED`

Option	Description
Scope:	Global
Data type:	Numeric

This variable shows the number of submitted buffered asynchronous I/O requests.

#### Other reading

- [Making full table scan 10x faster in InnoDB](#)
- [Bug #68659 InnoDB Linux native aio should submit more i/o requests at once](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

### 13.1.12 Prefix index queries optimization

*Percona Server for MySQL* has ported Prefix Index Queries Optimization feature from Facebook patch for MySQL.

Prior to this *InnoDB* would always fetch the clustered index for all prefix columns in an index, even when the value of a particular record was smaller than the prefix length. This implementation optimizes that case to use the record from the secondary index and avoid the extra lookup.

### Status variables

`INNODB_SECONDARY_INDEX_TRIGGERED_CLUSTER_READS`

Option	Description
Scope:	Global
Data type:	Numeric

This variable shows the number of times secondary index lookup triggered cluster lookup.

`INNODB_SECONDARY_INDEX_TRIGGERED_CLUSTER_READS_AVOIDED`

Option	Description
Scope:	Global
Data type:	Numeric

This variable shows the number of times prefix optimization avoided triggering cluster lookup.

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)    [:percona–logo: Get a Percona Expert](#)

Last update: 2023-09-25

## 13.1.13 Limit the estimation of records in a Query

### Important

This feature is a [tech preview](#). Before using this feature in production, we recommend that you test restoring production from physical backups in your environment, and also use the alternative backup method for redundancy.

This page describes an alternative when running queries against a large number of table partitions. When a query runs, InnoDB estimates the records in each partition. This process can result in more pages read and more disk I/O, if the buffer pool must fetch the pages from disk. This process increases the query time if there are a large number of partitions.

The addition of two variables makes it possible to override `records_in_range` which effectively bypasses the process.

**Warning**

The use of these variables may result in improper index selection by the optimizer.

**innodb\_records\_in\_range**

Option	Description
Command-line:	<code>--innodb-records-in-range</code>
Scope:	Global
Dynamic:	Yes
Data type:	Numeric
Default	0

**Important**

This feature is a [tech preview](#). Before using this feature in production, we recommend that you test restoring production from physical backups in your environment, and also use the alternative backup method for redundancy.

The variable provides a method to limit the number of records estimated for a query.

```
mysql> SET @@GLOBAL.innodb_records_in_range=100;
100
```

**INNODB\_FORCE\_INDEX\_RECORDS\_IN\_RANGE**

Option	Description
Command-line:	<code>--innodb-force-index-records-in-range</code>
Scope:	Global
Dynamic:	Yes
Data type:	Numeric
Default	0

**Important**

This feature is a [tech preview](#). Before using this feature in production, we recommend that you test restoring production from physical backups in your environment, and also use the alternative backup method for redundancy.

This variable provides a method to override the `records_in_range` result when a FORCE INDEX is used in a query.

```
mysql> SET @@GLOBAL.innodb_force_index_records_in_range=100;
100
```

## Using the favor\_range\_scan optimizer switch

### Important

This feature is a [tech preview](#). Before using this feature in production, we recommend that you test restoring production from physical backups in your environment, and also use the alternative backup method for redundancy.

In specific scenarios, the optimizer chooses to scan a table instead of using a range scan. The conditions are the following:

- Table with an extremely large number of rows
- Compound primary keys made of two or more columns
- WHERE clause contains multiple range conditions

The [optimizer\\_switch](#) controls the optimizer behavior. The favor\_range\_scan switch arbitrarily lowers the cost of a range scan by a factor of 10.

The available values are:

- ON
- OFF (Default)
- DEFAULT

```
mysql> SET optimizer_switch='favor_range_scan=on';
```

### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 13.1.14 InnoDB page fragmentation counters

*InnoDB* page fragmentation is caused by random insertion or deletion from a secondary index. This means that the physical ordering of the index pages on the disk is not same as the index ordering of the records on the pages. As a consequence this means that some pages take a lot more space and that queries which require a full table scan can take a long time to finish.

To provide more information about the *InnoDB* page fragmentation *Percona Server for MySQL* now provides the following counters as status variables: `Innodb_scan_pages_contiguous`, `Innodb_scan_pages_disjointed`, `Innodb_scan_data_size`, `Innodb_scan_deleted_recs_size`, and `Innodb_scan_pages_total_seek_distance`.

## Status variables

### INNODB\_SCAN\_PAGES\_CONTIGUOUS

Option	Description
Scope	Session
Data type	Numeric

This variable shows the number of contiguous page reads inside a query.

### INNODB\_SCAN\_PAGES\_DISJOINTED

Option	Description
Scope	Session
Data type	Numeric

This variable shows the number of disjointed page reads inside a query.

### INNODB\_SCAN\_DATA\_SIZE

Option	Description
Scope	Session
Data type	Numeric

This variable shows the size of data in all *InnoDB* pages read inside a query (in bytes) - calculated as the sum of `page_get_data_size(page)` for every page scanned.

### INNODB\_SCAN\_DELETED\_RECS\_SIZE

Option	Description
Scope	Session
Data type	Numeric

This variable shows the size of deleted records (marked as `deleted` in `page_delete_rec_list_end()`) in all *InnoDB* pages read inside a query (in bytes) - calculated as the sum of `page_header_get_field(page, PAGE_GARBAGE)` for every page scanned.

### INNODB\_SCAN\_PAGES\_TOTAL\_SEEK\_DISTANCE

Option	Description
Scope	Session
Data type	Numeric

This variable shows the total seek distance when moving between pages.

## Related reading

- [InnoDB: look after fragmentation](#)
- [Defragmenting a Table](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 13.2 Percona MyRocks

### 13.2.1 Percona MyRocks introduction

[MyRocks](#) is a storage engine for [MySQL](#) based on [RocksDB](#), an embeddable, persistent key-value store. *Percona MyRocks* is an implementation for [Percona Server for MySQL](#).

The RocksDB store is based on the log-structured merge-tree (or LSM tree). It is optimized for fast storage and combines outstanding space and write efficiency with acceptable read performance. As a result, MyRocks has the following advantages compared to other storage engines, if your workload uses fast storage, such as SSD:

- Requires less storage space
- Provides more storage endurance
- Ensures better IO capacity

[Percona MyRocks Installation Guide](#)

[MyRocks Limitations](#)

[Differences between Percona MyRocks and Facebook MyRocks](#)

[MyRocks Column Families](#)

[MyRocks Server Variables](#)

[MyRocks Information Schema Tables](#)

[Performance Schema MyRocks changes](#)

## Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-08-17

## 13.2.2 Percona MyRocks installation guide

Percona MyRocks is distributed as a separate package that can be enabled as a plugin for Percona Server for MySQL 8.1 and later versions.

### Note

File formats across different MyRocks variants may not be compatible. *Percona Server for MySQL* supports only *Percona MyRocks*. Migrating from one variant to another requires a logical data dump and reload.

- [Install Percona MyRocks](#)
- [Remove Percona MyRocks](#)

### Install Percona MyRocks

It is recommended to install Percona software from official repositories:

1. Configure Percona repositories as described in [Percona Software Repositories Documentation](#).
2. Install Percona MyRocks using the corresponding package manager:
  - For Debian or Ubuntu:

```
$ sudo apt install percona-server-rocksdb
```

- For RHEL or CentOS:

```
$ sudo yum install percona-server-rocksdb
```

After installation, you should see the following output:

### Expected output

```
* This release of |Percona Server| is distributed with RocksDB storage engine.
* Run the following script to enable the RocksDB storage engine in Percona Server:
```

```
$ ps-admin --enable-rocksdb -u <mysql_admin_user> -p[mysql_admin_pass] [-S <socket>]
[-h <host> -P <port>]
```

#### ENABLE MYROCKS WITH PS-ADMIN

Run the `ps-admin` script as system root user or with **sudo** and provide the MySQL root user credentials to properly enable the RocksDB (MyRocks) storage engine:

```
$ sudo ps-admin --enable-rocksdb -u root -pPassw0rd
```

#### Expected output

```
Checking if RocksDB plugin is available for installation ...
INFO: ha_rocksdb.so library for RocksDB found at /usr/lib64/mysql/plugin/ha_rocksdb.so.

Checking RocksDB engine plugin status...
INFO: RocksDB engine plugin is not installed.

Installing RocksDB engine...
INFO: Successfully installed RocksDB engine plugin.
```

#### Note

When you use the `ps-admin` script to enable Percona MyRocks, it performs the following:

- Disables Transparent huge pages
- Installs and enables the RocksDB plugin

If the script returns no errors, Percona MyRocks should be successfully enabled on the server. You can verify it as follows:

```
mysql> SHOW ENGINES;
```

#### Expected output

```
+-----+-----+-----+
+-----+-----+-----+
| Engine | Support | Comment
| Transactions | XA | Savepoints |
+-----+-----+-----+
+-----+-----+-----+
| ROCKSDB | YES | RocksDB storage engine
| YES | YES | YES |
...
| InnoDB | DEFAULT | Percona-XtraDB, Supports transactions, row-level locking, and foreign keys
| YES | YES | YES |
+-----+-----+-----+
+-----+-----+-----+
10 rows in set (0.00 sec)
```

Note that the RocksDB engine is not set to be default, new tables will still be created using the InnoDB (XtraDB) storage engine. To make RocksDB storage engine default, set `default-storage-engine=rocksdb` in the `[mysqld]` section of `my.cnf` and restart *Percona Server for MySQL*.

Alternatively, you can add `ENGINE=RocksDB` after the `CREATE TABLE` statement for every table that you create.

#### INSTALL MYROCKS PLUGINS

You can install MyRocks manually with a series of `INSTALL PLUGIN` statements. You must have the `INSERT` privilege for the `mysql.plugin` system table.

The following statements install MyRocks:

```
INSTALL PLUGIN ROCKSDB SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_CFSTATS SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_DBSTATS SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_PERF_CONTEXT SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_PERF_CONTEXT_GLOBAL SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_CF_OPTIONS SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_GLOBAL_INFO SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_COMPACTION_HISTORY SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_COMPACTION_STATS SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_ACTIVE_COMPACTION_STATS SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_DDL SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_INDEX_FILE_MAP SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_LOCKS SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_TRX SONAME 'ha_rocksdb.so';
INSTALL PLUGIN ROCKSDB_DEADLOCK SONAME 'ha_rocksdb.so';
```

#### Remove Percona MyRocks

It will not be possible to access tables created using the RocksDB engine with another storage engine after you remove Percona MyRocks. If you need this data, alter the tables to another storage engine. For example, to alter the `City` table to InnoDB, run the following:

```
mysql> ALTER TABLE City ENGINE=InnoDB;
```

To disable and uninstall the RocksDB engine plugins, use the `ps-admin` script as follows:

```
$ sudo ps-admin --disable-rocksdb -u root -pPassw0rd
```

#### Expected output

```
Checking RocksDB engine plugin status...
INFO: RocksDB engine plugin is installed.

Uninstalling RocksDB engine plugin...
INFO: Successfully uninstalled RocksDB engine plugin.
```

After the engine plugins have been uninstalled, remove the Percona MyRocks package:

- For Debian or Ubuntu:

```
$ sudo apt remove percona-server-rocksdb-8.1
```

- For RHEL or CentOS:

```
$ sudo yum remove percona-server-rocksdb-80.x86_64
```

Finally, remove all the [MyRocks Server Variables](#) from the configuration file (`my.cnf`) and restart *Percona Server for MySQL*.

#### UNINSTALL MYROCKS PLUGINS

You can [uninstall the plugins](#) for MyRocks. You must have the `DELETE` privilege for the `mysql.plugin` system table.

The following statements remove the MyRocks plugins:

```
UNINSTALL PLUGIN ROCKSDB;
UNINSTALL PLUGIN ROCKSDB_CFSTATS;
UNINSTALL PLUGIN ROCKSDB_DBSTATS;
UNINSTALL PLUGIN ROCKSDB_PERF_CONTEXT;
UNINSTALL PLUGIN ROCKSDB_PERF_CONTEXT_GLOBAL;
UNINSTALL PLUGIN ROCKSDB_CF_OPTIONS;
UNINSTALL PLUGIN ROCKSDB_GLOBAL_INFO;
UNINSTALL PLUGIN ROCKSDB_COMPACTON_HISTORY;
UNINSTALL PLUGIN ROCKSDB_COMPACTON_STATS;
UNINSTALL PLUGIN ROCKSDB_ACTIVE_COMPACTON_STATS;
UNINSTALL PLUGIN ROCKSDB_DDL;
UNINSTALL PLUGIN ROCKSDB_INDEX_FILE_MAP;
UNINSTALL PLUGIN ROCKSDB_LOCKS;
UNINSTALL PLUGIN ROCKSDB_TRX;
UNINSTALL PLUGIN ROCKSDB_DEADLOCK;
```

#### Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

### 13.2.3 Updated supported features

The following is a list of the latest supported features:

- Percona Server for MySQL supports `SELECT FOR UPDATE SKIP LOCKED/NOWAIT`. The transaction isolation level must be `READ COMMITTED`.
- Percona Server for MySQL adds the ability to cancel ongoing manual compactons. The cancel methods are the following:
  - Using either `Control+C` (from a session) or `KILL` (from another session) for client sessions running manual compactons by `SET GLOBAL rocksdb_compact_cf (variable)`.
  - Using a global variable `rocksdb_cancel_manual_compactions` to cancel all ongoing manual compactons.
- Percona Server for MySQL adds supported for [Generated Columns](#) and index are supported.
- Percona Server for MySQL adds support for [explicit DEFAULT value expressions](#).

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 13.2.4 MyRocks limitations

The MyRocks storage engine lacks the following features compared to InnoDB:

- [Online DDL](#) is not supported due to the lack of atomic DDL support.
  - There is no `ALTER TABLE ... ALGORITHM=INSTANT` functionality
  - A partition management operation only supports the `COPY` algorithms, which rebuilds the partition table and moves the data based on the new `PARTITION ... VALUE` definition. In the case of `DROP PARTITION`, the data not moved to another partition is deleted.
- [ALTER TABLE .. EXCHANGE PARTITION.](#)
- [SAVEPOINT](#)
- [Transportable tablespace](#)
- [Foreign keys](#)
- [Spatial indexes](#)
- [Fulltext indexes](#)
- [Gap locks](#)
- [Group Replication](#)
- [Partial Update of LOB in InnoDB](#)

You should also consider the following:

- All collations are supported on `CHAR` and `VARCHAR` indexed columns. By default, MyRocks prevents creating indexes with non-binary collations (including `latin1`). You can optionally use it by setting `rocksdb_strict_collation_exceptions` to `t1` (table names with regex format), but non-binary covering indexes other than `latin1` (excluding `german1`) still require a primary key lookup to return the `CHAR` or `VARCHAR` column.
- Either `ORDER BY DESC` or `ORDER BY ASC` is slow. This is because of “Prefix Key Encoding” feature in RocksDB. See <https://www.slideshare.net/matsunobu/myrocks-deep-dive/58> for details. By default, ascending scan is faster and descending scan is slower. If the “reverse column family” is configured, then descending scan will be faster and ascending scan will be slower. Note that InnoDB also imposes a cost when the index is scanned in the opposite order.
- When converting from large MyISAM/InnoDB tables, either by using the `ALTER` or `INSERT INTO ... SELECT` statements it’s recommended that you check the [Data loading](#) documentation and create MyRocks tables as below (in case the table is sufficiently big it will cause the server to consume all the memory and then be terminated by the OOM killer):

```
SET session sql_log_bin=0;
SET session rocksdb_bulk_load=1;
ALTER TABLE large_myisam_table ENGINE=RocksDB;
SET session rocksdb_bulk_load=0;
```

### Expected output

```
.. warning::

If you are loading large data without enabling :ref:`rocksdb_bulk_load`
or :ref:`rocksdb_commit_in_the_middle`, please make sure transaction
size is small enough. All modifications of the ongoing transactions are
kept in memory.
```

- With partitioned tables that use the *MyRocks* storage engine, the upgrade only works with native partitioning.

### See also

MySQL Documentation: [Preparing Your Installation for Upgrade](#)

- Percona Server for MySQL 8.1 and Unicode 9.0.0 standards have defined a change in the handling of binary collations. These collations are handled as NO PAD, trailing spaces are included in key comparisons. A binary collation comparison may result in two unique rows inserted and does not generate a `DUP\_ENTRY` error. MyRocks key encoding and comparison does not account for this character set attribute.

## Not supported on MyRocks

MyRocks does not support the following:

- Operating as either a source or a replica in any replication topology that is not exclusively row-based. Statement-based and mixed-format binary logging is not supported. For more information, see [Replication Formats](#).
- Using [multi-valued indexes](#). InnoDB supports this feature.
- Using [spatial data types](#).
- Using the [Clone Plugin](#) and the Clone Plugin API. InnoDB supports either these features.
- Using encryption in tables. At this time, during an `ALTER TABLE` operation, MyRocks mistakenly detects all InnoDB tables as encrypted. Therefore, any attempt to `ALTER` an InnoDB table to MyRocks fails.

As a workaround, we recommend a manual move of the table. The following steps are the same as the `ALTER TABLE ... ENGINE=...` process:

- Use `SHOW CREATE TABLE ...` to return the InnoDB table definition.
- With the table definition as the source, perform a `CREATE TABLE ... ENGINE=RocksDB`.
- In the new table, use `INSERT INTO <new table> SELECT \* FROM <old table>`.

### Note

With MyRocks and with large tables, it is recommended to set the session variable `rocksdb_bulk_load=1` during the load to prevent running out of memory. This recommendation is because of the MyRocks large transaction limitation. For more information, see [MyRocks Data Loading](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-26

### 13.2.5 Differences between Percona MyRocks and Facebook MyRocks

The original MyRocks was developed by Facebook and works with their implementation of MySQL. *Percona MyRocks* is a branch of MyRocks for *Percona Server for MySQL* and includes the following differences from the original implementation:

- The behavior of the `START TRANSACTION WITH CONSISTENT SNAPSHOT` statement depends on the [transaction isolation level](#).

Storage Engine	Transaction isolation level	
	READ COMMITTED	REPEATABLE
InnoDB	Success	Success
Facebook MyRocks	Fail	Success MyRocks
Percona MyRocks	Fail with any DML which would violate the read-only snapshot constraint	Success the engine

- Percona MyRocks includes the `lz4` and `zstd` statically linked libraries.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-07-27

### 13.2.6 MyRocks Information Schema tables

When you install the MyRocks plugin for *MySQL*, the Information Schema is extended to include the following tables:

**ROCKSDB\_GLOBAL\_INFO**

## COLUMNS

Column Name	Type
TYPE	varchar(513)
NAME	varchar(513)
VALUE	varchar(513)

**ROCKSDB\_CFSTATS**

## COLUMNS

Column Name	Type
CF_NAME	varchar(193)
STAT_TYPE	varchar(193)
VALUE	bigint(8)

**ROCKSDB\_TRX**

This table stores mappings of RocksDB transaction identifiers to *MySQL* client identifiers to enable associating a RocksDB transaction with a *MySQL* client operation.

