

# MySQL Information Schema

---

## Abstract

This is the MySQL Information Schema extract from the MySQL 5.6 Reference Manual.

For legal information, see the [Legal Notices](#).

For help with using MySQL, please visit the [MySQL Forums](#), where you can discuss your issues with other MySQL users.

Document generated on: 2021-05-13 (revision: 69769)

---

---

# Table of Contents

Preface and Legal Notices .....	v
1 INFORMATION_SCHEMA Tables .....	1
2 Introduction .....	3
3 INFORMATION_SCHEMA Table Reference .....	7
4 INFORMATION_SCHEMA General Tables .....	11
4.1 INFORMATION_SCHEMA General Table Reference .....	11
4.2 The INFORMATION_SCHEMA CHARACTER_SETS Table .....	12
4.3 The INFORMATION_SCHEMA COLLATIONS Table .....	13
4.4 The INFORMATION_SCHEMA COLLATION_CHARACTER_SET_APPLICABILITY Table .....	13
4.5 The INFORMATION_SCHEMA COLUMNS Table .....	14
4.6 The INFORMATION_SCHEMA COLUMN_PRIVILEGES Table .....	16
4.7 The INFORMATION_SCHEMA ENGINES Table .....	17
4.8 The INFORMATION_SCHEMA EVENTS Table .....	18
4.9 The INFORMATION_SCHEMA GLOBAL_STATUS and SESSION_STATUS Tables .....	21
4.10 The INFORMATION_SCHEMA GLOBAL_VARIABLES and SESSION_VARIABLES Tables .....	21
4.11 The INFORMATION_SCHEMA KEY_COLUMN_USAGE Table .....	22
4.12 The INFORMATION_SCHEMA OPTIMIZER_TRACE Table .....	23
4.13 The INFORMATION_SCHEMA PARAMETERS Table .....	23
4.14 The INFORMATION_SCHEMA PARTITIONS Table .....	25
4.15 The INFORMATION_SCHEMA PLUGINS Table .....	28
4.16 The INFORMATION_SCHEMA PROCESSLIST Table .....	29
4.17 The INFORMATION_SCHEMA PROFILING Table .....	30
4.18 The INFORMATION_SCHEMA REFERENTIAL_CONSTRAINTS Table .....	31
4.19 The INFORMATION_SCHEMA ROUTINES Table .....	32
4.20 The INFORMATION_SCHEMA SCHEMATA Table .....	35
4.21 The INFORMATION_SCHEMA SCHEMA_PRIVILEGES Table .....	36
4.22 The INFORMATION_SCHEMA STATISTICS Table .....	36
4.23 The INFORMATION_SCHEMA TABLES Table .....	38
4.24 The INFORMATION_SCHEMA TABLESPACES Table .....	41
4.25 The INFORMATION_SCHEMA TABLE_CONSTRAINTS Table .....	41
4.26 The INFORMATION_SCHEMA TABLE_PRIVILEGES Table .....	42
4.27 The INFORMATION_SCHEMA TRIGGERS Table .....	43
4.28 The INFORMATION_SCHEMA USER_PRIVILEGES Table .....	45
4.29 The INFORMATION_SCHEMA VIEWS Table .....	45
5 INFORMATION_SCHEMA InnoDB Tables .....	47
5.1 INFORMATION_SCHEMA InnoDB Table Reference .....	47
5.2 The INFORMATION_SCHEMA INNODB_BUFFER_PAGE Table .....	48
5.3 The INFORMATION_SCHEMA INNODB_BUFFER_PAGE_LRU Table .....	51
5.4 The INFORMATION_SCHEMA INNODB_BUFFER_POOL_STATS Table .....	54
5.5 The INFORMATION_SCHEMA INNODB_CMP and INNODB_CMP_RESET Tables .....	57
5.6 The INFORMATION_SCHEMA INNODB_CMPMEM and INNODB_CMPMEM_RESET Tables .....	58
5.7 The INFORMATION_SCHEMA INNODB_CMP_PER_INDEX and INNODB_CMP_PER_INDEX_RESET Tables .....	60
5.8 The INFORMATION_SCHEMA INNODB_FT_BEING_DELETED Table .....	61
5.9 The INFORMATION_SCHEMA INNODB_FT_CONFIG Table .....	62
5.10 The INFORMATION_SCHEMA INNODB_FT_DEFAULT_STOPWORD Table .....	63
5.11 The INFORMATION_SCHEMA INNODB_FT_DELETED Table .....	64
5.12 The INFORMATION_SCHEMA INNODB_FT_INDEX_CACHE Table .....	65
5.13 The INFORMATION_SCHEMA INNODB_FT_INDEX_TABLE Table .....	66
5.14 The INFORMATION_SCHEMA INNODB_LOCKS Table .....	68
5.15 The INFORMATION_SCHEMA INNODB_LOCK_WAITS Table .....	69
5.16 The INFORMATION_SCHEMA INNODB_METRICS Table .....	70

5.17	The INFORMATION_SCHEMA INNODB_SYS_COLUMNS Table .....	72
5.18	The INFORMATION_SCHEMA INNODB_SYS_DATAFILES Table .....	73
5.19	The INFORMATION_SCHEMA INNODB_SYS_FIELDS Table .....	74
5.20	The INFORMATION_SCHEMA INNODB_SYS_FOREIGN Table .....	74
5.21	The INFORMATION_SCHEMA INNODB_SYS_FOREIGN_COLS Table .....	75
5.22	The INFORMATION_SCHEMA INNODB_SYS_INDEXES Table .....	76
5.23	The INFORMATION_SCHEMA INNODB_SYS_TABLES Table .....	77
5.24	The INFORMATION_SCHEMA INNODB_SYS_TABLESPACES Table .....	78
5.25	The INFORMATION_SCHEMA INNODB_SYS_TABLESTATS View .....	79
5.26	The INFORMATION_SCHEMA INNODB_TRX Table .....	80
6	INFORMATION_SCHEMA NDB Cluster Tables .....	85
6.1	The INFORMATION_SCHEMA FILES Table .....	85
6.2	The INFORMATION_SCHEMA ndb_transid_mysql_connection_map Table .....	90
7	INFORMATION_SCHEMA Thread Pool Tables .....	93
7.1	INFORMATION_SCHEMA Thread Pool Table Reference .....	93
7.2	The INFORMATION_SCHEMA TP_THREAD_GROUP_STATE Table .....	93
7.3	The INFORMATION_SCHEMA TP_THREAD_GROUP_STATS Table .....	95
7.4	The INFORMATION_SCHEMA TP_THREAD_STATE Table .....	96
8	INFORMATION_SCHEMA Connection-Control Tables .....	99
8.1	INFORMATION_SCHEMA Connection-Control Table Reference .....	99
8.2	The INFORMATION_SCHEMA CONNECTION_CONTROL_FAILED_LOGIN_ATTEMPTS Table .....	99
9	INFORMATION_SCHEMA MySQL Enterprise Firewall Tables .....	101
9.1	INFORMATION_SCHEMA Firewall Table Reference .....	101
9.2	The INFORMATION_SCHEMA MYSQL_FIREWALL_USERS Table .....	101
9.3	The INFORMATION_SCHEMA MYSQL_FIREWALL_WHITELIST Table .....	101
10	Extensions to SHOW Statements .....	103
11	MySQL 5.6 FAQ: INFORMATION_SCHEMA .....	105