## COLUMNS

Column Name	Type
TRANSACTION_ID	bigint(8)
STATE	varchar(193)
NAME	varchar(193)
WRITE_COUNT	bigint(8)
LOCK_COUNT	bigint(8)
TIMEOUT_SEC	int(4)
WAITING_KEY	varchar(513)
WAITING_COLUMN_FAMILY_ID	int(4)
IS_REPLICATION	int(4)
SKIP_TRX_API	int(4)
READ_ONLY	int(4)
HAS_DEADLOCK_DETECTION	int(4)
NUM_ONGOING_BULKLOAD	int(4)
THREAD_ID	int(8)
QUERY	varchar(193)

**ROCKSDB\_CF\_OPTIONS**

## COLUMNS

Column Name	Type
CF_NAME	varchar(193)
OPTION_TYPE	varchar(193)
VALUE	varchar(193)

**ROCKSDB\_ACTIVE\_COMPACTION\_STATS**

## COLUMNS

Column Name	Type
THREAD_ID	bigint
CF_NAME	varchar(193)
INPUT_FILES	varchar(513)
OUTPUT_FILES	varchar(513)
COMPACTION_REASON	varchar(513)

**ROCKSDB\_COMPACTION\_HISTORY**

## COLUMNS

Column Name	Type
THREAD_ID	bigint
CF_NAME	varchar(513)
INPUT_LEVEL	integer
OUTPUT_LEVEL	integer
INPUT_FILES	varchar(513)
OUTPUT_FILES	varchar(513)
COMPACTION_REASON	varchar(513)
START_TIMESTAMP	bigint
END_TIMESTAMP	bigint

**ROCKSDB\_COMPACTION\_STATS**

## COLUMNS

Column Name	Type
CF_NAME	varchar(193)
LEVEL	varchar(513)
TYPE	varchar(513)

Column Name	Type
VALUE	double

**ROCKSDB\_DBSTATS**

## COLUMNS

Column Name	Type
STAT_TYPE	varchar(193)
VALUE	bigint(8)

**ROCKSDB\_DDL**

## COLUMNS

Column Name	Type
TABLE_SCHEMA	varchar(193)
TABLE_NAME	varchar(193)
PARTITION_NAME	varchar(193)
INDEX_NAME	varchar(193)
COLUMN_FAMILY	int(4)
INDEX_NUMBER	int(4)
INDEX_TYPE	smallint(2)
KV_FORMAT_VERSION	smallint(2)
TTL_DURATION	bigint(8)
INDEX_FLAGS	bigint(8)
CF	varchar(193)
AUTO_INCREMENT	bigint(8) unsigned

**ROCKSDB\_INDEX\_FILE\_MAP**

## COLUMNS

Column Name	Type
COLUMN_FAMILY	int(4)
INDEX_NUMBER	int(4)
SST_NAME	varchar(193)
NUM_ROWS	bigint(8)
DATA_SIZE	bigint(8)
ENTRY_DELETES	bigint(8)
ENTRY_SINGLEDELETES	bigint(8)

Column Name	Type
ENTRY_MERGES	bigint(8)
ENTRY_OTHERS	bigint(8)
DISTINCT_KEYS_PREFIX	varchar(400)

**ROCKSDB\_LIVE\_FILES\_METADATA**

Column Name	Type
CF_NAME	varchar(193)
LEVEL	varchar(513)
NAME	varchar(513)
DB_PATH	varchar(513)
FILE_NUMBER	bigint
FILE_TYPE	varchar(193)
SIZE	bigint
RELATIVE_FILENAME	varchar(193)
DIRECTORY	varchar(513)
TEMPERATURE	varchar(193)
FILE_CHECKSUM	varchar(513)
FILE_CHECKSUM_FUNC_NAME	varchar(193)
SMALLEST_SEQNO	bigint
LARGEST_SEQNO	bigint
SMALLEST_KEY	varchar(513)
LARGEST_KEY	varchar(513)
NUM_READS_SAMPLED	bigint
BEING_COMPACTED	tinyint
NUM_ENTRIES	bigint
NUM_DELETIONS	bigint
OLDEST_BLOB_FILE_NUMBER	bigint
OLDEST_ANCESTER_TIME	bigint
FILE_CREATION_TIME	bigint

**ROCKSDB\_LOCKS**

This table contains the set of locks granted to MyRocks transactions.

## COLUMNS

Column Name	Type
COLUMN_FAMILY_ID	int(4)
TRANSACTION_ID	int(4)
KEY	varchar(513)
MODE	varchar(32)

**ROCKSDB\_PERF\_CONTEXT**

## COLUMNS

Column Name	Type
TABLE_SCHEMA	varchar(193)
TABLE_NAME	varchar(193)
PARTITION_NAME	varchar(193)
STAT_TYPE	varchar(193)
VALUE	bigint(8)

**ROCKSDB\_PERF\_CONTEXT\_GLOBAL**

## COLUMNS

Column Name	Type
STAT_TYPE	varchar(193)
VALUE	bigint(8)

**ROCKSDB\_DEADLOCK**

This table records information about deadlocks.

## COLUMNS

Column Name	Type
DEADLOCK_ID	bigint(8)
TRANSACTION_ID	bigint(8)
CF_NAME	varchar(193)
WAITING_KEY	varchar(513)
LOCK_TYPE	varchar(193)
INDEX_NAME	varchar(193)
TABLE_NAME	varchar(193)
ROLLED_BACK	bigint(8)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-07-27

## 13.2.7 MyRocks server variables

The MyRocks server variables expose configuration of the underlying RocksDB engine. There several ways to set these variables:

- For production deployments, you should have all variables defined in the configuration file.
- *Dynamic* variables can be changed at runtime using the `SET` statement.
- If you want to test things out, you can set some of the variables when starting `mysqld` using corresponding command-line options.

If a variable was not set in either the configuration file or as a command-line option, the default value is used.

Also, all variables can exist in one or both of the following scopes:

- *Global* scope defines how the variable affects overall server operation.
- *Session* scope defines how the variable affects operation for individual client connections.

Variable Name
<code>rocksdb_access_hint_on_compaction_start</code>
<code>rocksdb_advise_random_on_open</code>
<code>rocksdb_allow_concurrent_memtable_write</code>
<code>rocksdb_allow_to_start_after_corruption</code>
<code>rocksdb_allow_mmap_reads</code>
<code>rocksdb_allow_mmap_writes</code>
<code>rocksdb_allow_unsafe_alter</code>
<code>rocksdb_alter_column_default_inplace</code>
<code>rocksdb_alter_table_comment_inplace</code>
<code>rocksdb_base_background_compactions</code>
<code>rocksdb_blind_delete_primary_key</code>
<code>rocksdb_block_cache_size</code>
<code>rocksdb_bulk_load_fail_if_not_bottommost_level</code>

Variable Name
<code>rocksdb_bulk_load_partial_index</code>
<code>rocksdb_bulk_load_use_sst_partitioner</code>
<code>rocksdb_block_restart_interval</code>
<code>rocksdb_block_size</code>
<code>rocksdb_block_size_deviation</code>
<code>rocksdb_bulk_load</code>
<code>rocksdb_bulk_load_allow_sk</code>
<code>rocksdb_bulk_load_allow_unsorted</code>
<code>rocksdb_bulk_load_size</code>
<code>rocksdb_bytes_per_sync</code>
<code>rocksdb_cache_dump</code>
<code>rocksdb_cache_high_pri_pool_ratio</code>
<code>rocksdb_cache_index_and_filter_blocks</code>
<code>rocksdb_cache_index_and_filter_with_high_priority</code>
<code>rocksdb_cancel_manual_compactions</code>
<code>rocksdb_charge_memory</code>
<code>rocksdb_checksums_pct</code>
<code>rocksdb_collect_sst_properties</code>
<code>rocksdb_column_default_value_as_expression</code>
<code>rocksdb_commit_in_the_middle</code>
<code>rocksdb_commit_time_batch_for_recovery</code>
<code>rocksdb_compact_cf</code>
<code>rocksdb_compaction_readahead_size</code>
<code>rocksdb_compaction_sequential_deletes</code>
<code>rocksdb_compaction_sequential_deletes_count_sd</code>
<code>rocksdb_compaction_sequential_deletes_file_size</code>
<code>rocksdb_compaction_sequential_deletes_window</code>
<code>rocksdb_concurrent_prepare</code>
<code>rocksdb_converter_record_cached_length</code>
<code>rocksdb_corrupt_data_action</code>
<code>rocksdb_create_checkpoint</code>
<code>rocksdb_create_if_missing</code>
<code>rocksdb_create_missing_column_families</code>
<code>rocksdb_create_temporary_checkpoint</code>

Variable Name
<code>rocksdb_datadir</code>
<code>rocksdb_db_write_buffer_size</code>
<code>rocksdb_deadlock_detect</code>
<code>rocksdb_deadlock_detect_depth</code>
<code>rocksdb_debug_cardinality_multipler</code>
<code>rocksdb_debug_manual_compaction_delay</code>
<code>rocksdb_debug_optimizer_no_zero_cardinality</code>
<code>rocksdb_debug_ttl_ignore_pk</code>
<code>rocksdb_debug_ttl_read_filter_ts</code>
<code>rocksdb_debug_ttl_rec_ts</code>
<code>rocksdb_debug_ttl_snapshot_ts</code>
<code>rocksdb_default_cf_options</code>
<code>rocksdb_delayed_write_rate</code>
<code>rocksdb_delete_cf</code>
<code>rocksdb_delete_obsolete_files_period_micros</code>
<code>rocksdb_disable_file_deletions</code>
<code>rocksdb_disable_instant_ddl</code>
<code>rocksdb_enable_bulk_load_api</code>
<code>rocksdb_enable_delete_range_for_drop_index</code>
<code>rocksdb_enable_insert_with_update_caching</code>
<code>rocksdb_enable_iterate_bounds</code>
<code>rocksdb_enable_pipelined_write</code>
<code>rocksdb_enable_remove_orphaned_dropped_cfs</code>
<code>rocksdb_enable_ttl</code>
<code>rocksdb_enable_ttl_read_filtering</code>
<code>rocksdb_enable_thread_tracking</code>
<code>rocksdb_enable_write_thread_adaptive_yield</code>
<code>rocksdb_error_if_exists</code>
<code>rocksdb_error_on_suboptimal_collation</code>
<code>rocksdb_flush_log_at_trx_commit</code>
<code>rocksdb_flush_memtable_on_analyze</code>
<code>rocksdb_force_compute_memtable_stats</code>
<code>rocksdb_force_compute_memtable_stats_cachetime</code>
<code>rocksdb_force_flush_memtable_and_lzero_now</code>

Variable Name
rocksdb_force_flush_memtable_now
rocksdb_force_index_records_in_range
rocksdb_hash_index_allow_collision
rocksdb_ignore_unknown_options
rocksdb_index_type
rocksdb_info_log_level
rocksdb_is_fd_close_on_exec
rocksdb_keep_log_file_num
rocksdb_large_prefix
rocksdb_lock_scanned_rows
rocksdb_lock_wait_timeout
rocksdb_log_file_time_to_roll
rocksdb_manifest_preallocation_size
rocksdb_manual_compaction_bottommost_level
rocksdb_manual_compaction_threads
rocksdb_manual_wal_flush
rocksdb_master_skip_tx_api
rocksdb_max_background_compactions
rocksdb_max_background_flushes
rocksdb_max_background_jobs
rocksdb_max_bottom_pri_background_compactions
rocksdb_max_compaction_history
rocksdb_max_latest_deadlocks
rocksdb_max_log_file_size
rocksdb_max_manifest_file_size
rocksdb_max_manual_compactions
rocksdb_max_open_files
rocksdb_max_row_locks
rocksdb_max_subcompactions
rocksdb_max_total_wal_size
rocksdb_merge_buf_size
rocksdb_merge_combine_read_size
rocksdb_merge_tmp_file_removal_delay_ms
rocksdb_new_table_reader_for_compaction_inputs

Variable Name
rocksdb_no_block_cache
rocksdb_no_create_column_family
rocksdb_override_cf_options
rocksdb_paranoid_checks
rocksdb_partial_index_sort_max_mem
rocksdb_pause_background_work
rocksdb_partial_index_blind_delete
rocksdb_perf_context_level
rocksdb_persistent_cache_path
rocksdb_persistent_cache_size_mb
rocksdb_pin_l0_filter_and_index_blocks_in_cache
rocksdb_print_snapshot_conflict_queries
rocksdb_protection_bytes_per_key
rocksdb_rate_limiter_bytes_per_sec
rocksdb_read_free_rpl
rocksdb_read_free_rpl_tables
rocksdb_records_in_range
rocksdb_reset_stats
rocksdb_rollback_on_timeout
rocksdb_rpl_skip_tx_api
rocksdb_seconds_between_stat_computes
rocksdb_signal_drop_index_thread
rocksdb_sim_cache_size
rocksdb_skip_bloom_filter_on_read
rocksdb_skip_fill_cache
rocksdb_skip_locks_if_skip_unique_check
rocksdb_sst_mgr_rate_bytes_per_sec
rocksdb_stats_dump_period_sec
rocksdb_stats_level
rocksdb_stats_recalc_rate
rocksdb_store_row_debug_checksums
rocksdb_strict_collation_check
rocksdb_strict_collation_exceptions
rocksdb_table_cache_numshardbits

Variable Name
rocksdb_table_stats_background_thread_nice_value
rocksdb_table_stats_max_num_rows_scanned
rocksdb_table_stats_recalc_threshold_count
rocksdb_table_stats_recalc_threshold_pct
rocksdb_table_stats_sampling_pct
rocksdb_table_stats_use_table_scan
rocksdb_tmpdir
rocksdb_two_write_queues
rocksdb_trace_block_cache_access
rocksdb_trace_queries
rocksdb_trace_sst_api
rocksdb_track_and_verify_wals_in_manifest
rocksdb_unsafe_for_binlog
rocksdb_update_cf_options
rocksdb_use_adaptive_mutex
rocksdb_use_default_sk_cf
rocksdb_use_direct_io_for_flush_and_compaction
rocksdb_use_direct_reads
rocksdb_use_fsync
rocksdb_use_hyper_clock_cache
rocksdb_use_write_buffer_manager
rocksdb_validate_tables
rocksdb_verify_row_debug_checksums
rocksdb_wal_bytes_per_sync
rocksdb_wal_dir
rocksdb_wal_recovery_mode
rocksdb_wal_size_limit_mb
rocksdb_wal_ttl_seconds
rocksdb_whole_key_filtering
rocksdb_write_batch_flush_threshold
rocksdb_write_batch_max_bytes
rocksdb_write_disable_wal
rocksdb_write_ignore_missing_column_families
rocksdb_write_policy

**ROCKSDB\_ACCESS\_HINT\_ON\_COMPACTON\_START**

Option	Description
Command-line	<code>-rocksdb-access-hint-on-compact</code>
Dynamic	No
Scope	Global
Data type	String or numeric
Default	NORMAL or 1

Specifies the file access pattern once a compaction is started, applied to all input files of a compaction. Possible values are:

- 0 = NONE
- 1 = NORMAL (default)
- 2 = SEQUENTIAL
- 3 = WILLNEED

**ROCKSDB\_ADVISE\_RANDOM\_ON\_OPEN**

Option	Description
Command-line	<code>-rocksdb-advise-random-on-open</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether to hint the underlying file system that the file access pattern is random, when a data file is opened. Enabled by default.

**ROCKSDB\_ALLOW\_CONCURRENT\_MEMTABLE\_WRITE**

Option	Description
Command-line	<code>-rocksdb-allow-concurrent-memtable</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to allow multiple writers to update memtables in parallel. Disabled by default.

**ROCKSDB\_ALLOW\_TO\_START\_AFTER\_CORRUPTION**

Option	Description
Command-line	<code>-rocksdb_allow_to_start_after_cor</code>

Option	Description
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to allow server to restart once MyRocks reported data corruption. Disabled by default.

Once corruption is detected server writes marker file (named ROCKSDB\_CORRUPTED) in the data directory and aborts. If marker file exists, then mysqld exits on startup with an error message. The restart failure will continue until the problem is solved or until mysqld is started with this variable turned on in the command line.

 **Note**

Not all memtables support concurrent writes.

#### ROCKSDB\_ALLOW\_MMAP\_READS

Option	Description
Command-line	<code>--rocksdb-allow-mmap-reads</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to allow the OS to map a data file into memory for reads. Disabled by default. If you enable this, make sure that `rocksdb_use_direct_reads` is disabled.

#### ROCKSDB\_ALLOW\_MMAP\_WRITES

Option	Description
Command-line	<code>--rocksdb-allow-mmap-writes</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to allow the OS to map a data file into memory for writes. Disabled by default.

**ROCKSDB\_ALLOW\_UNSAFE\_ALTER**

Option	Description
Command-line	<code>-rocksdb-allow-unsafe-alter</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Enables crash unsafe INPLACE ADD/DROP partition.

**ROCKSDB ALTER COLUMN DEFAULT INPLACE**

Option	Description
Command-line	<code>-rocksdb-alter-column-default-inplace</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	ON

Allows an inplace alter for the ALTER COLUMN default operation.

**ROCKSDB ALTER TABLE COMMENT INPLACE**

Option	Description
Command-line	<code>-rocksdb_alter_table_comment_inplace</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Allows changing ALTER TABLE COMMENT inplace.

This variable is disabled (OFF) by default.

**ROCKSDB\_BASE\_BACKGROUND\_COMPACTIONS**

Option	Description
Command-line	<code>-rocksdb-base-background-compactions</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	1

Specifies the suggested number of concurrent background compaction jobs, submitted to the default LOW priority thread pool in RocksDB. The default is `1`. The allowed range of values is from `-1` to `64`. The maximum value depends on the `rocksdb_max_background_compactions` variable. This variable was replaced with `rocksdb_max_background_jobs`, which automatically decides how many threads to allocate toward flush/compaction.

**ROCKSDB\_BLIND\_DELETE\_PRIMARY\_KEY**

Option	Description
Command-line	<code>--rocksdb-blind-delete-primary-key</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Skips verifying if rows exists before executing deletes. The following conditions must be met:

- The variable is enabled
- Only a single table listed in the `DELETE` statement
- The table has only a primary key with no secondary keys

**ROCKSDB\_BLOCK\_CACHE\_SIZE**

Option	Description
Command-line	<code>--rocksdb-block-cache-size</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	536870912

Specifies the size of the LRU block cache for RocksDB. This memory is reserved for the block cache, which is in addition to any filesystem caching that may occur.

Minimum value is `1024`, because that's the size of one block.

Default value is `536870912`.

Maximum value is `9223372036854775807`.

**ROCKSDB\_BULK\_LOAD\_FAIL\_IF\_NOT\_BOTTOMMOST\_LEVEL**

Option	Description
Command-line	<code>--rocksdb_bulk_load_fail_if_not_bottommost_level</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean

Option	Description
Default	OFF

When this variable is enabled, the bulk load fails if an sst file created during bulk load cannot be placed to the bottommost level in the rocksdb.

This variable can be enabled or disabled only when the `rocksdb_bulk_load` is OFF.

This variable is disabled (OFF) by default.