---

## Preface and Legal Notices

This is the MySQL Information Schema extract from the MySQL 5.6 Reference Manual.

**Licensing information—MySQL 5.6.** This product may include third-party software, used under license. If you are using a *Commercial* release of MySQL 5.6, see the [MySQL 5.6 Commercial Release License Information User Manual](#) for licensing information, including licensing information relating to third-party software that may be included in this Commercial release. If you are using a *Community* release of MySQL 5.6, see the [MySQL 5.6 Community Release License Information User Manual](#) for licensing information, including licensing information relating to third-party software that may be included in this Community release.

**Licensing information—MySQL NDB Cluster 7.3.** This product may include third-party software, used under license. If you are using a *Commercial* release of NDB Cluster 7.3, see the [MySQL NDB Cluster 7.3 Commercial Release License Information User Manual](#) for licensing information relating to third-party software that may be included in this Commercial release. If you are using a *Community* release of NDB Cluster 7.3, see the [MySQL NDB Cluster 7.3 Community Release License Information User Manual](#) for licensing information relating to third-party software that may be included in this Community release.

**Licensing information—MySQL NDB Cluster 7.4.** This product may include third-party software, used under license. If you are using a *Commercial* release of NDB Cluster 7.4, see the [MySQL NDB Cluster 7.4 Commercial Release License Information User Manual](#) for licensing information relating to third-party software that may be included in this Commercial release. If you are using a *Community* release of NDB Cluster 7.4, see the [MySQL NDB Cluster 7.4 Community Release License Information User Manual](#) for licensing information relating to third-party software that may be included in this Community release.

## Legal Notices

Copyright © 1997, 2021, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including

applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how Oracle disseminates it (that is, electronically for download on a Web site with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Oracle. Oracle and/or its affiliates reserve any and all rights to this documentation not expressly granted above.

## Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

<https://www.oracle.com/corporate/accessibility/>.

## Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

---

## Chapter 1 INFORMATION\_SCHEMA Tables

[INFORMATION\\_SCHEMA](#) provides access to database *metadata*, information about the MySQL server such as the name of a database or table, the data type of a column, or access privileges. Other terms that are sometimes used for this information are *data dictionary* and *system catalog*.



---

## Chapter 2 Introduction

`INFORMATION_SCHEMA` provides access to database *metadata*, information about the MySQL server such as the name of a database or table, the data type of a column, or access privileges. Other terms that are sometimes used for this information are *data dictionary* and *system catalog*.

- [INFORMATION\\_SCHEMA Usage Notes](#)
- [Character Set Considerations](#)
- [INFORMATION\\_SCHEMA as Alternative to SHOW Statements](#)
- [INFORMATION\\_SCHEMA and Privileges](#)
- [Performance Considerations](#)
- [Standards Considerations](#)
- [Conventions in the INFORMATION\\_SCHEMA Reference Sections](#)
- [Related Information](#)

## INFORMATION\_SCHEMA Usage Notes

`INFORMATION_SCHEMA` is a database within each MySQL instance, the place that stores information about all the other databases that the MySQL server maintains. The `INFORMATION_SCHEMA` database contains several read-only tables. They are actually views, not base tables, so there are no files associated with them, and you cannot set triggers on them. Also, there is no database directory with that name.

Although you can select `INFORMATION_SCHEMA` as the default database with a `USE` statement, you can only read the contents of tables, not perform `INSERT`, `UPDATE`, or `DELETE` operations on them.

Here is an example of a statement that retrieves information from `INFORMATION_SCHEMA`:

```
mysql> SELECT table_name, table_type, engine
        FROM information_schema.tables
        WHERE table_schema = 'db5'
        ORDER BY table_name;
```

table_name	table_type	engine
fk	BASE TABLE	InnoDB
fk2	BASE TABLE	InnoDB
goto	BASE TABLE	MyISAM
into	BASE TABLE	MyISAM
k	BASE TABLE	MyISAM
kurs	BASE TABLE	MyISAM
loop	BASE TABLE	MyISAM
pk	BASE TABLE	InnoDB
t	BASE TABLE	MyISAM
t2	BASE TABLE	MyISAM
t3	BASE TABLE	MyISAM
t7	BASE TABLE	MyISAM
tables	BASE TABLE	MyISAM
v	VIEW	NULL
v2	VIEW	NULL
v3	VIEW	NULL
v56	VIEW	NULL

```
17 rows in set (0.01 sec)
```

Explanation: The statement requests a list of all the tables in database `db5`, showing just three pieces of information: the name of the table, its type, and its storage engine.

## Character Set Considerations

The definition for character columns (for example, `TABLES.TABLE_NAME`) is generally `VARCHAR(N) CHARACTER SET utf8` where `N` is at least 64. MySQL uses the default collation for this character set (`utf8_general_ci`) for all searches, sorts, comparisons, and other string operations on such columns.

Because some MySQL objects are represented as files, searches in `INFORMATION_SCHEMA` string columns can be affected by file system case sensitivity. For more information, see [Using Collation in INFORMATION\\_SCHEMA Searches](#).

## INFORMATION\_SCHEMA as Alternative to SHOW Statements

The `SELECT ... FROM INFORMATION_SCHEMA` statement is intended as a more consistent way to provide access to the information provided by the various `SHOW` statements that MySQL supports (`SHOW DATABASES`, `SHOW TABLES`, and so forth). Using `SELECT` has these advantages, compared to `SHOW`:

- It conforms to Codd's rules, because all access is done on tables.
- You can use the familiar syntax of the `SELECT` statement, and only need to learn some table and column names.
- The implementor need not worry about adding keywords.
- You can filter, sort, concatenate, and transform the results from `INFORMATION_SCHEMA` queries into whatever format your application needs, such as a data structure or a text representation to parse.
- This technique is more interoperable with other database systems. For example, Oracle Database users are familiar with querying tables in the Oracle data dictionary.

Because `SHOW` is familiar and widely used, the `SHOW` statements remain as an alternative. In fact, along with the implementation of `INFORMATION_SCHEMA`, there are enhancements to `SHOW` as described in [Chapter 10, Extensions to SHOW Statements](#).

## INFORMATION\_SCHEMA and Privileges

For most `INFORMATION_SCHEMA` tables, each MySQL user has the right to access them, but can see only the rows in the tables that correspond to objects for which the user has the proper access privileges. In some cases (for example, the `ROUTINE_DEFINITION` column in the `INFORMATION_SCHEMA.ROUTINES` table), users who have insufficient privileges see `NULL`. Some tables have different privilege requirements; for these, the requirements are mentioned in the applicable table descriptions. For example, `InnoDB` tables (tables with names that begin with `INNODB_`) require the `PROCESS` privilege.

The same privileges apply to selecting information from `INFORMATION_SCHEMA` and viewing the same information through `SHOW` statements. In either case, you must have some privilege on an object to see information about it.

## Performance Considerations

`INFORMATION_SCHEMA` queries that search for information from more than one database might take a long time and impact performance. To check the efficiency of a query, you can use `EXPLAIN`. For information about using `EXPLAIN` output to tune `INFORMATION_SCHEMA` queries, see [Optimizing INFORMATION\\_SCHEMA Queries](#).

## Standards Considerations

The implementation for the `INFORMATION_SCHEMA` table structures in MySQL follows the ANSI/ISO SQL:2003 standard Part 11 *Schemata*. Our intent is approximate compliance with SQL:2003 core feature F021 *Basic information schema*.

Users of SQL Server 2000 (which also follows the standard) may notice a strong similarity. However, MySQL has omitted many columns that are not relevant for our implementation, and added columns that are MySQL-specific. One such added column is the `ENGINE` column in the `INFORMATION_SCHEMA TABLES` table.

Although other DBMSs use a variety of names, like `syscat` or `system`, the standard name is `INFORMATION_SCHEMA`.

To avoid using any name that is reserved in the standard or in DB2, SQL Server, or Oracle, we changed the names of some columns marked “MySQL extension”. (For example, we changed `COLLATION` to `TABLE_COLLATION` in the `TABLES` table.) See the list of reserved words near the end of this article: <https://web.archive.org/web/20070428032454/http://www.dbazine.com/db2/db2-disarticles/gulutzan5>.

## Conventions in the INFORMATION\_SCHEMA Reference Sections

The following sections describe each of the tables and columns in `INFORMATION_SCHEMA`. For each column, there are three pieces of information:

- “`INFORMATION_SCHEMA` Name” indicates the name for the column in the `INFORMATION_SCHEMA` table. This corresponds to the standard SQL name unless the “Remarks” field says “MySQL extension.”
- “`SHOW` Name” indicates the equivalent field name in the closest `SHOW` statement, if there is one.
- “Remarks” provides additional information where applicable. If this field is `NULL`, it means that the value of the column is always `NULL`. If this field says “MySQL extension,” the column is a MySQL extension to standard SQL.

Many sections indicate what `SHOW` statement is equivalent to a `SELECT` that retrieves information from `INFORMATION_SCHEMA`. For `SHOW` statements that display information for the default database if you omit a `FROM db_name` clause, you can often select information for the default database by adding an `AND TABLE_SCHEMA = SCHEMA()` condition to the `WHERE` clause of a query that retrieves information from an `INFORMATION_SCHEMA` table.

## Related Information

These sections discuss additional `INFORMATION_SCHEMA`-related topics:

- information about `INFORMATION_SCHEMA` tables specific to the `InnoDB` storage engine: [Chapter 5, `INFORMATION\_SCHEMA InnoDB Tables`](#)
- information about `INFORMATION_SCHEMA` tables specific to the `NDB` storage engine (NDB Cluster): [Chapter 6, `INFORMATION\_SCHEMA NDB Cluster Tables`](#)
- information about `INFORMATION_SCHEMA` tables specific to the thread pool plugin: [Chapter 7, `INFORMATION\_SCHEMA Thread Pool Tables`](#)
- information about `INFORMATION_SCHEMA` tables specific to the `CONNECTION_CONTROL` plugin: [Chapter 8, `INFORMATION\_SCHEMA Connection-Control Tables`](#)
- Answers to questions that are often asked concerning the `INFORMATION_SCHEMA` database: [Chapter 11, `MySQL 5.6 FAQ: INFORMATION\_SCHEMA`](#)

- `INFORMATION_SCHEMA` queries and the optimizer: [Optimizing INFORMATION\\_SCHEMA Queries](#)
- The effect of collation on `INFORMATION_SCHEMA` comparisons: [Using Collation in INFORMATION\\_SCHEMA Searches](#)
- As of MySQL 5.6.49, you must have the `PROCESS` privilege to query this table.

## Chapter 3 INFORMATION\_SCHEMA Table Reference

The following table summarizes all available [INFORMATION\\_SCHEMA](#) tables. For greater detail, see the individual table descriptions.