#### Warning

When `rocksdb_bulk_load_fail_if_not_bottommost_level` is disabled, it may cause severe performance impact.

#### ROCKSDB\_BLOCK\_RESTART\_INTERVAL

Option	Description
Command-line	<code>-rocksdb-block-restart-interval</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	16

Specifies the number of keys for each set of delta encoded data. Default value is 16. Allowed range is from 1 to 2147483647.

#### ROCKSDB\_BLOCK\_SIZE

Option	Description
Command-line	<code>-rocksdb-block-size</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	16 KB

Specifies the size of the data block for reading RocksDB data files. The default value is 16 KB. The allowed range is from 1024 to 18446744073709551615 bytes.

#### ROCKSDB\_BLOCK\_SIZE\_DEVIATION

Option	Description
Command-line	<code>-rocksdb-block-size-deviation</code>
Dynamic	No
Scope	Global
Data type	Numeric

Option	Description
Default	10

Specifies the threshold for free space allowed in a data block (see `rocksdb_block_size`). If there is less space remaining, close the block (and write to new block). Default value is `10`, meaning that the block is not closed until there is less than 10 bits of free space remaining.

Allowed range is from `1` to `2147483647`.

#### ROCKSDB\_BULK\_LOAD\_ALLOW\_SK

Option	Description
Command-line	<code>-rocksdb-bulk-load-allow-sk</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Enabling this variable allows secondary keys to be added using the bulk loading feature. This variable can be enabled or disabled only when the `rocksdb_bulk_load` is `OFF`.

#### ROCKSDB\_BULK\_LOAD\_ALLOW\_UNSORTED

Option	Description
Command-line	<code>-rocksdb-bulk-load-allow-unsorted</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

By default, the bulk loader requires its input to be sorted in the primary key order. If enabled, unsorted inputs are allowed too, which are then sorted by the bulkloader itself, at a performance penalty.

#### ROCKSDB\_BULK\_LOAD

Option	Description
Command-line	<code>-rocksdb-bulk-load</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Specifies whether to use bulk load: MyRocks will ignore checking keys for uniqueness or acquiring locks during transactions. Disabled by default. Enable this only if you are certain that there are no row conflicts, for example, when setting up a new MyRocks instance from a MySQL dump.

When the `rocksdb_bulk_load` variable is enabled, it behaves as if the variable `rocksdb_commit_in_the_middle` is enabled, even if the variable `rocksdb_commit_in_the_middle` is disabled.

**ROCKSDB\_BULK\_LOAD\_PARTIAL\_INDEX**

Option	Description
Command-line	<code>-rocksdb-bulk-load-partial-index</code>
Dynamic	Yes
Scope	Local
Data type	Boolean
Default	ON

Materializes partial index during bulk load instead of leaving the index empty.

**ROCKSDB\_BULK\_LOAD\_USE\_SST\_PARTITIONER**

Option	Description
Command-line	<code>-rocksdb_bulk_load_use_sst_parti</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

If enabled, this variable uses sst partitioner to split sst files to ensure bulk load sst files can be ingested to bottommost level.

This variable is disabled (OFF) by default.

**ROCKSDB\_BULK\_LOAD\_SIZE**

Option	Description
Command-line	<code>-rocksdb-bulk-load-size</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	1000

Specifies the number of keys to accumulate before committing them to the storage engine when bulk load is enabled (see `rocksdb_bulk_load`). Default value is `1000`, which means that a batch can contain up to 1000 records before they are implicitly committed. Allowed range is from `1` to `1073741824`.

## ROCKSDB\_BYTES\_PER\_SYNC

Option	Description
Command-line	<code>-rocksdb-bytes-per-sync</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

Specifies how often should the OS sync files to disk as they are being written, asynchronously, in the background. This operation can be used to smooth out write I/O over time. Default value is 0 meaning that files are never synced. Allowed range is up to 18446744073709551615.

## ROCKSDB\_CACHE\_DUMP

Option	Description
Command-line	<code>-rocksdb-cache-dump</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Includes RocksDB block cache content in core dump. This variable is enabled by default.

## ROCKSDB\_CACHE\_HIGH\_PRI\_POOL\_RATIO

Option	Description
Command-line	<code>-rocksdb-cache-high-pri-pool-ratio</code>
Dynamic	No
Scope	Global
Data type	Double
Default	0.0

This variable specifies the size of the block cache high-pri pool. The default value and minimum value is 0.0. The maximum value is 1.0.

## ROCKSDB\_CACHE\_INDEX\_AND\_FILTER\_BLOCKS

Option	Description
Command-line	<code>-rocksdb-cache-index-and-filter-bl</code>
Dynamic	No
Scope	Global
Data type	Boolean

Option	Description
Default	ON

Specifies whether RocksDB should use the block cache for caching the index and bloomfilter data blocks from each data file. Enabled by default. If you disable this feature, RocksDB allocates additional memory to maintain these data blocks.

#### ROCKSDB\_CACHE\_INDEX\_AND\_FILTER\_WITH\_HIGH\_PRIORITY

Option	Description
Command-line	<code>-rocksdb-cache-index-and-filter-w</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether RocksDB should use the block cache with high priority for caching the index and bloomfilter data blocks from each data file. Enabled by default. If you disable this feature, RocksDB allocates additional memory to maintain these data blocks.

#### ROCKSDB\_CANCEL\_MANUAL\_COMPACTIONS

Option	Description
Command-line	<code>-rocksdb-cancel-manual-compact</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Cancels all ongoing manual compactions.

#### ROCKSDB\_CHARGE\_MEMORY

Option	Description
Command-line	<code>-rocksdb_charge_memory</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

This variable is [tech preview](#) and may be removed in the future releases.

Turns on RocksDB memory-charging related features (BlockBasedTableOptions::cache\_usage\_options.options.charged) from `cnf` files. This variable is related to [rocksdb\\_use\\_write\\_buffer\\_manager](#).

This variable is disabled (OFF) by default.

**ROCKSDB\_CHECKSUMS\_PCT**

Option	Description
Command-line	<code>--rocksdb-checksums-pct</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	100

Specifies the percentage of rows to be checksummed. Default value is 100 (checksum all rows). Allowed range is from 0 to 100.

**ROCKSDB\_COLLECT\_SST\_PROPERTIES**

Option	Description
Command-line	<code>--rocksdb-collect-sst-properties</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether to collect statistics on each data file to improve optimizer behavior. Enabled by default.

**ROCKSDB\_COLUMN\_DEFAULT\_VALUE\_AS\_EXPRESSION**

Option	Description
Command-line	<code>--rocksdb_column_default_value_as_expression</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	ON

Allows to set a function as the default value for a column.

This variable is enabled (ON) by default.

**ROCKSDB\_COMMIT\_IN\_THE\_MIDDLE**

Option	Description
Command-line	<code>--rocksdb-commit-in-the-middle</code>
Dynamic	Yes
Scope	Global

Option	Description
Data type	Boolean
Default	OFF

Specifies whether to commit rows implicitly when a batch contains more than the value of `rocksdb_bulk_load_size`.

This option should only be enabled at the time of data import because it may cause locking errors.

This variable is disabled by default. When the `rocksdb_bulk_load` variable is enabled, it behaves as if the variable `rocksdb_commit_in_the_middle` is enabled, even if the variable `rocksdb_commit_in_the_middle` is disabled.

#### ROCKSDB\_COMMIT\_TIME\_BATCH\_FOR\_RECOVERY

Option	Description
Command-line	<code>-rocksdb-commit-time-batch-for-r</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Specifies whether to write the commit time write batch into the database or not.

#### Note

If the commit time write batch is only useful for recovery, then writing to WAL is enough.

#### ROCKSDB\_COMPACT\_CF

Option	Description
Command-line	<code>-rocksdb-compact-cf</code>
Dynamic	Yes
Scope	Global
Data type	String
Default	

Specifies the name of the column family to compact.

#### ROCKSDB\_COMPACTION\_READAHEAD\_SIZE

Option	Description
Command-line	<code>-rocksdb-compaction-readahead-s</code>
Dynamic	Yes
Scope	Global

Option	Description
Data type	Numeric
Default	0

Specifies the size of reads to perform ahead of compaction. Default value is 0. Set this to at least 2 megabytes ( 16777216 ) when using MyRocks with spinning disks to ensure sequential reads instead of random. Maximum allowed value is 18446744073709551615 .

 **Note**

If you set this variable to a non-zero value, rocksdb\_new\_table\_reader\_for\_compaction\_inputs is enabled.

#### ROCKSDB\_COMPACTION\_SEQUENTIAL\_DELETES

Option	Description
Command-line	-rocksdb-compaction-sequential-d
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

Specifies the threshold to trigger compaction on a file if it has more than this number of sequential delete markers. Default value is 0 meaning that compaction is not triggered regardless of the number of delete markers. Maximum allowed value is 2000000 (two million delete markers).

 **Note**

Depending on workload patterns, MyRocks can potentially maintain large numbers of delete markers, which increases latency of queries. This compaction feature will reduce latency, but may also increase the MyRocks write rate. Use this variable together with rocksdb\_compaction\_sequential\_deletes\_file\_size to only perform compaction on large files.

#### ROCKSDB\_COMPACTION\_SEQUENTIAL\_DELETES\_COUNT\_SD

Option	Description
Command-line	-rocksdb-compaction-sequential-d
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to count single deletes as delete markers recognized by rocksdb\_compaction\_sequential\_deletes. Disabled by default.

**ROCKSDB\_COMPACTON\_SEQUENTIAL\_DELETES\_FILE\_SIZE**

Option	Description
Command-line	<code>--rocksdb-compaction-sequential-d</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

Specifies the minimum file size required to trigger compaction on it by `rocksdb_compaction_sequential_deletes`. Default value is 0, meaning that compaction is triggered regardless of file size. Allowed range is from -1 to 9223372036854775807.

**ROCKSDB\_COMPACTON\_SEQUENTIAL\_DELETES\_WINDOW**

Option	Description
Command-line	<code>--rocksdb-compaction-sequential-d</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

Specifies the size of the window for counting delete markers by `rocksdb_compaction_sequential_deletes`. Default value is 0. Allowed range is up to 2000000 (two million).

**ROCKSDB\_CONCURRENT\_PREPARE**

Option	Description
Command-line	<code>--rocksdb-concurrent_prepare</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

When enabled this variable allows/encourages threads that are using two-phase commit to `prepare` in parallel. This variable was renamed in upstream to `rocksdb_two_write_queues`.

**ROCKSDB\_CORRUPT\_DATA\_ACTION**

Option	Description
Command-line	<code>--rocksdb_corrupt_data_action</code>
Dynamic	Yes
Scope	Global

Option	Description
Data type	enum { ERROR = 0, ABORT_SERVER, W
Default	ERROR

This variable controls the behavior when hitting the data corruption in MyRocks.

You can select one of the following actions:

- `ERROR` - fail the query with the error `HA_ERR_ROCKSDB_CORRUPT_DATA`
- `ABORT_SERVER` - crash the server
- `WARNING` - pass the query with warning

The default value is `ERROR` that means the query fails with the error `HA_ERR_ROCKSDB_CORRUPT_DATA`.

#### ROCKSDB\_CONVERTER\_RECORD\_CACHED\_LENGTH

Option	Description
Command-line	<code>-rocksdb_converter_record_cached</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

Specifies the maximum number of bytes to cache on table handler for encoding table record data.

If the used memory exceeds `rocksdb_converter_record_cached_length`, the memory is released when the handler is returned to the table handler cache.

The minimum value is `0` (zero) that means there is no limit. The maximum value is `UINT64_MAX` (`0xffffffffffffffff`).

The default value is `0` (zero) that means there is no limit.

#### ROCKSDB\_CREATE\_CHECKPOINT

Option	Description
Command-line	<code>-rocksdb-create-checkpoint</code>
Dynamic	Yes
Scope	Global
Data type	String
Default	

Specifies the directory where MyRocks should create a checkpoint. Empty by default.

**ROCKSDB\_CREATE\_IF\_MISSING**

Option	Description
Command-line	<code>-rocksdb-create-if-missing</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether MyRocks should create its database if it does not exist. Enabled by default.

**ROCKSDB\_CREATE\_MISSING\_COLUMN\_FAMILIES**

Option	Description
Command-line	<code>-rocksdb-create-missing-column-f</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether MyRocks should create new column families if they do not exist. Disabled by default.

**ROCKSDB\_CREATE\_TEMPORARY\_CHECKPOINT**

Option	Description
Command-line	<code>-rocksdb-create-temporary-checkp</code>
Dynamic	Yes
Scope	Session
Data type	String

When specified it will create a temporary RocksDB 'checkpoint' or 'snapshot' in the datadir. If the session ends with an existing checkpoint, or if the variable is reset to another value, the checkpoint will get removed. This variable should be used by backup tools. Prolonged use or other misuse can have serious side effects to the server instance.

**ROCKSDB\_DATADIR**

Option	Description
Command-line	<code>-rocksdb-datadir</code>
Dynamic	No
Scope	Global
Data type	String
Default	<code>./rocksdb</code>

Specifies the location of the MyRocks data directory. By default, it is created in the current working directory.

**ROCKSDB\_DB\_WRITE\_BUFFER\_SIZE**

Option	Description
Command-line	<code>-rocksdb-db-write-buffer-size</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	0

Specifies the maximum size of all memtables used to store writes in MyRocks across all column families. When this size is reached, the data is flushed to persistent media. The default value is 0. The allowed range is up to 18446744073709551615.

**ROCKSDB\_DEADLOCK\_DETECT**

Option	Description
Command-line	<code>-rocksdb-deadlock-detect</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Specifies whether MyRocks should detect deadlocks. Disabled by default.

**ROCKSDB\_DEADLOCK\_DETECT\_DEPTH**

Option	Description
Command-line	<code>-rocksdb-deadlock-detect-depth</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	50

Specifies the number of transactions deadlock detection will traverse through before assuming deadlock.

**ROCKSDB\_DEBUG\_CARDINALITY\_MULTIPLIER**

Option	Description
Command-line	<code>-rocksdb-debug-cardinality-multiplier</code>
Dynamic	Yes

Option	Description
Scope	Global
Data type	UINT
Default	2

The cardinality multiplier used in tests. The minimum value is 0. The maximum value is 2147483647 (INT\_MAX).

#### ROCKSDB\_DEBUG\_MANUAL\_COMPACTION\_DELAY

Option	Description
Command-line	-rocksdb-debug-manual-compaction-delay
Dynamic	Yes
Scope	Global
Data type	UINT
Default	0

Only use this variable when debugging.

This variable specifies a sleep, in seconds, to simulate long-running compactions. The minimum value is 0. The maximum value is 4292967295 (UINT\_MAX).

#### ROCKSDB\_DEBUG\_OPTIMIZER\_NO\_ZERO\_CARDINALITY

Option	Description
Command-line	-rocksdb-debug-optimizer-no-zero-cardinality
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	ON

Specifies whether MyRocks should prevent zero cardinality by always overriding it with some value.

#### ROCKSDB\_DEBUG\_TTL\_IGNORE\_PK

Option	Description
Command-line	-rocksdb-debug-ttl-ignore-pk
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

For debugging purposes only. If true, compaction filtering will not occur on Primary Key TTL data. This variable is a no-op in non-debug builds.

## ROCKSDB\_DEBUG\_TTL\_READ\_FILTER\_TS

Option	Description
Command-line	<code>-rocksdb_debug-ttl-read-filter-ts</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

For debugging purposes only. Overrides the TTL read filtering time to `time + debug_ttl_read_filter_ts`. A value of 0 denotes that the variable is not set. This variable is a no-op in non-debug builds.

## ROCKSDB\_DEBUG\_TTL\_REC\_TS

Option	Description
Command-line	<code>-rocksdb-debug-ttl-rec-ts</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

For debugging purposes only. Overrides the TTL of records to `now() + debug_ttl_rec_ts`. The value can be  $\pm$  to simulate a record inserted in the past vs a record inserted in the future . A value of 0 denotes that the variable is not set. This variable is a no-op in non-debug builds.

## ROCKSDB\_DEBUG\_TTL\_SNAPSHOT\_TS

Option	Description
Command-line	<code>-rocksdb-debug-ttl-snapshot-ts</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

For debugging purposes only. Sets the snapshot during compaction to `now() + rocksdb_debug_set_ttl_snapshot_ts`.

The value can be  $\pm$  to simulate a snapshot in the past vs a snapshot created in the future . A value of 0 denotes that the variable is not set. This variable is a no-op in non-debug builds.

## ROCKSDB\_DEFAULT\_CF\_OPTIONS

Option	Description
Command-line	<code>-rocksdb-default-cf-options</code>

Option	Description
Dynamic	No
Scope	Global
Data type	String

The default value is:

`block_based_table_factory={cache_index_and_filter_blocks=1;filter_policy=bloomfilter:10:false;whole_key_filtering=1};level_compaction_dynamic_level_bytes=true;optimize_filters_for_hits=true;compaction_concurrency=1`  
 the default column family options for MyRocks. On startup, the server applies this option to all existing column families. This option is read-only at runtime.

#### ROCKSDB\_DELAYED\_WRITE\_RATE

Option	Description
Command-line	<code>-rocksdb-delayed-write-rate</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	16777216

Specifies the write rate in bytes per second, which should be used if MyRocks hits a soft limit or threshold for writes. Default value is `16777216` (16 MB/sec). Allowed range is from `0` to `18446744073709551615`.

#### ROCKSDB\_DELETE\_CF

Option	Description
Command-line	<code>-rocksdb-delete-cf</code>
Dynamic	Yes
Scope	Global
Data type	String
Default	<code>""</code>

Deletes the column family by name. The default value is `""`, an empty string.

For example:

```
SET @@global.ROCKSDB_DELETE_CF = 'cf_primary_key';
```

#### ROCKSDB\_DELETE\_OBSOLETE\_FILES\_PERIOD\_MICROS

Option	Description
Command-line	<code>-rocksdb-delete-obsolete-files-period</code>
Dynamic	No

Option	Description
Scope	Global
Data type	Numeric
Default	21600000000

Specifies the period in microseconds to delete obsolete files regardless of files removed during compaction. Default value is 21600000000 (6 hours). Allowed range is up to 9223372036854775807.

**ROCKSDB\_DISABLE\_FILE\_DELETIONS**

Option	Description
Command-line	<code>--rocksdb-disable-file-deletions</code>
Dynamic	Yes
Scope	Session
Data type	Boolean
Default	OFF

It allows a client to temporarily disable RocksDB deletion of old `WAL` and `.sst` files for the purposes of making a consistent backup. If the client session terminates for any reason after disabling deletions and has not re-enabled deletions, they will be explicitly re-enabled. This variable should be used by backup tools. Prolonged use or other misuse can have serious side effects to the server instance.

**ROCKSDB\_DISABLE\_INSTANT\_DDL**

Option	Description
Command-line	<code>--rocksdb_disable_instant_ddl</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	ON

Disables instant DDL during `ALTER TABLE` operations.

This variable is enabled (ON) by default.

**ROCKSDB\_ENABLE\_BULK\_LOAD\_API**

Option	Description
Command-line	<code>--rocksdb-enable-bulk-load-api</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether to use the `SSTFileWriter` feature for bulk loading. This feature bypasses the memtable, but requires keys to be inserted into the table in either ascending or descending order. Enabled by default. If disabled, bulk loading uses the normal write path via the memtable and does not require keys to be inserted in any order.

**ROCKSDB\_ENABLE\_DELETE\_RANGE\_FOR\_DROP\_INDEX**

Option	Description
Command-line	<code>-rocksdb_enable_delete_range_for</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Enables drop table / index by calling the `DeleteRange`.

This option is disabled (OFF) by default.

**ROCKSDB\_ENABLE\_INSERT\_WITH\_UPDATE\_CACHING**

Option	Description
Command-line	<code>-rocksdb-enable-insert-with-updat</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	ON

Specifies whether to enable optimization where the read is cached from a failed insertion attempt in `INSERT ON DUPLICATE KEY UPDATE`.

**ROCKSDB\_ENABLE\_ITERATE\_BOUNDS**

Option	Description
Command-line	<code>-rocksdb-enable-iterate-bounds</code>
Dynamic	Yes
Scope	Global, Local
Data type	Boolean
Default	TRUE

Enables the rocksdb iterator upper bounds and lower bounds in read options.

**ROCKSDB\_ENABLE\_PIPELINED\_WRITE**

Option	Description
Command-line	<code>-rocksdb-enable-pipelined-write</code>

Option	Description
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

`DBOptions::enable_pipelined_write` for RocksDB.

If `enable_pipelined_write` is `ON`, a separate write thread is maintained for WAL write and memtable write. A write thread first enters the WAL writer queue and then the memtable writer queue. A pending thread on the WAL writer queue only waits for the previous WAL write operations but does not wait for memtable write operations. Enabling the feature may improve write throughput and reduce latency of the prepare phase of a two-phase commit.

#### `ROCKSDB_ENABLE_REMOVE_ORPHANED_DROPPED_CFS`

Option	Description
Command-line	<code>-rocksdb-enable-remove-orphaned</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	TRUE

Enables the removal of dropped column families (cfs) from metadata if the cfs do not exist in the cf manager.

The default value is `TRUE`.

#### `ROCKSDB_ENABLE_TTL`

Option	Description
Command-line	<code>-rocksdb-enable-ttl</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether to keep expired TTL records during compaction. Enabled by default. If disabled, expired TTL records will be dropped during compaction.

#### `ROCKSDB_ENABLE_TTL_READ_FILTERING`

Option	Description
Command-line	<code>-rocksdb-enable-ttl-read-filtering</code>
Dynamic	Yes

Option	Description
Scope	Global
Data type	Boolean
Default	ON

For tables with TTL, expired records are skipped/filtered out during processing and in query results. Disabling this will allow these records to be seen, but as a result rows may disappear in the middle of transactions as they are dropped during compaction. **Use with caution.**

**ROCKSDB\_ENABLE\_THREAD\_TRACKING**

Option	Description
Command-line	<code>--rocksdb-enable-thread-tracking</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to enable tracking the status of threads accessing the database. Disabled by default. If enabled, thread status will be available via `GetThreadList()`.

**ROCKSDB\_ENABLE\_WRITE\_THREAD\_ADAPTIVE\_YIELD**

Option	Description
Command-line	<code>--rocksdb-enable-write-thread-adaptive-yield</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether the MyRocks write batch group leader should wait up to the maximum allowed time before blocking on a mutex. Disabled by default. Enable it to increase throughput for concurrent workloads.

**ROCKSDB\_ERROR\_IF\_EXISTS**

Option	Description
Command-line	<code>--rocksdb-error-if-exists</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to report an error when a database already exists. Disabled by default.

**ROCKSDB\_ERROR\_ON\_SUBOPTIMAL\_COLLATION**

Option	Description
Command-line	<code>-rocksdb-error-on-suboptimal-coll</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether to report an error instead of a warning if an index is created on a char field where the table has a sub-optimal collation (case insensitive). Enabled by default.

**ROCKSDB\_FLUSH\_LOG\_AT\_TRX\_COMMIT**

Option	Description
Command-line	<code>-rocksdb-flush-log-at-trx-commit</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	1

Specifies whether to sync on every transaction commit, similar to `innodb_flush_log_at_trx_commit`. Enabled by default, which ensures ACID compliance.

Possible values:

- 0: Do not sync on transaction commit. This provides better performance, but may lead to data inconsistency in case of a crash.
- 1: Sync on every transaction commit. This is set by default and recommended as it ensures data consistency, but reduces performance.
- 2: Sync every second.

**ROCKSDB\_FLUSH\_MEMTABLE\_ON\_ANALYZE**

Option	Description
Command-line	<code>-rocksdb-flush-memtable-on-analy</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	ON

Specifies whether to flush the memtable when running `ANALYZE` on a table. Enabled by default. This ensures accurate cardinality by including data in the memtable for calculating stats.

**ROCKSDB\_FORCE\_COMPUTE\_MEMTABLE\_STATS**

Option	Description
Command-line	-rocksdb-force-compute-memtable
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	ON

Specifies whether data in the memtables should be included for calculating index statistics used by the query optimizer. Enabled by default. This provides better accuracy, but may reduce performance.

**ROCKSDB\_FORCE\_COMPUTE\_MEMTABLE\_STATS\_CACHETIME**

Option	Description
Command-line	-rocksdb-force-compute-memtable
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	60000000

Specifies for how long the cached value of memtable statistics should be used instead of computing it every time during the query plan analysis.

**ROCKSDB\_FORCE\_FLUSH\_MEMTABLE\_AND\_LZERO\_NOW**

Option	Description
Command-line	-rocksdb-force-flush-memtable-and
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Works similar to rocksdb\_force\_flush\_memtable\_now but also flushes all L0 files.

**ROCKSDB\_FORCE\_FLUSH\_MEMTABLE\_NOW**

Option	Description
Command-line	-rocksdb-force-flush-memtable-no
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Forces MyRocks to immediately flush all memtables out to data files.

**Warning**

Use with caution! Write requests will be blocked until all memtables are flushed.

**ROCKSDB\_FORCE\_INDEX\_RECORDS\_IN\_RANGE**

Option	Description
Command-line	<code>-rocksdb-force-index-records-in-range</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	1

Specifies the value used to override the number of rows returned to query optimizer when `FORCE_INDEX` is used. Default value is `1`. Allowed range is from `0` to `2147483647`. Set to `0` if you do not want to override the returned value.

**ROCKSDB\_HASH\_INDEX\_ALLOW\_COLLISION**

Option	Description
Command-line	<code>-rocksdb-hash-index-allow-collision</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether hash collisions are allowed. Enabled by default, which uses less memory. If disabled, full prefix is stored to prevent hash collisions.

**ROCKSDB\_IGNORE\_UNKNOWN\_OPTIONS**

Option	Description
Command-line	
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

When enabled, it allows RocksDB to receive unknown options and not exit.

**ROCKSDB\_INDEX\_TYPE**

Option	Description
Command-line	<code>--rocksdb-index-type</code>
Dynamic	No
Scope	Global
Data type	Enum
Default	<code>kBinarySearch</code>

Specifies the type of indexing used by MyRocks:

- `kBinarySearch`: Binary search (default).
- `kHashSearch`: Hash search.

**ROCKSDB\_INFO\_LOG\_LEVEL**

Option	Description
Command-line	<code>--rocksdb-info-log-level</code>
Dynamic	Yes
Scope	Global
Data type	Enum
Default	<code>error_level</code>

Specifies the level for filtering messages written by MyRocks to the `mysqld` log.

- `debug_level`: Maximum logging (everything including debugging log messages)
- `info_level`
- `warn_level`
- `error_level` (default)
- `fatal_level`: Minimum logging (only fatal error messages logged)

**ROCKSDB\_IS\_FD\_CLOSE\_ON\_EXEC**

Option	Description
Command-line	<code>--rocksdb-is-fd-close-on-exec</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether child processes should inherit open file handles. Enabled by default.

## ROCKSDB\_LARGE\_PREFIX

Option	Description
Command-line	<code>-rocksdb-large-prefix</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	TRUE

When enabled, this option allows index key prefixes longer than 767 bytes (up to 3072 bytes). The values for `rocksdb_large_prefix` should be the same between source and replica.

## ROCKSDB\_KEEP\_LOG\_FILE\_NUM

Option	Description
Command-line	<code>-rocksdb-keep-log-file-num</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	1000

Specifies the maximum number of info log files to keep. Default value is `1000`. Allowed range is from `1` to `18446744073709551615`.

## ROCKSDB\_LOCK\_SCANNED\_ROWS

Option	Description
Command-line	<code>-rocksdb-lock-scanned-rows</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Specifies whether to hold the lock on rows that are scanned during `UPDATE` and not actually updated. Disabled by default.

## ROCKSDB\_LOCK\_WAIT\_TIMEOUT

Option	Description
Command-line	<code>-rocksdb-lock-wait-timeout</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric

Option	Description
Default	1

Specifies the number of seconds MyRocks should wait to acquire a row lock before aborting the request. Default value is 1. Allowed range is up to 1073741824.

## ROCKSDB\_LOG\_FILE\_TIME\_TO\_ROLL

Option	Description
Command-line	-rocksdb-log-file-time-to-roll
Dynamic	No
Scope	Global
Data type	Numeric
Default	0

Specifies the period (in seconds) for rotating the info log files. Default value is 0, meaning that the log file is not rotated. Allowed range is up to 18446744073709551615.

## ROCKSDB\_MANIFEST\_PREALLOCATION\_SIZE

Option	Description
Command-line	-rocksdb-manifest-preallocation-size
Dynamic	No
Scope	Global
Data type	Numeric
Default	0

Specifies the number of bytes to preallocate for the MANIFEST file used by MyRocks to store information about column families, levels, active files, etc. Default value is 0. Allowed range is up to 18446744073709551615.

 **Note**

A value of 4194304 (4 MB) is reasonable to reduce random I/O on XFS.

## ROCKSDB\_MANUAL\_COMPACTION\_BOTTOMMOST\_LEVEL

Option	Description
Command-line	-rocksdb-manual-compaction-bottommost-level
Dynamic	Yes
Scope	Local
Data type	Enum
Default	kForceOptimized

Option for bottommost level compaction during manual compaction:

- `kSkip` – Skip bottommost level compaction
- `kIfHaveCompactionFilter` – Only compact bottommost level if there is a compaction filter
- `kForce` – Always compact bottommost level
- `kForceOptimized` – Always compact bottommost level but in bottommost level avoid double-compacting files created in the same compaction

#### ROCKSDB\_MANUAL\_COMPACTON\_THREADS

Option	Description
Command-line	<code>-rocksdb-manual-compaction-thre</code>
Dynamic	Yes
Scope	Local
Data type	INT
Default	0

The variable defines the number of RocksDB threads to run for a manual compaction. The minimum value is 0. The maximum value is 120.

#### ROCKSDB\_MANUAL\_WAL\_FLUSH

Option	Description
Command-line	<code>-rocksdb-manual-wal-flush</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

This variable can be used to disable automatic/timed WAL flushing and instead rely on the application to do the flushing.

#### ROCKSDB\_MASTER\_SKIP\_TX\_API

Option	Description
Command-line	
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

When enabled, uses the WriteBatch API, which is faster. The session does not hold any lock on row access. This variable is not effective on replica.

 **Note**

Due to the disabled row locks, improper use of the variable can cause data corruption or inconsistency.

**ROCKSDB\_MAX\_BACKGROUND\_COMPACTIONS**

Option	Description
Command-line	<code>--rocksdb-max-background-compa</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	-1

Sets `DBOptions::max_background_compactions` for RocksDB. The default value is `-1`. The allowed range is `-1` to `64`. This variable was replaced by `rocksdb_max_background_jobs`, which automatically decides how many threads to allocate towards flush/compaction.

**ROCKSDB\_MAX\_BACKGROUND\_FLUSHES**

Option	Description
Command-line	<code>--rocksdb-max-background-flushes</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	-1

Sets `DBOptions::max_background_flushes` for RocksDB. The default value is `-1`. The allowed range is `-1` to `64`. This variable has been replaced by `rocksdb_max_background_jobs`, which automatically decides how many threads to allocate towards flush/compaction.

**ROCKSDB\_MAX\_BACKGROUND\_JOBS**

Option	Description
Command-line	<code>--rocksdb-max-background-jobs</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	2

This variable replaced `rocksdb_base_background_compactions`, `rocksdb_max_background_compactions`, and `rocksdb_max_background_flushes` variables. This variable specifies the maximum number of background jobs. It automatically decides how many threads to allocate towards flush/compaction. It was implemented to reduce the number of (confusing) options users and can tweak and push the responsibility down to RocksDB level.

## ROCKSDB\_MAX\_BOTTOM\_PRI\_BACKGROUND\_COMPACTIONS

Option	Description
Command-line	<code>-rocksdb_max_bottom_pri_backgr</code>
Dynamic	No
Data type	Unsigned integer
Default	0

Creates a specified number of threads, sets a lower CPU priority, and letting compactions use them. The maximum compaction concurrency is capped by `rocksdb_max_background_compactions` or `rocksdb_max_background_jobs`

The minimum value is 0 and the maximum value is 64.

## ROCKSDB\_MAX\_COMPACTON\_HISTORY

Option	Description
Command-line	<code>-rocksdb-max-compaction-history</code>
Dynamic	Yes
Scope	Global
Data type	Unsigned integer
Default	64

The minimum value is 0 and the maximum value is `UINT64_MAX`.

Tracks the history for at most `rockdb_mx_compaction_history` completed compactions. The history is in the `INFORMATION_SCHEMA.ROCKSDB_COMPACTON_HISTORY` table.

## ROCKSDB\_MAX\_LATEST\_DEADLOCKS

Option	Description
Command-line	<code>-rocksdb-max-latest-deadlocks</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	5

Specifies the maximum number of recent deadlocks to store.

## ROCKSDB\_MAX\_LOG\_FILE\_SIZE

Option	Description
Command-line	<code>-rocksdb-max-log-file-size</code>
Dynamic	No
Scope	Global

Option	Description
Data type	Numeric
Default	0

Specifies the maximum size for info log files, after which the log is rotated. Default value is 0, meaning that only one log file is used. Allowed range is up to 18446744073709551615 .

Also see `rocksdb_log_file_time_to_roll`.

#### ROCKSDB\_MAX\_MANIFEST\_FILE\_SIZE

Option	Description
Command-line	<code>-rocksdb-manifest-log-file-size</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	18446744073709551615

Specifies the maximum size of the MANIFEST data file, after which it is rotated. Default value is also the maximum, making it practically unlimited: only one manifest file is used.

#### ROCKSDB\_MAX\_MANUAL\_COMPACTIONS

Option	Description
Command-line	<code>-rocksdb-max-manual-compaction</code>
Dynamic	Yes
Scope	Global
Data type	UINT
Default	10

The variable defines the maximum number of pending plus ongoing manual compactions. The default value and the minimum value is 0. The maximum value is 4294967295 (`UNIT_MAX`).

#### ROCKSDB\_MAX\_OPEN\_FILES

Option	Description
Command-line	<code>-rocksdb-max-open-files</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	1000

Specifies the maximum number of file handles opened by MyRocks. Values in the range between 0 and `open_files_limit` are taken as they are. If `rocksdb_max_open_files` value is greater than

`open_files_limit`, it will be reset to  $\frac{1}{2}$  of `open_files_limit`, and a warning will be emitted to the `mysqld` error log. A value of `-2` denotes auto tuning: just sets `rocksdb_max_open_files` value to  $\frac{1}{2}$  of `open_files_limit`. Finally, `-1` means no limit, i.e. an infinite number of file handles.

#### Warning

Setting `rocksdb_max_open_files` to `-1` is dangerous, as the server may quickly run out of file handles in this case.

#### ROCKSDB\_MAX\_ROW\_LOCKS

Option	Description
Command-line	<code>--rocksdb-max-row-locks</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	1048576

Specifies the limit on the maximum number of row locks a transaction can have before it fails. Default value is also the maximum, making it practically unlimited: transactions never fail due to row locks.

#### ROCKSDB\_MAX\_SUBCOMPACTIONS

Option	Description
Command-line	<code>--rocksdb-max-subcompactions</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	1

Specifies the maximum number of threads allowed for each compaction job. Default value of `1` means no subcompactions (one thread per compaction job). Allowed range is up to `64`.

#### ROCKSDB\_MAX\_TOTAL\_WAL\_SIZE

Option	Description
Command-line	<code>--rocksdb-max-total-wal-size</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	2 GB

Specifies the maximum total size of WAL (write-ahead log) files, after which memtables are flushed. Default value is `2 GB`. The allowed range is up to `9223372036854775807`.

**ROCKSDB\_MERGE\_BUF\_SIZE**

Option	Description
Command-line	<code>-rocksdb-merge-buf-size</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	67108864

Specifies the size (in bytes) of the merge-sort buffers used to accumulate data during secondary key creation. New entries are written directly to the lowest level in the database, instead of updating indexes through the memtable and L0. These values are sorted using merge-sort, with buffers set to 64 MB by default ( 67108864 ). Allowed range is from 100 to 18446744073709551615 .

**ROCKSDB\_MERGE\_COMBINE\_READ\_SIZE**

Option	Description
Command-line	<code>-rocksdb-merge-combine-read-size</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	1073741824

Specifies the size (in bytes) of the merge-combine buffer used for the merge-sort algorithm as described in `rocksdb_merge_buf_size`. Default size is 1 GB ( 1073741824 ). Allowed range is from 100 to 18446744073709551615 .

**ROCKSDB\_MERGE\_TMP\_FILE\_REMOVAL\_DELAY\_MS**

Option	Description
Command-line	<code>-rocksdb_merge_tmp_file_removal</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	0

Fast secondary index creation creates merge files when needed. After finishing secondary index creation, merge files are removed. By default, the file removal is done without any sleep, so removing GBs of merge files within <1s may happen, which will cause trim stalls on Flash. This variable can be used to rate limit the delay in milliseconds.

**ROCKSDB\_NEW\_TABLE\_READER\_FOR\_COMPACTION\_INPUTS**

Option	Description
Command-line	<code>-rocksdb-new-table-reader-for-comp</code>

Option	Description
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether MyRocks should create a new file descriptor and table reader for each compaction input. Disabled by default. Enabling this may increase memory consumption, but will also allow pre-fetch options to be specified for compaction input files without impacting table readers used for user queries.

**ROCKSDB\_NO\_BLOCK\_CACHE**

Option	Description
Command-line	<code>-rocksdb-no-block-cache</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to disable the block cache for column families. Variable is disabled by default, meaning that using the block cache is allowed.

**ROCKSDB\_NO\_CREATE\_COLUMN\_FAMILY**

Option	Description
Command-line	<code>-rocksdb-no-create-column-family</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Controls the processing of the column family name given in the `COMMENT` clause in the `CREATE TABLE` or `ALTER TABLE` statement in case the column family name does not refer to an existing column family.