**Table 3.1 INFORMATION\_SCHEMA Tables**

Name	Description	Introduced
<a href="#">CHARACTER_SETS</a>	Available character sets	
<a href="#">COLLATION_CHARACTER_SET_APPLICABLE</a>	Character sets applicable to each collation	
<a href="#">COLLATIONS</a>	Collations for each character set	
<a href="#">COLUMN_PRIVILEGES</a>	Privileges defined on columns	
<a href="#">COLUMNS</a>	Columns in each table	
<a href="#">CONNECTION_CONTROL_FAILED</a>	Current number of consecutive failed connection attempts per account	5.6.35
<a href="#">ENGINES</a>	Storage engine properties	
<a href="#">EVENTS</a>	Event Manager events	
<a href="#">FILES</a>	Files that store tablespace data	
<a href="#">GLOBAL_STATUS</a>	Global status variables	
<a href="#">GLOBAL_VARIABLES</a>	Global system variables	
<a href="#">INNODB_BUFFER_PAGE</a>	Pages in InnoDB buffer pool	
<a href="#">INNODB_BUFFER_PAGE_LRU</a>	LRU ordering of pages in InnoDB buffer pool	
<a href="#">INNODB_BUFFER_POOL_STATS</a>	InnoDB buffer pool statistics	
<a href="#">INNODB_CMP</a>	Status for operations related to compressed InnoDB tables	
<a href="#">INNODB_CMP_PER_INDEX</a>	Status for operations related to compressed InnoDB tables and indexes	
<a href="#">INNODB_CMP_PER_INDEX_RESET</a>	Status for operations related to compressed InnoDB tables and indexes	
<a href="#">INNODB_CMP_RESET</a>	Status for operations related to compressed InnoDB tables	
<a href="#">INNODB_CMPMEM</a>	Status for compressed pages within InnoDB buffer pool	
<a href="#">INNODB_CMPMEM_RESET</a>	Status for compressed pages within InnoDB buffer pool	
<a href="#">INNODB_FT_BEING_DELETED</a>	Snapshot of INNODB_FT_DELETED table	
<a href="#">INNODB_FT_CONFIG</a>	Metadata for InnoDB table FULLTEXT index and associated processing	
<a href="#">INNODB_FT_DEFAULT_STOPWORDS</a>	Default list of stopwords for InnoDB FULLTEXT indexes	
<a href="#">INNODB_FT_DELETED</a>	Rows deleted from InnoDB table FULLTEXT index	

Name	Description	Introduced
<code>INNODB_FT_INDEX_CACHE</code>	Token information for newly inserted rows in InnoDB FULLTEXT index	
<code>INNODB_FT_INDEX_TABLE</code>	Inverted index information for processing text searches against InnoDB table FULLTEXT index	
<code>INNODB_LOCK_WAITS</code>	InnoDB transaction lock-wait information	
<code>INNODB_LOCKS</code>	InnoDB transaction lock information	
<code>INNODB_METRICS</code>	InnoDB performance information	
<code>INNODB_SYS_COLUMNS</code>	Columns in each InnoDB table	
<code>INNODB_SYS_DATAFILES</code>	Data file path information for InnoDB file-per-table and general tablespaces	
<code>INNODB_SYS_FIELDS</code>	Key columns of InnoDB indexes	
<code>INNODB_SYS_FOREIGN</code>	InnoDB foreign-key metadata	
<code>INNODB_SYS_FOREIGN_COLS</code>	InnoDB foreign-key column status information	
<code>INNODB_SYS_INDEXES</code>	InnoDB index metadata	
<code>INNODB_SYS_TABLES</code>	InnoDB table metadata	
<code>INNODB_SYS_TABLESPACES</code>	InnoDB file-per-table, general, and undo tablespace metadata	
<code>INNODB_SYS_TABLESTATS</code>	InnoDB table low-level status information	
<code>INNODB_TRX</code>	Active InnoDB transaction information	
<code>KEY_COLUMN_USAGE</code>	Which key columns have constraints	
<code>MYSQL_FIREWALL_USERS</code>	Firewall in-memory data for account profiles	
<code>MYSQL_FIREWALL_WHITELIST</code>	Firewall in-memory data for account profile allowlists	
<code>ndb_transid_mysql_connect</code>	NDB transaction information	
<code>OPTIMIZER_TRACE</code>	Information produced by optimizer trace activity	
<code>PARAMETERS</code>	Stored routine parameters and stored function return values	
<code>PARTITIONS</code>	Table partition information	
<code>PLUGINS</code>	Plugin information	
<code>PROCESSLIST</code>	Information about currently executing threads	
<code>PROFILING</code>	Statement profiling information	
<code>REFERENTIAL_CONSTRAINTS</code>	Foreign key information	
<code>ROUTINES</code>	Stored routine information	
<code>SCHEMA_PRIVILEGES</code>	Privileges defined on schemas	

<b>Name</b>	<b>Description</b>	<b>Introduced</b>
<a href="#">SCHEMATA</a>	Schema information	
<a href="#">SESSION_STATUS</a>	Status variables for current session	
<a href="#">SESSION_VARIABLES</a>	System variables for current session	
<a href="#">STATISTICS</a>	Table index statistics	
<a href="#">TABLE_CONSTRAINTS</a>	Which tables have constraints	
<a href="#">TABLE_PRIVILEGES</a>	Privileges defined on tables	
<a href="#">TABLES</a>	Table information	
<a href="#">TABLESPACES</a>	Tablespace information	
<a href="#">TP_THREAD_GROUP_STATE</a>	Thread pool thread group states	
<a href="#">TP_THREAD_GROUP_STATS</a>	Thread pool thread group statistics	
<a href="#">TP_THREAD_STATE</a>	Thread pool thread information	
<a href="#">TRIGGERS</a>	Trigger information	
<a href="#">USER_PRIVILEGES</a>	Privileges defined globally per user	
<a href="#">VIEWS</a>	View information	



---

# Chapter 4 INFORMATION\_SCHEMA General Tables

## Table of Contents

4.1 INFORMATION_SCHEMA General Table Reference .....	11
4.2 The INFORMATION_SCHEMA CHARACTER_SETS Table .....	12
4.3 The INFORMATION_SCHEMA COLLATIONS Table .....	13
4.4 The INFORMATION_SCHEMA COLLATION_CHARACTER_SET_APPLICABILITY Table .....	13
4.5 The INFORMATION_SCHEMA COLUMNS Table .....	14
4.6 The INFORMATION_SCHEMA COLUMN_PRIVILEGES Table .....	16
4.7 The INFORMATION_SCHEMA ENGINES Table .....	17
4.8 The INFORMATION_SCHEMA EVENTS Table .....	18
4.9 The INFORMATION_SCHEMA GLOBAL_STATUS and SESSION_STATUS Tables .....	21
4.10 The INFORMATION_SCHEMA GLOBAL_VARIABLES and SESSION_VARIABLES Tables ....	21
4.11 The INFORMATION_SCHEMA KEY_COLUMN_USAGE Table .....	22
4.12 The INFORMATION_SCHEMA OPTIMIZER_TRACE Table .....	23
4.13 The INFORMATION_SCHEMA PARAMETERS Table .....	23
4.14 The INFORMATION_SCHEMA PARTITIONS Table .....	25
4.15 The INFORMATION_SCHEMA PLUGINS Table .....	28
4.16 The INFORMATION_SCHEMA PROCESSLIST Table .....	29
4.17 The INFORMATION_SCHEMA PROFILING Table .....	30
4.18 The INFORMATION_SCHEMA REFERENTIAL_CONSTRAINTS Table .....	31
4.19 The INFORMATION_SCHEMA ROUTINES Table .....	32
4.20 The INFORMATION_SCHEMA SCHEMATA Table .....	35
4.21 The INFORMATION_SCHEMA SCHEMA_PRIVILEGES Table .....	36
4.22 The INFORMATION_SCHEMA STATISTICS Table .....	36
4.23 The INFORMATION_SCHEMA TABLES Table .....	38
4.24 The INFORMATION_SCHEMA TABLESPACES Table .....	41
4.25 The INFORMATION_SCHEMA TABLE_CONSTRAINTS Table .....	41
4.26 The INFORMATION_SCHEMA TABLE_PRIVILEGES Table .....	42
4.27 The INFORMATION_SCHEMA TRIGGERS Table .....	43
4.28 The INFORMATION_SCHEMA USER_PRIVILEGES Table .....	45
4.29 The INFORMATION_SCHEMA VIEWS Table .....	45

The following sections describe what may be denoted as the “general” set of [INFORMATION\\_SCHEMA](#) tables. These are the tables not associated with particular storage engines, components, or plugins.

## 4.1 INFORMATION\_SCHEMA General Table Reference

The following table summarizes [INFORMATION\\_SCHEMA](#) general tables. For greater detail, see the individual table descriptions.

**Table 4.1 INFORMATION\_SCHEMA General Tables**

Name	Description
<a href="#">CHARACTER_SETS</a>	Available character sets
<a href="#">COLLATION_CHARACTER_SET_APPLICABILITY</a>	Character set applicable to each collation
<a href="#">COLLATIONS</a>	Collations for each character set
<a href="#">COLUMN_PRIVILEGES</a>	Privileges defined on columns
<a href="#">COLUMNS</a>	Columns in each table
<a href="#">ENGINES</a>	Storage engine properties

Name	Description
EVENTS	Event Manager events
FILES	Files that store tablespace data
GLOBAL_STATUS	Global status variables
GLOBAL_VARIABLES	Global system variables
KEY_COLUMN_USAGE	Which key columns have constraints
ndb_transid_mysql_connection_map	NDB transaction information
OPTIMIZER_TRACE	Information produced by optimizer trace activity
PARAMETERS	Stored routine parameters and stored function return values
PARTITIONS	Table partition information
PLUGINS	Plugin information
PROCESSLIST	Information about currently executing threads
PROFILING	Statement profiling information
REFERENTIAL_CONSTRAINTS	Foreign key information
ROUTINES	Stored routine information
SCHEMA_PRIVILEGES	Privileges defined on schemas
SCHEMATA	Schema information
SESSION_STATUS	Status variables for current session
SESSION_VARIABLES	System variables for current session
STATISTICS	Table index statistics
TABLE_CONSTRAINTS	Which tables have constraints
TABLE_PRIVILEGES	Privileges defined on tables
TABLES	Table information
TABLESPACES	Tablespace information
TRIGGERS	Trigger information
USER_PRIVILEGES	Privileges defined globally per user
VIEWS	View information

## 4.2 The INFORMATION\_SCHEMA CHARACTER\_SETS Table

The `CHARACTER_SETS` table provides information about available character sets.

The `CHARACTER_SETS` table has these columns:

- `CHARACTER_SET_NAME`

The character set name.

- `DEFAULT_COLLATE_NAME`

The default collation for the character set.

- `DESCRIPTION`

A description of the character set.

- `MAXLEN`

The maximum number of bytes required to store one character.

## Notes

Character set information is also available from the [SHOW CHARACTER SET](#) statement. See [SHOW CHARACTER SET Statement](#). The following statements are equivalent:

```
SELECT * FROM INFORMATION_SCHEMA.CHARACTER_SETS
  [WHERE CHARACTER_SET_NAME LIKE 'wild']
SHOW CHARACTER SET
  [LIKE 'wild']
```

## 4.3 The INFORMATION\_SCHEMA COLLATIONS Table

The [COLLATIONS](#) table provides information about collations for each character set.

The [COLLATIONS](#) table has these columns:

- [COLLATION\\_NAME](#)

The collation name.

- [CHARACTER\\_SET\\_NAME](#)

The name of the character set with which the collation is associated.

- [ID](#)

The collation ID.

- [IS\\_DEFAULT](#)

Whether the collation is the default for its character set.

- [IS\\_COMPILED](#)

Whether the character set is compiled into the server.

- [SORTLEN](#)

This is related to the amount of memory required to sort strings expressed in the character set.

## Notes

Collation information is also available from the [SHOW COLLATION](#) statement. See [SHOW COLLATION Statement](#). The following statements are equivalent:

```
SELECT COLLATION_NAME FROM INFORMATION_SCHEMA.COLLATIONS
  [WHERE COLLATION_NAME LIKE 'wild']
SHOW COLLATION
  [LIKE 'wild']
```

## 4.4 The INFORMATION\_SCHEMA COLLATION\_CHARACTER\_SET\_APPLICABILITY Table

The [COLLATION\\_CHARACTER\\_SET\\_APPLICABILITY](#) table indicates what character set is applicable for what collation.

The [COLLATION\\_CHARACTER\\_SET\\_APPLICABILITY](#) table has these columns:

- [COLLATION\\_NAME](#)

The collation name.

- `CHARACTER_SET_NAME`

The name of the character set with which the collation is associated.