If `rocksdb_no_create_column_family` is set to NO, a new column family will be created and the new index will be placed into it.

If `rocksdb_no_create_column_family` is set to YES, no new column family will be created and the index will be placed into the default column family. A warning is issued in this case informing that the specified column family does not exist and cannot be created.

## ROCKSDB\_OVERRIDE\_CF\_OPTIONS

Option	Description
Command-line	<code>-rocksdb-override-cf-options</code>
Dynamic	No
Scope	Global
Data type	String
Default	

Specifies option overrides for each column family. Empty by default.

## ROCKSDB\_PARANOID\_CHECKS

Option	Description
Command-line	<code>-rocksdb-paranoid-checks</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether MyRocks should re-read the data file as soon as it is created to verify correctness. Enabled by default.

## ROCKSDB\_PARTIAL\_INDEX\_SORT\_MAX\_MEM

Option	Description
Command-line	<code>-rocksdb-partial-index-sort-max-m</code>
Dynamic	Yes
Scope	Local
Data type	Unsigned Integer
Default	0

Maximum memory to use when sorting an unmaterialized group for partial indexes. The 0(zero) value is defined as no limit.

## ROCKSDB\_PAUSE\_BACKGROUND\_WORK

Option	Description
Command-line	<code>-rocksdb-pause-background-work</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether MyRocks should pause all background operations. Disabled by default. There is no practical reason for a user to ever use this variable because it is intended as a test synchronization tool for the MyRocks MTR test suites.

#### Warning

If someone were to set a `rocksdb_force_flush_memtable_now` to `1` while `rocksdb_pause_background_work` is set to `1`, the client that issued the `rocksdb_force_flush_memtable_now=1` will be blocked indefinitely until `rocksdb_pause_background_work` is set to `0`.

#### ROCKSDB\_PARTIAL\_INDEX\_BLIND\_DELETE

Option	Description
Command-line	<code>-rocksdb_partial_index_blind_delete</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	ON

If enabled, the server does not read from the partial index to check if the key exists before deleting the partial index and the delete marker is unconditionally written.

If the variable is disabled (OFF), the server always reads from partial index to check if key exists before deleting the partial index.

This variable is enabled (ON) by default.

#### ROCKSDB\_PERF\_CONTEXT\_LEVEL

Option	Description
Command-line	<code>-rocksdb-perf-context-level</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	0

Specifies the level of information to capture with the Perf Context plugins. The default value is `0`. The allowed range is up to `5`.

Value	Description
1	Disable perf stats
2	Enable only count stats
3	Enable count stats and time stats ex
4	Enable count stats and time stats, ex mutexes

Value	Description
5	Enable all count stats and time stats

**ROCKSDB\_PERSISTENT\_CACHE\_PATH**

Option	Description
Command-line	<code>-rocksdb-persistent-cache-path</code>
Dynamic	No
Scope	Global
Data type	String
Default	

Specifies the path to the persistent cache. Set this together with `rocksdb_persistent_cache_size_mb`.

**ROCKSDB\_PERSISTENT\_CACHE\_SIZE\_MB**

Option	Description
Command-line	<code>-rocksdb-persistent-cache-size-mb</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	0

Specifies the size of the persistent cache in megabytes. Default is 0 (persistent cache disabled). Allowed range is up to 18446744073709551615. Set this together with `rocksdb_persistent_cache_path`.

**ROCKSDB\_PIN\_LO\_FILTER\_AND\_INDEX\_BLOCKS\_IN\_CACHE**

Option	Description
Command-line	<code>-rocksdb-pin-l0-filter-and-index-bl</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether MyRocks pins the filter and index blocks in the cache if `rocksdb_cache_index_and_filter_blocks` is enabled. Enabled by default.

**ROCKSDB\_PRINT\_SNAPSHOT\_CONFLICT\_QUERIES**

Option	Description
Command-line	<code>-rocksdb-print-snapshot-conflict-q</code>
Dynamic	Yes

Option	Description
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether queries that generate snapshot conflicts should be logged to the error log. Disabled by default.

**ROCKSDB\_PROTECTION\_BYTES\_PER\_KEY**

Option	Description
Command-line	<code>-rocksdb_protection_bytes_per_key</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	0

This variable is used to configure `WriteOptions::protection_bytes_per_key`. The default value is 0 (disabled). When this variable is set to 1, 2, 4, or 8, it uses that number of bytes per key value to protect entries in the WriteBatch.

The minimum value is 0.

The maximum value is `ULONG_MAX` (0xFFFFFFFF).

**ROCKSDB\_RATE\_LIMITER\_BYTES\_PER\_SEC**

Option	Description
Command-line	<code>-rocksdb-rate-limiter-bytes-per-sec</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

Specifies the maximum rate at which MyRocks can write to media via memtable flushes and compaction. Default value is 0 (write rate is not limited). Allowed range is up to 9223372036854775807.

**ROCKSDB\_READ\_FREE\_RPL**

Option	Description
Command-line	<code>-rocksdb-read-free-rpl</code>
Dynamic	Yes
Scope	Global
Data type	Enum

Option	Description
Default	OFF

Uses read-free replication, which allows no row lookup during replication, on the replica.

The options are the following:

- OFF - Disables the variable
- PK\_SK - Enables the variable on all tables with a primary key
- PK\_ONLY - Enables the variable on tables where the only key is the primary key

#### ROCKSDB\_READ\_FREE\_RPL\_TABLES

Option	Description
Command-line	<code>--rocksdb-read-free-rpl-tables</code>
Dynamic	Yes
Scope	Global, Session
Data type	String
Default	

We recommend that you use `rocksdb_read_free_rpl` instead of this variable.

This variable lists tables (as a regular expression) that should use read-free replication on the replica (that is, replication without row lookups). Empty by default.

#### ROCKSDB\_RECORDS\_IN\_RANGE

Option	Description
Command-line	<code>--rocksdb-records-in-range</code>
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	0

Specifies the value to override the result of `records_in_range()`. Default value is 0. Allowed range is up to 2147483647.

#### ROCKSDB\_RESET\_STATS

Option	Description
Command-line	<code>--rocksdb-reset-stats</code>
Dynamic	Yes
Scope	Global
Data type	Boolean

Option	Description
Default	OFF

Resets MyRocks internal statistics dynamically (without restarting the server).

#### ROCKSDB\_ROLLBACK\_ON\_TIMEOUT

Option	Description
Command-line	<code>--rocksdb-rollback-on-timeout</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

By default, only the last statement on a transaction is rolled back. If `--rocksdb-rollback-on-timeout=ON`, a transaction timeout causes a rollback of the entire transaction.

#### ROCKSDB\_RPL\_SKIP\_TX\_API

Option	Description
Command-line	<code>--rocksdb-rpl-skip-tx-api</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether write batches should be used for replication thread instead of the transaction API. Disabled by default.

There are two conditions which are necessary to use it: row replication format and replica operating in super read only mode.

#### ROCKSDB\_SECONDS\_BETWEEN\_STAT\_COMPUTES

Option	Description
Command-line	<code>--rocksdb-seconds-between-stat-c</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	3600

Specifies the number of seconds to wait between recomputation of table statistics for the optimizer. During that time, only changed indexes are updated. Default value is 3600. Allowed is from 0 to 4294967295.

**ROCKSDB\_SIGNAL\_DROP\_INDEX\_THREAD**

Option	Description
Command-line	<code>-rocksdb-signal-drop-index-thread</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Signals the MyRocks drop index thread to wake up.

**ROCKSDB\_SIM\_CACHE\_SIZE**

Option	Description
Command-line	<code>-rocksdb-sim-cache-size</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	0

Enables the simulated cache, which allows us to figure out the hit/miss rate with a specific cache size without changing the real block cache.

**ROCKSDB\_SKIP\_BLOOM\_FILTER\_ON\_READ**

Option	Description
Command-line	<code>-rocksdb-skip-bloom-filter-on_read</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Specifies whether bloom filters should be skipped on reads. Disabled by default (bloom filters are not skipped).

**ROCKSDB\_SKIP\_FILL\_CACHE**

Option	Description
Command-line	<code>-rocksdb-skip-fill-cache</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Specifies whether to skip caching data on read requests. Disabled by default (caching is not skipped).

**ROCKSDB\_SKIP\_LOCKS\_IF\_SKIP\_UNIQUE\_CHECK**

Option	Description
Command-line	rocksdb_skip_locks_if_skip_unique
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Skip row locking when unique checks are disabled.

**ROCKSDB\_SST\_MGR\_RATE\_BYTES\_PER\_SEC**

Option	Description
Command-line	-rocksdb-sst-mgr-rate-bytes-per-s
Dynamic	Yes
Scope	Global, Session
Data type	Numeric
Default	0

Specifies the maximum rate for writing to data files. Default value is 0. This option is not effective on HDD. Allowed range is from 0 to 18446744073709551615.

**ROCKSDB\_STATS\_DUMP\_PERIOD\_SEC**

Option	Description
Command-line	-rocksdb-stats-dump-period-sec
Dynamic	No
Scope	Global
Data type	Numeric
Default	600

Specifies the period in seconds for performing a dump of the MyRocks statistics to the info log. Default value is 600. Allowed range is up to 2147483647.

**ROCKSDB\_STATS\_LEVEL**

Option	Description
Command-line	-rocksdb-stats-level
Dynamic	Yes
Scope	Global

Option	Description
Data type	Numeric
Default	0

Controls the RocksDB statistics level. The default value is "0" (kExceptHistogramOrTimers), which is the fastest level. The maximum value is "4".

## ROCKSDB\_STATS\_RECALC\_RATE

Option	Description
Command-line	-rocksdb-stats-recalc-rate
Dynamic	No
Scope	Global
Data type	Numeric
Default	0

Specifies the number of indexes to recalculate per second. Recalculating index statistics periodically ensures it to match the actual sum from SST files. Default value is 0. Allowed range is up to 4294967295.

## ROCKSDB\_STORE\_ROW\_DEBUG\_CHECKSUMS

Option	Description
Command-line	-rocksdb-store-row-debug-checksums
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to include checksums when writing index or table records. Disabled by default.

## ROCKSDB\_STRICT\_COLLATION\_CHECK

Option	Description
Command-line	-rocksdb-strict-collation-check
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	ON

This variable is considered **deprecated** in version 8.0.23-14.

Specifies whether to check and verify that table indexes have proper collation settings. Enabled by default.

## ROCKSDB\_STRICT\_COLLATION\_EXCEPTIONS

Option	Description
Command-line	<code>-rocksdb-strict-collation-exceptions</code>
Dynamic	Yes
Scope	Global
Data type	String
Default	

This variable is considered **deprecated** in version 8.0.23-14.

Lists tables (as a regular expression) that should be excluded from verifying case-sensitive collation enforced by `rocksdb_strict_collation_check`. Empty by default.

## ROCKSDB\_TABLE\_CACHE\_NUMSHARDBITS

Option	Description
Command-line	<code>-rocksdb-table-cache-numshardbits</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	6

Specifies the number of table caches. The default value is 6. The allowed range is from 0 to 19.

## ROCKSDB\_TABLE\_STATS\_BACKGROUND\_THREAD\_NICE\_VALUE

Option	Description
Command-line	<code>-rocksdb-table-stats-background-t</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	19

The nice value for index stats. The minimum = -20 (THREAD\_PRIO\_MIN) The maximum = 19 (THREAD\_PRIO\_MAX)

## ROCKSDB\_TABLE\_STATS\_MAX\_NUM\_ROWS\_SCANNED

Option	Description
Command-line	<code>-rocksdb-table-stats-max-num-rows</code>
Dynamic	Yes
Scope	Global
Data type	Numeric

Option	Description
Default	0

The maximum number of rows to scan in a table scan based on a cardinality calculation. The minimum is 0 (every modification triggers a stats recalculation). The maximum is 18,446,744,073,709,551,615.

## ROCKSDB\_TABLE\_STATS\_RECALC\_THRESHOLD\_COUNT

Option	Description
Command-line	-rocksdb-table-stats-recalc-thresh
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	100

The number of modified rows to trigger a stats recalculation. This is a dependent variable for stats recalculation. The minimum is 0. The maximum is 18,446,744,073,709,551,615.

## ROCKSDB\_TABLE\_STATS\_RECALC\_THRESHOLD\_PCT

Option	Description
Command-line	-rocksdb-table-stats-recalc-thresh
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	10

The percentage of the number of modified rows over the total number of rows to trigger stats recalculations. This is a dependent variable for stats recalculation. The minimum value is 0. The maximum value is 100 (RDB\_TBL\_STATS\_RECALC\_THRESHOLD\_PCT\_MAX).

## ROCKSDB\_TABLE\_STATS\_SAMPLING\_PCT

Option	Description
Command-line	-rocksdb-table-stats-sampling-pct
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	10

Specifies the percentage of entries to sample when collecting statistics about table properties. Default value is 10. Allowed range is from 0 to 100.

**ROCKSDB\_TABLE\_STATS\_USE\_TABLE\_SCAN**

Option	Description
Command-line	<code>-rocksdb-table-stats-use-table-scan</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	FALSE

Enables table-scan-based index calculations. The default value is `FALSE`.

**ROCKSDB\_TMPDIR**

Option	Description
Command-line	<code>-rocksdb-tmpdir</code>
Dynamic	Yes
Scope	Global, Session
Data type	String
Default	

Specifies the path to the directory for temporary files during DDL operations.

**ROCKSDB\_TRACE\_BLOCK\_CACHE\_ACCESS**

Option	Description
Command-line	<code>-rocksdb-trace-block-cache-access</code>
Dynamic	Yes
Scope	Global
Data type	String
Default	""

Defines the block cache trace option string. The format is `sampling frequency: max_trace_file_size:trace_file_name`. The sampling frequency value and `max_trace_file_size` value are positive integers. The block accesses are saved to the `rocksdb_datadir/block_cache_traces/trace_file_name`. The default value is an empty string.

**ROCKSDB\_TRACE\_QUERIES**

Option	Description
Command-line	<code>-rocksdb-trace-queries</code>
Dynamic	Yes
Scope	Global
Data type	String

Option	Description
Default	""

This variable is a trace option string. The format is `sampling_frequency:max_trace_file_size:trace_file_name`. The `sampling_frequency` and `max_trace_file_size` are positive integers. The queries are saved to the `rocksdb_datadir/queries_traces/trace_file_name`.

**ROCKSDB\_TRACE\_SST\_API**

Option	Description
Command-line	<code>-rocksdb-trace-sst-api</code>
Dynamic	Yes
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to generate trace output in the log for each call to `SstFileWriter`. Disabled by default.

**ROCKSDB\_TRACK\_AND\_VERIFY\_WALS\_IN\_MANIFEST**

Option	Description
Command-line	<code>-rocksdb-track-and-verify-wals-in-</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

`DBOptions::track_and_verify_wals_in_manifest` for RocksDB.

**ROCKSDB\_TWO\_WRITE\_QUEUES**

Option	Description
Command-line	<code>-rocksdb-track-and-verify-wals-in-</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

When enabled this variable allows/encourages threads that are using two-phase commit to `prepare` in parallel.

**ROCKSDB\_UNSAFE\_FOR\_BINLOG**

Option	Description
Command-line	<code>-rocksdb-unsafe-for-binlog</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Specifies whether to allow statement-based binary logging which may break consistency. Disabled by default.

**ROCKSDB\_UPDATE\_CF\_OPTIONS**

Option	Description
Command-line	<code>-rocksdb-update-cf-options</code>
Dynamic	No
Scope	Global
Data type	String
Default	

Specifies option updates for each column family. Empty by default.

**ROCKSDB\_USE\_ADAPTIVE\_MUTEX**

Option	Description
Command-line	<code>-rocksdb-use-adaptive-mutex</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to use adaptive mutex which spins in user space before resorting to the kernel. Disabled by default.

**ROCKSDB\_USE\_DEFAULT\_SK\_CF**

Option	Description
Command-line	<code>-rocksdb-use-default-sk-cf</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Use `default_sk` column family for secondary keys.

**ROCKSDB\_USE\_DIRECT\_IO\_FOR\_FLUSH\_AND\_COMPACTION**

Option	Description
Command-line	<code>-rocksdb-use-direct-io-for-flush-an</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to write to data files directly, without caches or buffers. Disabled by default.

**ROCKSDB\_USE\_DIRECT\_READS**

Option	Description
Command-line	<code>-rocksdb-use-direct-reads</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether to read data files directly, without caches or buffers. Disabled by default. If you enable this, make sure that `rocksdb_allow_mmap_reads` is disabled.

**ROCKSDB\_USE\_FSYNC**

Option	Description
Command-line	<code>-rocksdb-use-fsync</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

Specifies whether MyRocks should use `fsync` instead of `fdatasync` when requesting a sync of a data file. Disabled by default.

**ROCKSDB\_USE\_HYPER\_CLOCK\_CACHE**

Option	Description
Command-line	<code>-rocksdb_use_hyper_clock_cache</code>
Dynamic	No
Scope	Global
Data type	Boolean

Option	Description
Default	OFF

If enabled, this variable uses HyperClockCache instead of default LRUCache for RocksDB.

This variable is disabled (OFF) by default.

#### ROCKSDB\_USE\_WRITE\_BUFFER\_MANAGER

Option	Description
Command-line	<code>--rocksdb_use_write_buffer_manag</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	OFF

This variable is [tech preview](#) and may be removed in the future releases.

Allows to turn on the write buffer manager (WriteBufferManager) from `cnf` files. This variable is related to [rocksdb\\_charge\\_memory](#).

#### ROCKSDB\_VALIDATE\_TABLES

Option	Description
Command-line	<code>--rocksdb-validate-tables</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	1

Specifies whether to verify that MySQL data dictionary is equal to the MyRocks data dictionary.

- 0 : do not verify.
- 1 : verify and fail on error (default).
- 2 : verify and continue with error.

#### ROCKSDB\_VERIFY\_ROW\_DEBUG\_CHECKSUMS

Option	Description
Command-line	<code>--rocksdb-verify-row-debug-checks</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Specifies whether to verify checksums when reading index or table records. Disabled by default.

**ROCKSDB\_WAL\_BYTES\_PER\_SYNC**

Option	Description
Command-line	<code>-rocksdb-wal-bytes-per-sync</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

Specifies how often should the OS sync WAL (write-ahead log) files to disk as they are being written, asynchronously, in the background. This operation can be used to smooth out write I/O over time. Default value is 0, meaning that files are never synced. Allowed range is up to 18446744073709551615.

**ROCKSDB\_WAL\_DIR**

Option	Description
Command-line	<code>-rocksdb-wal-dir</code>
Dynamic	No
Scope	Global
Data type	String
Default	

Specifies the path to the directory where MyRocks stores WAL files.

**ROCKSDB\_WAL\_RECOVERY\_MODE**

Option	Description
Command-line	<code>-rocksdb-wal-recovery-mode</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	2

Specifies the level of tolerance when recovering write-ahead logs (WAL) files after a system crash.