## Notes

The `COLLATION_CHARACTER_SET_APPLICABILITY` columns are equivalent to the first two columns displayed by the `SHOW COLLATION` statement.

## 4.5 The `INFORMATION_SCHEMA COLUMNS` Table

The `COLUMNS` table provides information about columns in tables.

The `COLUMNS` table has these columns:

- `TABLE_CATALOG`

The name of the catalog to which the table containing the column belongs. This value is always `def`.

- `TABLE_SCHEMA`

The name of the schema (database) to which the table containing the column belongs.

- `TABLE_NAME`

The name of the table containing the column.

- `COLUMN_NAME`

The name of the column.

- `ORDINAL_POSITION`

The position of the column within the table. `ORDINAL_POSITION` is necessary because you might want to say `ORDER BY ORDINAL_POSITION`. Unlike `SHOW COLUMNS`, `SELECT` from the `COLUMNS` table does not have automatic ordering.

- `COLUMN_DEFAULT`

The default value for the column. This is `NULL` if the column has an explicit default of `NULL`, or if the column definition includes no `DEFAULT` clause.

- `IS_NULLABLE`

The column nullability. The value is `YES` if `NULL` values can be stored in the column, `NO` if not.

- `DATA_TYPE`

The column data type.

The `DATA_TYPE` value is the type name only with no other information. The `COLUMN_TYPE` value contains the type name and possibly other information such as the precision or length.

- `CHARACTER_MAXIMUM_LENGTH`

For string columns, the maximum length in characters.

- `CHARACTER_OCTET_LENGTH`

For string columns, the maximum length in bytes.

- `NUMERIC_PRECISION`

For numeric columns, the numeric precision.

- [NUMERIC\\_SCALE](#)

For numeric columns, the numeric scale.

- [DATETIME\\_PRECISION](#)

For temporal columns, the fractional seconds precision.

- [CHARACTER\\_SET\\_NAME](#)

For character string columns, the character set name.

- [COLLATION\\_NAME](#)

For character string columns, the collation name.

- [COLUMN\\_TYPE](#)

The column data type.

The [DATA\\_TYPE](#) value is the type name only with no other information. The [COLUMN\\_TYPE](#) value contains the type name and possibly other information such as the precision or length.

- [COLUMN\\_KEY](#)

Whether the column is indexed:

- If [COLUMN\\_KEY](#) is empty, the column either is not indexed or is indexed only as a secondary column in a multiple-column, nonunique index.
- If [COLUMN\\_KEY](#) is [PRI](#), the column is a [PRIMARY KEY](#) or is one of the columns in a multiple-column [PRIMARY KEY](#).
- If [COLUMN\\_KEY](#) is [UNI](#), the column is the first column of a [UNIQUE](#) index. (A [UNIQUE](#) index permits multiple [NULL](#) values, but you can tell whether the column permits [NULL](#) by checking the [Null](#) column.)
- If [COLUMN\\_KEY](#) is [MUL](#), the column is the first column of a nonunique index in which multiple occurrences of a given value are permitted within the column.

If more than one of the [COLUMN\\_KEY](#) values applies to a given column of a table, [COLUMN\\_KEY](#) displays the one with the highest priority, in the order [PRI](#), [UNI](#), [MUL](#).

A [UNIQUE](#) index may be displayed as [PRI](#) if it cannot contain [NULL](#) values and there is no [PRIMARY KEY](#) in the table. A [UNIQUE](#) index may display as [MUL](#) if several columns form a composite [UNIQUE](#) index; although the combination of the columns is unique, each column can still hold multiple occurrences of a given value.

- [EXTRA](#)

Any additional information that is available about a given column. The value is nonempty in these cases: [auto\\_increment](#) for columns that have the [AUTO\\_INCREMENT](#) attribute; [on update CURRENT\\_TIMESTAMP](#) for [TIMESTAMP](#) or [DATETIME](#) columns that have the [ON UPDATE CURRENT\\_TIMESTAMP](#) attribute.

- [PRIVILEGES](#)

The privileges you have for the column.

- [COLUMN\\_COMMENT](#)

Any comment included in the column definition.

## Notes

- In `SHOW COLUMNS`, the `Type` display includes values from several different `COLUMNS` columns.
- `CHARACTER_OCTET_LENGTH` should be the same as `CHARACTER_MAXIMUM_LENGTH`, except for multibyte character sets.
- `CHARACTER_SET_NAME` can be derived from `COLLATION_NAME`. For example, if you say `SHOW FULL COLUMNS FROM t`, and you see in the `COLLATION_NAME` column a value of `latin1_swedish_ci`, the character set is what is before the first underscore: `latin1`.

Column information is also available from the `SHOW COLUMNS` statement. See [SHOW COLUMNS Statement](#). The following statements are nearly equivalent:

```
SELECT COLUMN_NAME, DATA_TYPE, IS_NULLABLE, COLUMN_DEFAULT
FROM INFORMATION_SCHEMA.COLUMNS
WHERE table_name = 'tbl_name'
[AND table_schema = 'db_name']
[AND column_name LIKE 'wild']
SHOW COLUMNS
FROM tbl_name
[FROM db_name]
[LIKE 'wild']
```

## 4.6 The INFORMATION\_SCHEMA COLUMN\_PRIVILEGES Table

The `COLUMN_PRIVILEGES` table provides information about column privileges. It takes its values from the `mysql.columns_priv` system table.

The `COLUMN_PRIVILEGES` table has these columns:

- `GRANTEE`

The name of the account to which the privilege is granted, in `'user_name'@'host_name'` format.

- `TABLE_CATALOG`

The name of the catalog to which the table containing the column belongs. This value is always `def`.

- `TABLE_SCHEMA`

The name of the schema (database) to which the table containing the column belongs.

- `TABLE_NAME`

The name of the table containing the column.

- `COLUMN_NAME`

The name of the column.

- `PRIVILEGE_TYPE`

The privilege granted. The value can be any privilege that can be granted at the column level; see [GRANT Statement](#). Each row lists a single privilege, so there is one row per column privilege held by the grantee.

In the output from `SHOW FULL COLUMNS`, the privileges are all in one column and in lowercase, for example, `select,insert,update,references`. In `COLUMN_PRIVILEGES`, there is one privilege per row, in uppercase.

- [IS\\_GRANTABLE](#)

**YES** if the user has the [GRANT OPTION](#) privilege, **NO** otherwise. The output does not list [GRANT OPTION](#) as a separate row with [PRIVILEGE\\_TYPE='GRANT OPTION'](#).

## Notes

- [COLUMN\\_PRIVILEGES](#) is a nonstandard [INFORMATION\\_SCHEMA](#) table.

The following statements are *not* equivalent:

```
SELECT ... FROM INFORMATION_SCHEMA.COLUMN_PRIVILEGES
SHOW GRANTS ...
```

## 4.7 The INFORMATION\_SCHEMA ENGINES Table

The [ENGINES](#) table provides information about storage engines. This is particularly useful for checking whether a storage engine is supported, or to see what the default engine is.

The [ENGINES](#) table has these columns:

- [ENGINE](#)

The name of the storage engine.

- [SUPPORT](#)

The server's level of support for the storage engine, as shown in the following table.

Value	Meaning
<a href="#">YES</a>	The engine is supported and is active
<a href="#">DEFAULT</a>	Like <a href="#">YES</a> , plus this is the default engine
<a href="#">NO</a>	The engine is not supported
<a href="#">DISABLED</a>	The engine is supported but has been disabled

A value of [NO](#) means that the server was compiled without support for the engine, so it cannot be enabled at runtime.

A value of [DISABLED](#) occurs either because the server was started with an option that disables the engine, or because not all options required to enable it were given. In the latter case, the error log should contain a reason indicating why the option is disabled. See [The Error Log](#).

You might also see [DISABLED](#) for a storage engine if the server was compiled to support it, but was started with a `--skip-engine_name` option. For the [NDB](#) storage engine, [DISABLED](#) means the server was compiled with support for NDB Cluster, but was not started with the `--ndbcluster` option.

All MySQL servers support [MyISAM](#) tables. It is not possible to disable [MyISAM](#).

- [COMMENT](#)

A brief description of the storage engine.

- [TRANSACTIONS](#)

Whether the storage engine supports transactions.

- [XA](#)

Whether the storage engine supports XA transactions.

- [SAVEPOINTS](#)

Whether the storage engine supports savepoints.

## Notes

- [ENGINES](#) is a nonstandard [INFORMATION\\_SCHEMA](#) table.

Storage engine information is also available from the [SHOW ENGINES](#) statement. See [SHOW ENGINES Statement](#). The following statements are equivalent:

```
SELECT * FROM INFORMATION_SCHEMA.ENGINES
SHOW ENGINES
```

## 4.8 The INFORMATION\_SCHEMA EVENTS Table

The [EVENTS](#) table provides information about Event Manager events, which are discussed in [Using the Event Scheduler](#).

The [EVENTS](#) table has these columns:

- [EVENT\\_CATALOG](#)

The name of the catalog to which the event belongs. This value is always `def`.

- [EVENT\\_SCHEMA](#)

The name of the schema (database) to which the event belongs.

- [EVENT\\_NAME](#)

The name of the event.

- [DEFINER](#)

The account named in the [DEFINER](#) clause (often the user who created the event), in `'user_name'@'host_name'` format.

- [TIME\\_ZONE](#)

The event time zone, which is the time zone used for scheduling the event and that is in effect within the event as it executes. The default value is `SYSTEM`.

- [EVENT\\_BODY](#)

The language used for the statements in the event's [DO](#) clause. The value is always `SQL`.

- [EVENT\\_DEFINITION](#)

The text of the SQL statement making up the event's [DO](#) clause; in other words, the statement executed by this event.

- [EVENT\\_TYPE](#)

The event repetition type, either `ONE TIME` (transient) or `RECURRING` (repeating).

- [EXECUTE\\_AT](#)

For a one-time event, this is the `DATETIME` value specified in the `AT` clause of the `CREATE EVENT` statement used to create the event, or of the last `ALTER EVENT` statement that modified the event. The value shown in this column reflects the addition or subtraction of any `INTERVAL` value included in the event's `AT` clause. For example, if an event is created using `ON SCHEDULE AT`

`CURRENT_TIMESTAMP + '1:6' DAY_HOUR`, and the event was created at 2018-02-09 14:05:30, the value shown in this column would be `'2018-02-10 20:05:30'`. If the event's timing is determined by an `EVERY` clause instead of an `AT` clause (that is, if the event is recurring), the value of this column is `NULL`.

- `INTERVAL_VALUE`

For a recurring event, the number of intervals to wait between event executions. For a transient event, the value is always `NULL`.

- `INTERVAL_FIELD`

The time units used for the interval which a recurring event waits before repeating. For a transient event, the value is always `NULL`.

- `SQL_MODE`

The SQL mode in effect when the event was created or altered, and under which the event executes. For the permitted values, see [Server SQL Modes](#).

- `STARTS`

The start date and time for a recurring event. This is displayed as a `DATETIME` value, and is `NULL` if no start date and time are defined for the event. For a transient event, this column is always `NULL`. For a recurring event whose definition includes a `STARTS` clause, this column contains the corresponding `DATETIME` value. As with the `EXECUTE_AT` column, this value resolves any expressions used. If there is no `STARTS` clause affecting the timing of the event, this column is `NULL`.

- `ENDS`

For a recurring event whose definition includes a `ENDS` clause, this column contains the corresponding `DATETIME` value. As with the `EXECUTE_AT` column, this value resolves any expressions used. If there is no `ENDS` clause affecting the timing of the event, this column is `NULL`.

- `STATUS`

The event status. One of `ENABLED`, `DISABLED`, or `SLAVESIDE_DISABLED`. `SLAVESIDE_DISABLED` indicates that the creation of the event occurred on another MySQL server acting as a replication source and replicated to the current MySQL server which is acting as a replica, but the event is not presently being executed on the replica. For more information, see [Replication of Invoked Features](#) information.

- `ON_COMPLETION`

One of the two values `PRESERVE` or `NOT PRESERVE`.

- `CREATED`

The date and time when the event was created. This is a `TIMESTAMP` value.

- `LAST_ALTERED`

The date and time when the event was last modified. This is a `TIMESTAMP` value. If the event has not been modified since its creation, this value is the same as the `CREATED` value.