The following are the options:

- 0: if the last WAL entry is corrupted, truncate the entry and either start the server normally or refuse to start.
- 1: if a WAL entry is corrupted, the server fails to start and does not recover from the crash.
- 2 (default): if a corrupted WAL entry is detected, truncate all entries after the detected corrupted entry. You can select this setting for replication replicas.
- 3: If a corrupted WAL entry is detected, skip only the corrupted entry and continue the apply WAL entries. This option can be dangerous.

#### ROCKSDB\_WAL\_SIZE\_LIMIT\_MB

Option	Description
Command-line	<code>-rocksdb-wal-size-limit-mb</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	0

Specifies the maximum size of all WAL files in megabytes before attempting to flush memtables and delete the oldest files. Default value is 0 (never rotated). Allowed range is up to 9223372036854775807.

#### ROCKSDB\_WAL\_TTL\_SECONDS

Option	Description
Command-line	<code>-rocksdb-wal-ttl-seconds</code>
Dynamic	No
Scope	Global
Data type	Numeric
Default	0

Specifies the timeout in seconds before deleting archived WAL files. Default is 0 (archived WAL files are never deleted). Allowed range is up to 9223372036854775807.

#### ROCKSDB\_WHOLE\_KEY\_FILTERING

Option	Description
Command-line	<code>-rocksdb-whole-key-filtering</code>
Dynamic	No
Scope	Global
Data type	Boolean
Default	ON

Specifies whether the bloomfilter should use the whole key for filtering instead of just the prefix. Enabled by default. Make sure that lookups use the whole key for matching.

**ROCKSDB\_WRITE\_BATCH\_FLUSH\_THRESHOLD**

Option	Description
Command-line	<code>-rocksdb-write-batch-flush-threshold</code>
Dynamic	Yes
Scope	Local
Data type	Integer
Default	0

This variable specifies the maximum size of the write batch in bytes before flushing. Only valid if `rocksdb_write_policy` is `WRITE_UNPREPARED`. There is no limit if the variable is set to the default setting.

**ROCKSDB\_WRITE\_BATCH\_MAX\_BYTES**

Option	Description
Command-line	<code>-rocksdb-write-batch-max-bytes</code>
Dynamic	Yes
Scope	Global
Data type	Numeric
Default	0

Specifies the maximum size of a RocksDB write batch in bytes. `0` means no limit. In case user exceeds the limit following error will be shown: `ERROR HY000: Status error 10 received from RocksDB: Operation aborted: Memory limit reached.`

**ROCKSDB\_WRITE\_DISABLE\_WAL**

Option	Description
Command-line	<code>-rocksdb-write-disable-wal</code>
Dynamic	Yes
Scope	Global, Session
Data type	Boolean
Default	OFF

Lets you temporarily disable writes to WAL files, which can be useful for bulk loading.

**ROCKSDB\_WRITE\_IGNORE\_MISSING\_COLUMN\_FAMILIES**

Option	Description
Command-line	<code>-rocksdb-write-ignore-missing-columns</code>
Dynamic	Yes

Option	Description
Scope	Global, Session
Data type	Boolean
Default	OFF

Specifies whether to ignore writes to column families that do not exist. Disabled by default (writes to non-existent column families are not ignored).

#### ROCKSDB\_WRITE\_POLICY

Option	Description
Command-line	<code>--rocksdb-write-policy</code>
Dynamic	No
Scope	Global
Data type	String
Default	<code>write_committed</code>

Specifies when two-phase commit data are written into the database. Allowed values are `write_committed`, `write_prepared`, and `write_unprepared`.

Value	Description
<code>write_committed</code>	Data written at commit time
<code>write_prepared</code>	Data written after the prepare phase
<code>write_unprepared</code>	Data written before the prepare phase

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)    [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-31

## 13.2.8 MyRocks status variables

MyRocks status variables provide details about the inner workings of the storage engine and they can be useful in tuning the storage engine to a particular environment.

You can view these variables and their values by running:

```
mysql> SHOW STATUS LIKE 'rocksdb%';
```

The following global status variables are available:

Name	Var Type
rocksdb_rows_deleted	Numeric
rocksdb_rows_inserted	Numeric
rocksdb_rows_read	Numeric
rocksdb_rows_unfiltered_no_snapshot	Numeric
rocksdb_rows_updated	Numeric
rocksdb_rows_expired	Numeric
rocksdb_system_rows_deleted	Numeric
rocksdb_system_rows_inserted	Numeric
rocksdb_system_rows_read	Numeric
rocksdb_system_rows_updated	Numeric
rocksdb_memtable_total	Numeric
rocksdb_memtable_unflushed	Numeric
rocksdb_queries_point	Numeric
rocksdb_queries_range	Numeric
rocksdb_covered_secondary_key_lookups	Numeric
rocksdb_additional_compactions_trigger	Numeric
rocksdb_block_cache_add	Numeric
rocksdb_block_cache_add_failures	Numeric
rocksdb_block_cache_bytes_read	Numeric
rocksdb_block_cache_bytes_write	Numeric
rocksdb_block_cache_data_add	Numeric
rocksdb_block_cache_data_bytes_insert	Numeric
rocksdb_block_cache_data_hit	Numeric
rocksdb_block_cache_data_miss	Numeric
rocksdb_block_cache_filter_add	Numeric
rocksdb_block_cache_filter_bytes_evict	Numeric
rocksdb_block_cache_filter_bytes_insert	Numeric
rocksdb_block_cache_filter_hit	Numeric
rocksdb_block_cache_filter_miss	Numeric
rocksdb_block_cache_hit	Numeric
rocksdb_block_cache_index_add	Numeric
rocksdb_block_cache_index_bytes_evict	Numeric

Name	Var Type
rocksdb_block_cache_index_bytes_insert	Numeric
rocksdb_block_cache_index_hit	Numeric
rocksdb_block_cache_index_miss	Numeric
rocksdb_block_cache_miss	Numeric
rocksdb_block_cache_compressed_hit	Numeric
rocksdb_block_cache_compressed_miss	Numeric
rocksdb_bloom_filter_prefix_checked	Numeric
rocksdb_bloom_filter_prefix_useful	Numeric
rocksdb_bloom_filter_useful	Numeric
rocksdb_bytes_read	Numeric
rocksdb_bytes_written	Numeric
rocksdb_compact_read_bytes	Numeric
rocksdb_compact_write_bytes	Numeric
rocksdb_compaction_key_drop_new	Numeric
rocksdb_compaction_key_drop_obsolete	Numeric
rocksdb_compaction_key_drop_user	Numeric
rocksdb_flush_write_bytes	Numeric
rocksdb_get_hit_l0	Numeric
rocksdb_get_hit_l1	Numeric
rocksdb_get_hit_l2_and_up	Numeric
rocksdb_get_updates_since_calls	Numeric
rocksdb_iter_bytes_read	Numeric
rocksdb_memtable_hit	Numeric
rocksdb_memtable_miss	Numeric
rocksdb_no_file_closes	Numeric
rocksdb_no_file_errors	Numeric
rocksdb_no_file_opens	Numeric
rocksdb_num_iterators	Numeric
rocksdb_number_block_not_compressed	Numeric
rocksdb_number_db_next	Numeric
rocksdb_number_db_next_found	Numeric
rocksdb_number_db_prev	Numeric
rocksdb_number_db_prev_found	Numeric
rocksdb_number_db_seek	Numeric

Name	Var Type
rocksdb_number_db_seek_found	Numeric
rocksdb_number_deletes_filtered	Numeric
rocksdb_number_keys_read	Numeric
rocksdb_number_keys_updated	Numeric
rocksdb_number_keys_written	Numeric
rocksdb_number_merge_failures	Numeric
rocksdb_number_multiget_bytes_read	Numeric
rocksdb_number_multiget_get	Numeric
rocksdb_number_multiget_keys_read	Numeric
rocksdb_number_reseeks_iteration	Numeric
rocksdb_number_sst_entry_delete	Numeric
rocksdb_number_sst_entry_merge	Numeric
rocksdb_number_sst_entry_other	Numeric
rocksdb_number_sst_entry_put	Numeric
rocksdb_number_sst_entry_singledelete	Numeric
rocksdb_number_stat_computes	Numeric
rocksdb_number_superversion_acquires	Numeric
rocksdb_number_superversion_cleanups	Numeric
rocksdb_number_superversion_releases	Numeric
rocksdb_rate_limit_delay_millis	Numeric
rocksdb_row_lock_deadlocks	Numeric
rocksdb_row_lock_wait_timeouts	Numeric
rocksdb_snapshot_conflict_errors	Numeric
rocksdb_stall_l0_file_count_limit_slowdowns	Numeric
rocksdb_stall_locked_l0_file_count_limit_slowdowns	Numeric
rocksdb_stall_l0_file_count_limit_stops	Numeric
rocksdb_stall_locked_l0_file_count_limit_stops	Numeric
rocksdb_stall_pending_compaction_limit_stops	Numeric
rocksdb_stall_pending_compaction_limit_slowdowns	Numeric
rocksdb_stall_memtable_limit_stops	Numeric
rocksdb_stall_memtable_limit_slowdowns	Numeric
rocksdb_stall_total_stops	Numeric
rocksdb_stall_total_slowdowns	Numeric
rocksdb_stall_micros	Numeric

Name	Var Type
<code>rocksdb_wal_bytes</code>	Numeric
<code>rocksdb_wal_group_syncs</code>	Numeric
<code>rocksdb_wal_synced</code>	Numeric
<code>rocksdb_write_other</code>	Numeric
<code>rocksdb_write_self</code>	Numeric
<code>rocksdb_write_timeout</code>	Numeric
<code>rocksdb_write_wal</code>	Numeric

**ROCKSDB\_ROWS\_DELETED**

This variable shows the number of rows that were deleted from MyRocks tables.

**ROCKSDB\_ROWS\_INSERTED**

This variable shows the number of rows that were inserted into MyRocks tables.

**ROCKSDB\_ROWS\_READ**

This variable shows the number of rows that were read from MyRocks tables.

**ROCKSDB\_ROWS\_UNFILTERED\_NO\_SNAPSHOT**

This variable shows how many reads need TTL and have no snapshot timestamp.

**ROCKSDB\_ROWS\_UPDATED**

This variable shows the number of rows that were updated in MyRocks tables.

**ROCKSDB\_ROWS\_EXPIRED**

This variable shows the number of expired rows in MyRocks tables.

**ROCKSDB\_SYSTEM\_ROWS\_DELETED**

This variable shows the number of rows that were deleted from MyRocks system tables.

**ROCKSDB\_SYSTEM\_ROWS\_INSERTED**

This variable shows the number of rows that were inserted into MyRocks system tables.

**ROCKSDB\_SYSTEM\_ROWS\_READ**

This variable shows the number of rows that were read from MyRocks system tables.

**ROCKSDB\_SYSTEM\_ROWS\_UPDATED**

This variable shows the number of rows that were updated in MyRocks system tables.

**ROCKSDB\_MEMTABLE\_TOTAL**

This variable shows the memory usage, in bytes, of all memtables.

**ROCKSDB\_MEMTABLE\_UNFLUSHED**

This variable shows the memory usage, in bytes, of all unflushed memtables.

**ROCKSDB\_QUERIES\_POINT**

This variable shows the number of single row queries.

**ROCKSDB\_QUERIES\_RANGE**

This variable shows the number of multi/range row queries.

**ROCKSDB\_COVERED\_SECONDARY\_KEY\_LOOKUPS**

This variable shows the number of lookups via the secondary index that returned all fields requested directly from the secondary index.

**ROCKSDB\_ADDITIONAL\_COMPACTIONS\_TRIGGER**

This variable shows the number of triggered additional compactions. MyRocks triggers an additional compaction if  $(\text{number of deletions} / \text{number of entries}) > (\text{rocksdb\_compaction\_sequential\_deletes} / \text{rocksdb\_compaction\_sequential\_deletes\_window})$  in the SST file.

**ROCKSDB\_BLOCK\_CACHE\_ADD**

This variable shows the number of blocks added to block cache.

**ROCKSDB\_BLOCK\_CACHE\_ADD\_FAILURES**

This variable shows the number of failures when adding blocks to block cache.

**ROCKSDB\_BLOCK\_CACHE\_BYTES\_READ**

This variable shows the number of bytes read from cache.

**ROCKSDB\_BLOCK\_CACHE\_BYTES\_WRITE**

This variable shows the number of bytes written into cache.

**ROCKSDB\_BLOCK\_CACHE\_DATA\_ADD**

This variable shows the number of data blocks added to block cache.

**ROCKSDB\_BLOCK\_CACHE\_DATA\_BYTES\_INSERT**

This variable shows the number of bytes of data blocks inserted into cache.

**ROCKSDB\_BLOCK\_CACHE\_DATA\_HIT**

This variable shows the number of cache hits when accessing the data block from the block cache.

**ROCKSDB\_BLOCK\_CACHE\_DATA\_MISS**

This variable shows the number of cache misses when accessing the data block from the block cache.

**ROCKSDB\_BLOCK\_CACHE\_FILTER\_ADD**

This variable shows the number of filter blocks added to block cache.

**ROCKSDB\_BLOCK\_CACHE\_FILTER\_BYTES\_EVICT**

This variable shows the number of bytes of bloom filter blocks removed from cache.

**ROCKSDB\_BLOCK\_CACHE\_FILTER\_BYTES\_INSERT**

This variable shows the number of bytes of bloom filter blocks inserted into cache.

`ROCKSDB_BLOCK_CACHE_FILTER_HIT`

This variable shows the number of times cache hit when accessing filter block from block cache.

`ROCKSDB_BLOCK_CACHE_FILTER_MISS`

This variable shows the number of times cache miss when accessing filter block from block cache.

`ROCKSDB_BLOCK_CACHE_HIT`

This variable shows the total number of block cache hits.

`ROCKSDB_BLOCK_CACHE_INDEX_ADD`

This variable shows the number of index blocks added to block cache.

`ROCKSDB_BLOCK_CACHE_INDEX_BYTES_EVICT`

This variable shows the number of bytes of index block erased from cache.

`ROCKSDB_BLOCK_CACHE_INDEX_BYTES_INSERT`

This variable shows the number of bytes of index blocks inserted into cache.

`ROCKSDB_BLOCK_CACHE_INDEX_HIT`

This variable shows the total number of block cache index hits.

`ROCKSDB_BLOCK_CACHE_INDEX_MISS`

This variable shows the number of times cache hit when accessing index block from block cache.

`ROCKSDB_BLOCK_CACHE_MISS`

This variable shows the total number of block cache misses.

`ROCKSDB_BLOCK_CACHE_COMPRESSED_HIT`

This variable shows the number of hits in the compressed block cache.

`ROCKSDB_BLOCK_CACHE_COMPRESSED_MISS`

This variable shows the number of misses in the compressed block cache.

`ROCKSDB_BLOOM_FILTER_PREFIX_CHECKED`

This variable shows the number of times bloom was checked before creating iterator on a file.

`ROCKSDB_BLOOM_FILTER_PREFIX_USEFUL`

This variable shows the number of times the check was useful in avoiding iterator creation (and thus likely IOPs).

`ROCKSDB_BLOOM_FILTER_USEFUL`

This variable shows the number of times bloom filter has avoided file reads.

`ROCKSDB_BYTES_READ`

This variable shows the total number of uncompressed bytes read. It could be either from memtables, cache, or table files.

`ROCKSDB_BYTES_WRITTEN`

This variable shows the total number of uncompressed bytes written.

`ROCKSDB_COMPACT_READ_BYTES`

This variable shows the number of bytes read during compaction

`ROCKSDB_COMPACT_WRITE_BYTES`

This variable shows the number of bytes written during compaction.

`ROCKSDB_COMPACTION_KEY_DROP_NEW`

This variable shows the number of key drops during compaction because it was overwritten with a newer value.

`ROCKSDB_COMPACTION_KEY_DROP_OBSOLETE`

This variable shows the number of key drops during compaction because it was obsolete.

`ROCKSDB_COMPACTION_KEY_DROP_USER`

This variable shows the number of key drops during compaction because user compaction function has dropped the key.

`ROCKSDB_FLUSH_WRITE_BYTES`

This variable shows the number of bytes written during flush.

`ROCKSDB_GET_HIT_L0`

This variable shows the number of `Get()` queries served by L0.

`ROCKSDB_GET_HIT_L1`

This variable shows the number of `Get()` queries served by L1.

`ROCKSDB_GET_HIT_L2_AND_UP`

This variable shows the number of `Get()` queries served by L2 and up.

`ROCKSDB_GET_UPDATES_SINCE_CALLS`

This variable shows the number of calls to `GetUpdatesSince` function. Useful to keep track of transaction log iterator refreshes

`ROCKSDB_ITER_BYTES_READ`

This variable shows the number of uncompressed bytes read from an iterator. It includes size of key and value.

`ROCKSDB_MEMTABLE_HIT`

This variable shows the number of memtable hits.

`ROCKSDB_MEMTABLE_MISS`

This variable shows the number of memtable misses.

`ROCKSDB_NO_FILE_CLOSES`

This variable shows the number of time file were closed.

`ROCKSDB_NO_FILE_ERRORS`

This variable shows number of errors trying to read in data from an sst file.

**ROCKSDB\_NO\_FILE\_OPENS**

This variable shows the number of time file were opened.

**ROCKSDB\_NUM\_ITERATORS**

This variable shows the number of currently open iterators.

**ROCKSDB\_NUMBER\_BLOCK\_NOT\_COMPRESSED**

This variable shows the number of uncompressed blocks.

**ROCKSDB\_NUMBER\_DB\_NEXT**

This variable shows the number of calls to `next`.

**ROCKSDB\_NUMBER\_DB\_NEXT\_FOUND**

This variable shows the number of calls to `next` that returned data.

**ROCKSDB\_NUMBER\_DB\_PREV**

This variable shows the number of calls to `prev`.

**ROCKSDB\_NUMBER\_DB\_PREV\_FOUND**

This variable shows the number of calls to `prev` that returned data.

**ROCKSDB\_NUMBER\_DB\_SEEK**

This variable shows the number of calls to `seek`.

**ROCKSDB\_NUMBER\_DB\_SEEK\_FOUND**

This variable shows the number of calls to `seek` that returned data.

**ROCKSDB\_NUMBER\_DELETES\_FILTERED**

This variable shows the number of deleted records that were not required to be written to storage because key did not exist.

**ROCKSDB\_NUMBER\_KEYS\_READ**

This variable shows the number of keys read.

**ROCKSDB\_NUMBER\_KEYS\_UPDATED**

This variable shows the number of keys updated, if inplace update is enabled.

**ROCKSDB\_NUMBER\_KEYS\_WRITTEN**

This variable shows the number of keys written to the database.

**ROCKSDB\_NUMBER\_MERGE\_FAILURES**

This variable shows the number of failures performing merge operator actions in RocksDB.

**ROCKSDB\_NUMBER\_MULTIGET\_BYTES\_READ**

This variable shows the number of bytes read during RocksDB `MultiGet()` calls.

**ROCKSDB\_NUMBER\_MULTIGET\_GET**

This variable shows the number `MultiGet()` requests to RocksDB.

**ROCKSDB\_NUMBER\_MULTIGET\_KEYS\_READ**

This variable shows the keys read via `MultiGet()`.

**ROCKSDB\_NUMBER\_RESEEEKS\_ITERATION**

This variable shows the number of times reseek happened inside an iteration to skip over large number of keys with same userkey.

**ROCKSDB\_NUMBER\_SST\_ENTRY\_DELETE**

This variable shows the total number of delete markers written by MyRocks.

**ROCKSDB\_NUMBER\_SST\_ENTRY\_MERGE**

This variable shows the total number of merge keys written by MyRocks.

**ROCKSDB\_NUMBER\_SST\_ENTRY\_OTHER**

This variable shows the total number of non-delete, non-merge, non-put keys written by MyRocks.

**ROCKSDB\_NUMBER\_SST\_ENTRY\_PUT**

This variable shows the total number of put keys written by MyRocks.

**ROCKSDB\_NUMBER\_SST\_ENTRY\_SINGLEDELETE**

This variable shows the total number of single delete keys written by MyRocks.

**ROCKSDB\_NUMBER\_STAT\_COMPUTES**

This variable isn't used anymore and will be removed in future releases.

**ROCKSDB\_NUMBER\_SUPERVERSION\_ACQUIRES**

This variable shows the number of times the supervision structure has been acquired in RocksDB, this is used for tracking all of the files for the database.

**ROCKSDB\_NUMBER\_SUPERVERSION\_CLEANUPS****ROCKSDB\_NUMBER\_SUPERVERSION\_RELEASES****ROCKSDB\_RATE\_LIMIT\_DELAY\_MILLIS**

This variable was removed in *Percona Server for MySQL Percona Server 5.7.23-23*.

**ROCKSDB\_ROW\_LOCK\_DEADLOCKS**

This variable shows the total number of deadlocks that have been detected since the instance was started.

**ROCKSDB\_ROW\_LOCK\_WAIT\_TIMEOUTS**

This variable shows the total number of row lock wait timeouts that have been detected since the instance was started.

**ROCKSDB\_SNAPSHOT\_CONFLICT\_ERRORS**

This variable shows the number of snapshot conflict errors occurring during write transactions that forces the transaction to rollback.

**ROCKSDB\_STALL\_L0\_FILE\_COUNT\_LIMIT\_SLOWDOWNS**

This variable shows the slowdowns in write due to L0 being close to full.

`ROCKSDB_STALL_LOCKED_L0_FILE_COUNT_LIMIT_SLOWDOWNS`

This variable shows the slowdowns in write due to L0 being close to full and compaction for L0 is already in progress.

`ROCKSDB_STALL_L0_FILE_COUNT_LIMIT_STOPS`

This variable shows the stalls in write due to L0 being full.

`ROCKSDB_STALL_LOCKED_L0_FILE_COUNT_LIMIT_STOPS`

This variable shows the stalls in write due to L0 being full and compaction for L0 is already in progress.

`ROCKSDB_STALL_PENDING_COMPACTION_LIMIT_STOPS`

This variable shows the stalls in write due to hitting limits set for max number of pending compaction bytes.

`ROCKSDB_STALL_PENDING_COMPACTION_LIMIT_SLOWDOWNS`

This variable shows the slowdowns in write due to getting close to limits set for max number of pending compaction bytes.

`ROCKSDB_STALL_MEMTABLE_LIMIT_STOPS`

This variable shows the stalls in write due to hitting max number of `memTables` allowed.

`ROCKSDB_STALL_MEMTABLE_LIMIT_SLOWDOWNS`

This variable shows the slowdowns in writes due to getting close to max number of memtables allowed.

`ROCKSDB_STALL_TOTAL_STOPS`

This variable shows the total number of write stalls.

`ROCKSDB_STALL_TOTAL_SLOWDOWNS`

This variable shows the total number of write slowdowns.

`ROCKSDB_STALL_MICROS`

This variable shows how long (in microseconds) the writer had to wait for compaction or flush to finish.

`ROCKSDB_WAL_BYTES`

This variables shows the number of bytes written to WAL.

`ROCKSDB_WAL_GROUP_SYNCS`

This variable shows the number of group commit WAL file syncs that have occurred.

`ROCKSDB_WAL_SYNCED`

This variable shows the number of times WAL sync was done.

`ROCKSDB_WRITE_OTHER`

This variable shows the number of writes processed by another thread.

**ROCKSDB\_WRITE\_SELF**

This variable shows the number of writes that were processed by a requesting thread.

**ROCKSDB\_WRITE\_TIMEDOUT**

This variable shows the number of writes ending up with timed-out.

**ROCKSDB\_WRITE\_WAL**

This variable shows the number of Write calls that request WAL.

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-07-27

## 13.2.9 Gap locks detection

The [Gap locks](#) detection is based on a Facebook *MySQL* patch.

If a transactional storage engine does not support gap locks (for example MyRocks) and a gap lock is being attempted while the transaction isolation level is either `REPEATABLE READ` or `SERIALIZABLE`, the following SQL error will be returned to the client and no actual gap lock will be taken on the effected rows.

#### Error message

```
ERROR HY000: Using Gap Lock without full unique key in multi-table or multi-statement transactions is not allowed. You need to either rewrite queries to use all unique key columns in WHERE equal conditions, or rewrite to single-table, single-statement transaction.
```

#### Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 13.2.10 Data loading

By default, MyRocks configurations are optimized for short transactions, and not for data loading. MyRocks has a couple of special session variables to speed up data loading dramatically.

### Sorted bulk loading

If your data is guaranteed to be loaded in primary key order, then this method is recommended. This method works by dropping any secondary keys first, loading data into your table in primary key order, and then restoring the secondary keys via Fast Secondary Index Creation.

#### CREATING SECONDARY INDEXES

When loading data into empty tables, it is highly recommended to drop all secondary indexes first, then loading data, and adding all secondary indexes after finishing loading data. MyRocks has a feature called `Fast Secondary`

`Index Creation`. Fast Secondary Index Creation is automatically used when executing `CREATE INDEX` or `ALTER TABLE ... ADD INDEX`. With Fast Secondary Index Creation, the secondary index entries are directly written to bottommost RocksDB levels and bypassing compaction. This significantly reduces total write volume and CPU time for decompressing and compressing data on higher levels.

#### LOADING DATA

As described above, loading data is highly recommended for tables with primary key only (no secondary keys), with all secondary indexes added after loading data.

When loading data into MyRocks tables, there are two recommended session variables:

```
SET session sql_log_bin=0;
SET session rocksdb_bulk_load=1;
```

When converting from large MyISAM/InnoDB tables, either by using the `ALTER` or `INSERT INTO SELECT` statements it's recommended that you create MyRocks tables as below (in case the table is sufficiently big it will cause the server to consume all the memory and then be terminated by the OOM killer):

```
SET session sql_log_bin=0;
SET session rocksdb_bulk_load=1;
ALTER TABLE large_myisam_table ENGINE=RocksDB;
SET session rocksdb_bulk_load=0;
```

Using `sql_log_bin=0` avoids writing to binary logs.

With `rocksdb_bulk_load` set to `1`, MyRocks enters special mode to write all inserts into bottommost RocksDB levels, and skips writing data into MemTable and the following compactions. This is very efficient way to load data.

The `rocksdb_bulk_load` mode operates with a few conditions:

- None of the data being bulk loaded can overlap with existing data in the table. The easiest way to ensure this is to always bulk load into an empty table, but the mode will allow loading some data into the table, doing other operations, and then returning and bulk loading additional data if there is no overlap between what is being loaded and what already exists.
- The data may not be visible until bulk load mode is ended (i.e. the `rocksdb_bulk_load` is set to zero again). The method that is used is building up SST files which will later be added as-is to the database. Until a particular SST has been added the data will not be visible to the rest of the system, thus issuing a `SELECT` on the table currently being bulk loaded will only show older data and will likely not show the most recently added rows. Ending the bulk load mode will cause the most recent SST file to be added. When bulk loading multiple tables, starting a new table will trigger the code to add the most recent SST file to the system – as a result, it is inadvisable to interleave `INSERT` statements to two or more tables during bulk load mode.

By default, the `rocksdb_bulk_load` mode expects all data be inserted in primary key order (or reversed order). If the data is in the reverse order (i.e. the data is descending on a normally ordered primary key or is ascending on a reverse ordered primary key), the rows are cached in chunks to switch the order to match the expected order.

Inserting one or more rows out of order will result in an error and may result in some of the data being inserted in the table and some not. To resolve the problem, one can either fix the data order of the insert, truncate the table, and restart.

### Unsorted bulk loading

If your data is not ordered in primary key order, then this method is recommended. With this method, secondary keys do not need to be dropped and restored. However, writing to the primary key no longer goes directly to SST files, and are written to temporary files for sorted first, so there is extra cost to this method.

To allow for loading unsorted data:

```
SET session sql_log_bin=0;
SET session rocksdb_bulk_load_allow_unsorted=1;
SET session rocksdb_bulk_load=1;
...
SET session rocksdb_bulk_load=0;
SET session rocksdb_bulk_load_allow_unsorted=0;
```

Note that `rocksdb_bulk_load_allow_unsorted` can only be changed when `rocksdb_bulk_load` is disabled (set to 0). In this case, all input data will go through an intermediate step that writes the rows to temporary SST files, sorts them rows in the primary key order, and then writes to final SST files in the correct order.

### Other approaches

If `rocksdb_commit_in_the_middle` is enabled, MyRocks implicitly commits every `rocksdb_bulk_load_size` records (default is 1,000) in the middle of your transaction. If your data loading fails in the middle of the statement (`LOAD DATA` or bulk `INSERT`), rows are not entirely rolled back, but some of rows are stored in the table. To restart data loading, you'll need to truncate the table and loading data again.

**Warning**

If you are loading large data without enabling `rocksdb_bulk_load` or `rocksdb_commit_in_the_middle`, please make sure transaction size is small enough. All modifications of the ongoing transactions are kept in memory.

**Other reading**

- [Data Loading](#) - this document has been used as a source for writing this documentation
- [ALTER TABLE ... ENGINE=ROCKSDB uses too much memory](#)

**Get expert help**

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-07-27

## 14. Reference

### 14.1 Percona Server for MySQL 8.1 variables

#### 14.1.1 System variables

Name	Cmd-Line	Option File
csv_mode	Yes	Yes
enforce_storage_engine	Yes	Yes
expand_fast_index_creation	Yes	No
extra_max_connections	Yes	Yes
extra_port	Yes	Yes
have_backup_locks	Yes	No
have_backup_safe_binlog_info	Yes	No
have_snapshot_cloning	Yes	No
innodb_cleaner_lsn_age_factor	Yes	Yes
innodb_corrupt_table_action	Yes	Yes
innodb_empty_free_list_algorithm	Yes	Yes
innodb_encrypt_online_alter_logs	Yes	Yes
innodb_encrypt_tables	Yes	Yes
innodb_kill_idle_transaction	Yes	Yes
innodb_max_bitmap_file_size	Yes	Yes
innodb_max_changed_pages	Yes	Yes
innodb_print_lock_wait_timeout_info	Yes	Yes
innodb_show_locks_held	Yes	Yes
innodb_temp_tablespace_encrypt	Yes	Yes
innodb_track_changed_pages	Yes	Yes
keyring_vault_config	Yes	Yes
keyring_vault_timeout	Yes	Yes
log_slow_filter	Yes	Yes
log_slow_rate_limit	Yes	Yes
log_slow_rate_type	Yes	Yes
log_slow_sp_statements	Yes	Yes
log_slow_verbosity	Yes	Yes
log_warnings_suppress	Yes	Yes

Name	Cmd-Line	Option File
proxy_protocol_networks	Yes	Yes
query_response_time_flush	Yes	No
query_response_time_range_base	Yes	Yes
query_response_time_stats	Yes	Yes
secure_log_path	Yes	Yes
slow_query_log_always_write_time	Yes	Yes
slow_query_log_use_global_control	Yes	Yes
thread_pool_high_prio_mode	Yes	Yes
thread_pool_high_prio_tickets	Yes	Yes
thread_pool_idle_timeout	Yes	Yes
thread_pool_max_threads	Yes	Yes
thread_pool_oversubscribe	Yes	Yes
thread_pool_size	Yes	Yes
thread_pool_stall_limit	Yes	Yes
thread_statistics	Yes	Yes
userstat	Yes	Yes
version_comment	Yes	Yes
version_suffix	Yes	Yes

## 14.1.2 Status variables

Name	Var Type	Var Scope
Binlog_snapshot_file	String	Global
Binlog_snapshot_position	Numeric	Global
Com_lock_binlog_for_backup	Numeric	Both
Com_lock_tables_for_backup	Numeric	Both
Com_show_client_statistics	Numeric	Both
Com_show_index_statistics	Numeric	Both
Com_show_table_statistics	Numeric	Both
Com_show_thread_statistics	Numeric	Both
Com_show_user_statistics	Numeric	Both
Com_unlock_binlog	Numeric	Both
InnoDB_background_log_sync	Numeric	Global
InnoDB_buffer_pool_pages_LRU_flushed	Numeric	Global

Name	Var Type	Var Scope
InnoDB_buffer_pool_pages_made_not_young	Numeric	Global
InnoDB_buffer_pool_pages_made_young	Numeric	Global
InnoDB_buffer_pool_pages_old	Numeric	Global
InnoDB_checkpoint_age	Numeric	Global
InnoDB_checkpoint_max_age	Numeric	Global
InnoDB_ibuf_free_list	Numeric	Global
InnoDB_ibuf_segment_size	Numeric	Global
InnoDB_lsn_current	Numeric	Global
InnoDB_lsn_flushed	Numeric	Global
InnoDB_lsn_last_checkpoint	Numeric	Global
InnoDB_master_thread_active_loops	Numeric	Global
InnoDB_master_thread_idle_loops	Numeric	Global
InnoDB_max_trx_id	Numeric	Global
InnoDB_mem_adaptive_hash	Numeric	Global
InnoDB_mem_dictionary	Numeric	Global
InnoDB_oldest_view_low_limit_trx_id	Numeric	Global
InnoDB_purge_trx_id	Numeric	Global
InnoDB_purge_undo_no	Numeric	Global
Threadpool_idle_threads	Numeric	Global
Threadpool_threads	Numeric	Global

### 14.1.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 14.2 List of features available in Percona Server for MySQL releases

Percona Server for MySQL 5.7	Percona Server for MySQL 8.0
Improved Buffer Pool Scalability	Improved Buffer Pool Scalability

Percona Server for MySQL 5.7	Percona Server for MySQL 8.0
Improved InnoDB I/O Scalability	Improved InnoDB I/O Scalability
Multiple Adaptive Hash Search Partitions	Multiple Adaptive Hash Search Partitions
Atomic write support for Fusion-io devices	Atomic write support for Fusion-io devices
Query Cache Enhancements	Feature not implemented
Improved NUMA support	Improved NUMA support
Thread Pool	Thread Pool
Suppress Warning Messages	Suppress Warning Messages
Ability to change the database for mysqlbinlog	Ability to change the database for mysqlbinlog
Fixed Size for the Read Ahead Area	Fixed Size for the Read Ahead Area
Improved MEMORY Storage Engine	Improved MEMORY Storage Engine
Restricting the number of binlog files	Restricting the number of binlog files
Ignoring missing tables in mysqldump	Ignoring missing tables in mysqldump
Too Many Connections Warning	Too Many Connections Warning
Handle Corrupted Tables	Handle Corrupted Tables
Lock-Free SHOW SLAVE STATUS	Lock-Free SHOW REPLICA STATUS
Expanded Fast Index Creation	Expanded Fast Index Creation
Percona Toolkit UDFs	Percona Toolkit UDFs
Support for Fake Changes	Support for Fake Changes
Kill Idle Transactions	Kill Idle Transactions
XtraDB changed page tracking	XtraDB changed page tracking
Enforcing Storage Engine	Replaced with upstream implementation
Utility user	Utility user
Extending the secure-file-priv server option	Extending the secure-file-priv server option
Expanded Program Option Modifiers	Feature not implemented
PAM Authentication Plugin	PAM Authentication Plugin
Log Archiving for XtraDB	Log Archiving for XtraDB
User Statistics	User Statistics
Slow Query Log	Slow Query Log
Count InnoDB Deadlocks	Count InnoDB Deadlocks
Log All Client Commands (syslog)	Log All Client Commands (syslog)
Response Time Distribution	Feature not implemented
Show Storage Engines	Show Storage Engines
Show Lock Names	Show Lock Names
Process List	Process List

Percona Server for MySQL 5.7	Percona Server for MySQL 8.0
Misc. INFORMATION_SCHEMA Tables	Misc. INFORMATION_SCHEMA Tables
Extended Show Engine InnoDB Status	Extended Show Engine InnoDB Status
Thread Based Profiling	Thread Based Profiling
XtraDB Performance Improvements for I/O-Bound Highly-Concurrent Workloads	XtraDB Performance Improvements for I/O-Bound Highly-Concurrent Workloads
Page cleaner thread tuning	Page cleaner thread tuning
Statement Timeout	Statement Timeout
Extended SELECT INTO OUTFILE/DUMPFILE	Extended SELECT INTO OUTFILE/DUMPFILE
Per-query variable statement	Per-query variable statement
Extended mysqlbinlog	Extended mysqlbinlog
Slow Query Log Rotation and Expiration	Slow Query Log Rotation and Expiration
Metrics for scalability measurement	Feature not implemented
Audit Log	Audit Log
Backup Locks	Backup Locks
CSV engine mode for a standard-compliant quote and comma parsing	CSV engine mode for a standard-compliant quote and comma parsing
Super read-only	Super read-only

## 14.2.1 Other reading

- [What Is New in MySQL 5.7](#)
- [What Is New in MySQL 8.0](#)

## 14.2.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 14.3 Percona Server for MySQL feature comparison

*Percona Server for MySQL* is a free, fully compatible, enhanced, and open source drop-in replacement for any MySQL database. It provides superior performance, scalability, and instrumentation.

*Percona Server for MySQL* is trusted by thousands of enterprises to provide better performance and concurrency for their most demanding workloads. It delivers higher value to MySQL server users with optimized performance, greater performance scalability and availability, enhanced backups, and increased visibility.

We provide these benefits by significantly enhancing *Percona Server for MySQL* as compared to the standard *MySQL* database server:

Features	Percona Server for MySQL 8.1	MySQL 8.0
Open Source	Yes	Yes
ACID Compliance	Yes	Yes
Multi-Version Concurrency Control	Yes	Yes
Row-Level Locking	Yes	Yes
Automatic Crash Recovery	Yes	Yes
Table Partitioning	Yes	Yes
Views	Yes	Yes
Subqueries	Yes	Yes
Triggers	Yes	Yes
Stored Procedures	Yes	Yes
Foreign Keys	Yes	Yes
Window Functions	Yes	Yes
Common Table Expressions	Yes	Yes
Geospatial Features (GIS, SRS)	Yes	Yes
GTID Replication	Yes	Yes
Group Replication	Yes	Yes
<a href="#">MyRocks Storage Engine</a>	Yes	No

Improvements for Developers	Percona Server for MySQL 8.1	MySQL 8.0
NoSQL Socket-Level Interface	Yes	Yes
X API Support	Yes	Yes
JSON Functions	Yes	Yes
<a href="#">InnoDB Full-Text Search Improvements</a>	Yes	No
Extra Hash/Digest Functions	Yes	No

Instrumentation and Troubleshooting Features	Percona Server for MySQL 8.1	MySQL 8.0
INFORMATION_SCHEMA Tables	95	65
Global Performance and Status Counters	853	434
Optimizer Histograms	Yes	Yes
Per-Table Performance Counters	Yes	No

Instrumentation and Troubleshooting Features	Percona Server for MySQL 8.1	MySQL 8.0
Per-Index Performance Counters	Yes	No
Per-User Performance Counters	Yes	No
Per-Client Performance Counters	Yes	No
Per-Thread Performance Counters	Yes	No
Global Query Response Time Statistics	Yes	No
<a href="#">Enhanced SHOW INNODB ENGINE STATUS</a>	Yes	No
Undo Segment Information	Yes	No
Temporary Tables Information	Yes	No
<a href="#">Extended Slow Query Logging</a>	Yes	No
<a href="#">User Statistics</a>	Yes	No

Performance and Scalability Features	Percona Server for MySQL 8.1	MySQL 8.0
InnoDB Resource Groups	Yes	Yes
Configurable Page Sizes	Yes	Yes
Contention-Aware Transaction Scheduling	Yes	Yes
Improved Scalability By Splitting Mutexes	Yes	No
<a href="#">Improved MEMORY Storage Engine</a>	Yes	No
<a href="#">Improved Flushing</a>	Yes	No
Parallel Doublewrite Buffer	Yes	Yes
<a href="#">Configurable Fast Index Creation</a>	Yes	No
Per-Column Compression for VARCHAR/BLOB and JSON	Yes	No
<a href="#">Compressed Columns with Dictionaries</a>	Yes	No

Security Features	Percona Server for MySQL 8.1	MySQL 8.0
SQL Roles	Yes	Yes
SHA-2 Based Password Hashing	Yes	Yes
Password Rotation Policy	Yes	Yes
<a href="#">PAM Authentication Plugin</a>	Yes	Enterprise

Encryption Features	Percona Server for MySQL 8.1	MySQL 8.0
Storing Keyring in a File	Yes	Yes
Storing Keyring in Hashicorp Vault	Yes	Enterprise
Encrypt InnoDB Data	Yes	Yes
Encrypt InnoDB Logs	Yes	Yes
	Yes	Yes

Encryption Features	Percona Server for MySQL 8.1	MySQL 8.0
Encrypt Built-In InnoDB Tablespace (General, System, Undo, Temp)		
<a href="#">Encrypt Binary Logs</a>	Yes	No
<a href="#">Encrypt Temporary Files</a>	Yes	No
Enforce Encryption	Yes	No

  

Operational Improvements	Percona Server for MySQL 8.1	MySQL 8.0
Atomic DDL	Yes	Yes
Transactional Data Dictionary	Yes	Yes
Instant DDL	Yes	Yes
SET PERSIST	Yes	Yes
Invisible Indexes	Yes	Yes
<a href="#">Threadpool</a>	Yes	Enterprise
<a href="#">Backup Locks</a>	Yes	No
<a href="#">Extended SHOW GRANTS</a>	Yes	No
<a href="#">Improved Handling of Corrupted Tables</a>	Yes	No
<a href="#">Ability to Kill Idle Transactions</a>	Yes	No
<a href="#">Improvements to START TRANSACTION WITH CONSISTENT SNAPSHOT</a>	Yes	No

  

Features for Running Database as a Service (DBaaS)	Percona Server for MySQL 8.1	MySQL 8.0
<a href="#">Enforce a Specific Storage Engine</a>	Yes	Yes

### 14.3.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-31

## 14.4 Understand version numbers

A version number identifies the innovation product release. The product contains the latest features, improvements, and bug fixes at the time of that release.

| 8.1.0 | -1 | | - | - | | Base version | Minor build version |

Percona uses semantic version numbering, which follows the pattern of base version and build version. Percona assigns unique, non-negative integers in increasing order for each version release. The version number combines the base [MySQL 8.1](#) version number and the minor build version.

The version numbers for Percona Server for MySQL 8.1.0-1 define the following information:

- Base version – the leftmost numbers indicate [MySQL 8.1](#) version used as a base.
- Minor build version – an internal number that increases by one every time Percona Server for MySQL is released.

## 14.4.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 14.5 Development of Percona Server for MySQL

Percona Server for MySQL is an open source project to produce a distribution of the *MySQL* Server with improved performance, scalability and diagnostics.

### 14.5.1 Submit changes

We keep the trunk in a constant state of stability to allow for a release at any time and to minimize wasted time by developers due to broken code.

#### Overview

At Percona we use [Git](#) for source control, [GitHub](#) for code hosting, and [Jira](#) for release management.

We change our software to implement new features and/or to fix bugs. Refactoring could be classed either as a new feature or a bug depending on the scope of work.

New features and bugs are targeted to specific releases. A release is part of a series. For example, 2.4 is a series in Percona XtraBackup and 2.4.15, 2.4.16, and 2.4.17 are releases in this series.

Code is proposed for merging in the form of pull requests on GitHub.

For *Percona Server for MySQL*, we have several Git branches on which development occurs: 5.5, 5.6, 5.7, and 8.0. As *Percona Server for MySQL* is not a traditional project, instead of being a set of patches against an existing product, these branches are not related. In other words, we do not merge from one release branch to another. To have your changes in several branches, you must propose branches to each release branch.

## Making a change to a project

In this case, we are going to use `percona-xtrabackup` as an example. The workflow is similar for *Percona Server for MySQL*, but the patch will need to be modified in all release branches of *Percona Server for MySQL*.

- `git branch https://github.com/percona/percona-xtrabackup/featureX` (where 'featureX' is a sensible name for the task at hand)
- (developer makes changes in featureX, testing locally)
- The Developer pushes to `https://github.com/percona/username/percona-xtrabackup/featureX`
- The developer can submit a pull request to <https://github.com/percona/percona-xtrabackup>,
- Code undergoes a review
- Once code is accepted, it can be merged

If the change also applies to a stable release (e.g. 2.4) then changes should be made on a branch of 2.4 and merged to a branch of the trunk. In this case, there should be two branches run through the param build and two merge proposals (one for the stable release and one with the changes merged to the trunk). This prevents somebody else from having to guess how to merge your changes.

## Percona Server for MySQL

The same process for *Percona Server for MySQL*, but we have several different branches (and merge requests).

### 14.5.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)    [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-10-26

## 14.6 Telemetry on Percona Server for MySQL

Percona telemetry fills in the gaps in our understanding of how you use Percona Server for MySQL to improve our products. Participation in the anonymous program is optional. You can opt-out if you prefer to not share this information.

### 14.6.1 What information is collected

At this time, telemetry is added only to the Percona packages and Docker images. Percona Server for MySQL collects only information about the installation environment. Future releases may add additional metrics.

Be assured that access to this raw data is rigorously controlled. Percona does not collect personal data. All data is anonymous and cannot be traced to a specific user. To learn more about our privacy practices, read our [Percona Privacy statement](#).

An example of the data collected is the following:

```
[{"id" : "c416c3ee-48cd-471c-9733-37c2886f8231",
  "product_family" : "PRODUCT_FAMILY_PS",
  "instanceId" : "6aef422e-56a7-4530-af9d-94cc02198343",
  "createTime" : "2023-10-16T10:46:23Z",
  "metrics":
  [{"key" : "deployment","value" : "PACKAGE"},
   {"key" : "pillar_version","value" : "8.1.0-1"},
   {"key" : "OS","value" : "Oracle Linux Server 8.8"},
   {"key" : "hardware_arch","value" : "x86_64 x86_64"}]}
```

## 14.6.2 Disable telemetry

Telemetry is enabled by default. If you decide not to send usage data to Percona, you can set the `PERCONA_TELEMETRY_DISABLE=1` environment variable for either the root user or in the operating system prior to the installation process.

[Debian-derived distribution](#)   [Red Hat-derived distribution](#)   [DOCKER](#)

Add the environment variable before the install process.

```
$ sudo PERCONA_TELEMETRY_DISABLE=1 apt install percona-server-server
```

Add the environment variable before the install process.

```
$ sudo PERCONA_TELEMETRY_DISABLE=1 yum install percona-server-server
```

Add the environment variable when running a command in a new container.

```
$ docker run -d -e MYSQL_ROOT_PASSWORD=test1234# -e PERCONA_TELEMETRY_DISABLE=1 -e --
name=percona-server percona/percona-server:8.1
```

## 14.6.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

Last update: 2023-11-27

## 14.7 Trademark policy

This [Trademark Policy](#) is to ensure that users of Percona-branded products or services know that what they receive has really been developed, approved, tested, and maintained by Percona. Trademarks help to prevent confusion in the marketplace, by distinguishing one company's or person's products and services from another's.

Percona owns a number of marks, including but not limited to Percona, XtraDB, Percona XtraDB, XtraBackup, Percona XtraBackup, *Percona Server for MySQL*, and Percona Live, plus the distinctive visual icons and logos associated with these marks. Both the unregistered and registered marks of Percona are protected.

Use of any Percona trademark in the name, URL, or another identifying characteristic of any product, service, website, or other use is not permitted without Percona's written permission with the following three limited exceptions.

*First*, you may use the appropriate Percona mark when making a nominative fair use reference to a bona fide Percona product.

*Second*, when Percona has released a product under a version of the GNU General Public License ("GPL"), you may use the appropriate Percona mark when distributing a verbatim copy of that product in accordance with the terms and conditions of the GPL.

*Third*, you may use the appropriate Percona mark to refer to a distribution of GPL-released Percona software that has been modified with minor changes for the sole purpose of allowing the software to operate on an operating system or hardware platform for which Percona has not yet released the software, provided that those third party changes do not affect the behavior, functionality, features, design or performance of the software. Users who acquire this Percona-branded software receive substantially exact implementations of the Percona software.

Percona reserves the right to revoke this authorization at any time in its sole discretion. For example, if Percona believes that your modification is beyond the scope of the limited license granted in this Policy or that your use of the Percona mark is detrimental to Percona, Percona will revoke this authorization. Upon revocation, you must immediately cease using the applicable Percona mark. If you do not immediately cease using the Percona mark upon revocation, Percona may take action to protect its rights and interests in the Percona mark. Percona does not grant any license to use any Percona mark for any other modified versions of Percona software; such use will require our prior written permission.

Neither trademark law nor any of the exceptions set forth in this Trademark Policy permit you to truncate, modify, or otherwise use any Percona mark as part of your own brand. For example, if XYZ creates a modified version of the *Percona Server for MySQL*, XYZ may not brand that modification as "XYZ Percona Server" or "Percona XYZ Server", even if that modification otherwise complies with the third exception noted above.

In all cases, you must comply with applicable law, the underlying license, and this Trademark Policy, as amended from time to time. For instance, any mention of Percona trademarks should include the full trademarked name, with proper spelling and capitalization, along with attribution of ownership to Percona Inc. For example, the full proper name for XtraBackup is Percona XtraBackup. However, it is acceptable to omit the word "Percona" for brevity on the second and subsequent uses, where such omission does not cause confusion.

In the event of doubt as to any of the conditions or exceptions outlined in this Trademark Policy, please contact [trademarks@percona.com](mailto:trademarks@percona.com) for assistance and we will do our very best to be helpful.

## 14.7.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-06-15

## 14.8 Index of INFORMATION\_SCHEMA tables

This is a list of the `INFORMATION_SCHEMA` TABLES that exist in Percona Server for MySQL with XtraDB. The entry for each table points to the page in the documentation where it's described.

- [INFORMATION\\_SCHEMA.CLIENT\\_STATISTICS](#)
- [INFORMATION\\_SCHEMA.GLOBAL\\_TEMPORARY\\_TABLES](#)
- [\[INFORMATION\\_SCHEMA.INDEX\\_STATISTICS\]](#)
- [PROCFS](#)
- [INFORMATION\\_SCHEMA.QUERY\\_RESPONSE\\_TIME](#)
- [INFORMATION\\_SCHEMA.TABLE\\_STATISTICS](#)
- [INFORMATION\\_SCHEMA.TEMPORARY\\_TABLES](#)
- [THREAD\\_STATISTICS](#)
- [INFORMATION\\_SCHEMA.USER\\_STATISTICS](#)
- [XTRADB\\_INTERNAL\\_HASH\\_TABLES](#)
- [XTRADB\\_READ\\_VIEW](#)
- [INFORMATION\\_SCHEMA.XTRADB\\_RSEG](#)
- [INFORMATION\\_SCHEMA.XTRADB\\_ZIP\\_DICT](#)
- [INFORMATION\\_SCHEMA.XTRADB\\_ZIP\\_DICT\\_COLS](#)

### 14.8.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

Last update: 2023-10-26

## 14.9 Frequently asked questions

### 14.9.1 Q: Will *Percona Server for MySQL with XtraDB* invalidate our MySQL support?

A: We don't know the details of your support contract. You should check with your *Oracle* representative. We have heard anecdotal stories from MySQL Support team members that they have customers who use *Percona Server for MySQL with XtraDB*, but you should not base your decision on that.

### 14.9.2 Q: Will we have to *GPL* our whole application if we use *Percona Server for MySQL with XtraDB*?

A: This is a common misconception about the *GPL*. We suggest reading the *Free Software Foundation*'s excellent reference material on the [GPL Version 2](#), which is the license that applies to MySQL and therefore to *Percona Server for MySQL with XtraDB*. That document contains links to many other documents which should answer your questions. *Percona* is unable to give legal advice about the *GPL*.

### 14.9.3 Q: Do I need to install *Percona* client libraries?

A: No, you don't need to change anything on the clients. *Percona Server for MySQL* is 100% compatible with all existing client libraries and connectors.

### 14.9.4 Q: When using the *Percona XtraBackup* to set up a replication replica on Debian-based systems I'm getting: "ERROR 1045 (28000): Access denied for user 'debian-sys-maint'@'localhost' (using password: YES)"

A: In case you're using the init script on Debian-based system to start `mysqld`, be sure that the password for `debian-sys-maint` user has been updated and it's the same as that user's password from the server that the backup has been taken from. The password can be seen and updated in `/etc/mysql/debian.cnf`. For more information on how to set up a replication replica using *Percona XtraBackup* see [this how-to](#).

### 14.9.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-11-23

## 14.10 Copyright and licensing information

### 14.10.1 Documentation licensing

Percona Server for MySQL documentation is (C)2009–2023 Percona LLC and/or its affiliates and is distributed under the [Creative Commons Attribution 4.0 International License](#).

### 14.10.2 Software license

*Percona Server for MySQL* is built upon MySQL from Oracle. Along with making our own modifications, we merge in changes from other sources such as community contributions and changes from MariaDB.

The original SHOW USER/TABLE/INDEX statistics code came from Google.

Percona does not require copyright assignment.

See the COPYING files accompanying the software distribution.

### 14.10.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material–forum–outline: Community Forum](#)   [:percona–logo: Get a Percona Expert](#)

---

Last update: 2023-02-14

## 14.11 Glossary

### 14.11.1 ACID

Set of properties that guarantee database transactions are processed reliably. Stands for [Atomicity](#), [Consistency](#), [Isolation](#), [Durability](#).

### 14.11.2 Atomicity

Atomicity means that database operations are applied following a “all or nothing” rule. A transaction is either fully applied or not at all.

### 14.11.3 Consistency

Consistency means that each transaction that modifies the database takes it from one consistent state to another.

### 14.11.4 Durability

Once a transaction is committed, it will remain so.

## 14.11.5 Foreign Key

A referential constraint between two tables. Example: A purchase order in the purchase\_orders table must have been made by a customer that exists in the customers table.

## 14.11.6 General Availability (GA)

A finalized version of the product which is made available to the general public. It is the final stage in the software release cycle.

## 14.11.7 Isolation

The Isolation requirement means that no transaction can interfere with another.

## 14.11.8 InnoDB

A [Storage Engine](#) for MySQL and derivatives ([Percona Server](#), [MariaDB](#)) originally written by Innobase Oy, since acquired by Oracle. It provides [ACID](#) compliant storage engine with [foreign key](#) support. As of [MySQL](#) version 5.5, InnoDB became the default storage engine on all platforms.

## 14.11.9 Jenkins

[Jenkins](#) is a continuous integration system that we use to help ensure the continued quality of the software we produce. It helps us achieve the aims of:

- no failed tests in the trunk on any platform
- aid developers in ensuring merge requests build and test on all platform
- no known performance regressions (without a damn good explanation).

## 14.11.10 LSN

The Log Sequence Number (LSN) is an 8-byte number. Every data change adds an entry to the redo log and generates an LSN. The server increments the LSN with every change.

## 14.11.11 MariaDB

A fork of [MySQL](#) that is maintained primarily by Monty Program AB. It aims to add features, and fix bugs while maintaining 100% backward compatibility with MySQL.

## 14.11.12 my.cnf

The file name of the default MySQL configuration file.

## 14.11.13 MyISAM

A [MySQL Storage Engine](#) that was the default until MySQL 5.5.

### 14.11.14 MySQL

An open source database that has spawned several distributions and forks. MySQL AB was the primary maintainer and distributor until bought by Sun Microsystems, which was then acquired by Oracle. As Oracle owns the MySQL trademark, the term MySQL is often used for the Oracle distribution of MySQL as distinct from the drop-in replacements such as [MariaDB](#) and [Percona Server for MySQL](#).

### 14.11.15 NUMA

Non-Uniform Memory Access ([NUMA](#)) is a computer memory design used in multiprocessing, where the memory access time depends on the memory location relative to a processor. Under NUMA, a processor can access its own local memory faster than non-local memory, that is, memory local to another processor or memory shared between processors. The whole system may still operate as one unit, and all memory is basically accessible from everywhere but at a potentially higher latency and lower performance.

### 14.11.16 Percona Server for MySQL

The Percona branch of [MySQL](#) with performance and management improvements.

### 14.11.17 Storage Engine

A storage engine is a piece of software that implements the details of data storage and retrieval for a database system. This term is primarily used within the [MySQL](#) ecosystem due to it being the first widely used relational database to have an abstraction layer around storage. It is analogous to a Virtual File System layer in an Operating System. A VFS layer allows an operating system to read and write multiple file systems (e.g. FAT, NTFS, XFS, ext3) and a Storage Engine layer allows a database server to access tables stored in different engines (for example, [MyISAM](#) or [InnoDB](#)).

### 14.11.18 Tech Preview

A tech preview item can be a feature, a variable, or a value within a variable. The term designates that the item is not yet ready for production use and is not included in support by SLA. A tech preview item is included in a release so that users can provide feedback. The item is either updated and released as [general availability\(GA\)](#) or removed if not useful. The item's functionality can change from tech preview to GA.

### 14.11.19 XtraDB

The Percona improved version of [InnoDB](#) provides performance, features, and reliability above what is shipped by Oracle in [InnoDB](#).

### 14.11.20 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[:material-forum-outline: Community Forum](#)   [:percona-logo: Get a Percona Expert](#)

---

Last update: 2023-01-12