- `LAST_EXECUTED`

The date and time when the event last executed. This is a `DATETIME` value. If the event has never executed, this column is `NULL`.

`LAST_EXECUTED` indicates when the event started. As a result, the `ENDS` column is never less than `LAST_EXECUTED`.

- [EVENT\\_COMMENT](#)

The text of the comment, if the event has one. If not, this value is empty.

- [ORIGINATOR](#)

The server ID of the MySQL server on which the event was created; used in replication. This value may be updated by [ALTER EVENT](#) to the server ID of the server on which that statement occurs, if executed on a replication source. The default value is 0.

- [CHARACTER\\_SET\\_CLIENT](#)

The session value of the [character\\_set\\_client](#) system variable when the event was created.

- [COLLATION\\_CONNECTION](#)

The session value of the [collation\\_connection](#) system variable when the event was created.

- [DATABASE\\_COLLATION](#)

The collation of the database with which the event is associated.

## Notes

- [EVENTS](#) is a nonstandard [INFORMATION\\_SCHEMA](#) table.
- Times in the [EVENTS](#) table are displayed using the event time zone, the current session time zone, or UTC, as described in [Event Metadata](#).
- For more information about [SLAVESIDE\\_DISABLED](#) and the [ORIGINATOR](#) column, see [Replication of Invoked Features](#).

## Example

Suppose that the user 'jon'@'ghidora' creates an event named `e_daily`, and then modifies it a few minutes later using an [ALTER EVENT](#) statement, as shown here:

```
DELIMITER |
CREATE EVENT e_daily
  ON SCHEDULE
    EVERY 1 DAY
  COMMENT 'Saves total number of sessions then clears the table each day'
  DO
  BEGIN
    INSERT INTO site_activity.totals (time, total)
      SELECT CURRENT_TIMESTAMP, COUNT(*)
        FROM site_activity.sessions;
    DELETE FROM site_activity.sessions;
  END |
DELIMITER ;
ALTER EVENT e_daily
  ENABLE;
```

(Note that comments can span multiple lines.)

This user can then run the following [SELECT](#) statement, and obtain the output shown:

```
mysql> SELECT * FROM INFORMATION_SCHEMA.EVENTS
  WHERE EVENT_NAME = 'e_daily'
  AND EVENT_SCHEMA = 'myschema'\G
***** 1. row *****
EVENT_CATALOG: def
EVENT_SCHEMA: myschema
EVENT_NAME: e_daily
DEFINER: jon@ghidora
TIME_ZONE: SYSTEM
EVENT_BODY: SQL
```

```

EVENT_DEFINITION: BEGIN
  INSERT INTO site_activity.totals (time, total)
  SELECT CURRENT_TIMESTAMP, COUNT(*)
  FROM site_activity.sessions;
  DELETE FROM site_activity.sessions;
END
EVENT_TYPE: RECURRING
EXECUTE_AT: NULL
INTERVAL_VALUE: 1
INTERVAL_FIELD: DAY
SQL_MODE: NO_ENGINE_SUBSTITUTION
STARTS: 2018-08-08 11:06:34
ENDS: NULL
STATUS: ENABLED
ON_COMPLETION: NOT PRESERVE
CREATED: 2018-08-08 11:06:34
LAST_ALTERED: 2018-08-08 11:06:34
LAST_EXECUTED: 2018-08-08 16:06:34
EVENT_COMMENT: Saves total number of sessions then clears the
table each day
ORIGINATOR: 1
CHARACTER_SET_CLIENT: utf8
COLLATION_CONNECTION: utf8_general_ci
DATABASE_COLLATION: latin1_swedish_ci

```

Event information is also available from the [SHOW EVENTS](#) statement. See [SHOW EVENTS Statement](#). The following statements are equivalent:

```

SELECT
  EVENT_SCHEMA, EVENT_NAME, DEFINER, TIME_ZONE, EVENT_TYPE, EXECUTE_AT,
  INTERVAL_VALUE, INTERVAL_FIELD, STARTS, ENDS, STATUS, ORIGINATOR,
  CHARACTER_SET_CLIENT, COLLATION_CONNECTION, DATABASE_COLLATION
FROM INFORMATION_SCHEMA.EVENTS
WHERE table_schema = 'db_name'
[AND column_name LIKE 'wild']
SHOW EVENTS
[FROM db_name]
[LIKE 'wild']

```

## 4.9 The INFORMATION\_SCHEMA GLOBAL\_STATUS and SESSION\_STATUS Tables

The [GLOBAL\\_STATUS](#) and [SESSION\\_STATUS](#) tables provide information about server status variables. Their contents correspond to the information produced by the [SHOW GLOBAL STATUS](#) and [SHOW SESSION STATUS](#) statements (see [SHOW STATUS Statement](#)).

### Notes

- The [VARIABLE\\_VALUE](#) column for each of these tables is defined as [VARCHAR\(1024\)](#).

## 4.10 The INFORMATION\_SCHEMA GLOBAL\_VARIABLES and SESSION\_VARIABLES Tables

The [GLOBAL\\_VARIABLES](#) and [SESSION\\_VARIABLES](#) tables provide information about server status variables. Their contents correspond to the information produced by the [SHOW GLOBAL VARIABLES](#) and [SHOW SESSION VARIABLES](#) statements (see [SHOW VARIABLES Statement](#)).

### Notes

- The [VARIABLE\\_VALUE](#) column for each of these tables is defined as [VARCHAR\(1024\)](#). For variables with very long values that are not completely displayed, use [SELECT](#) as a workaround. For example:

```
SELECT @@GLOBAL.innodb_data_file_path;
```

## 4.11 The INFORMATION\_SCHEMA KEY\_COLUMN\_USAGE Table

The `KEY_COLUMN_USAGE` table describes which key columns have constraints.

The `KEY_COLUMN_USAGE` table has these columns:

- `CONSTRAINT_CATALOG`

The name of the catalog to which the constraint belongs. This value is always `def`.

- `CONSTRAINT_SCHEMA`

The name of the schema (database) to which the constraint belongs.

- `CONSTRAINT_NAME`

The name of the constraint.

- `TABLE_CATALOG`

The name of the catalog to which the table belongs. This value is always `def`.

- `TABLE_SCHEMA`

The name of the schema (database) to which the table belongs.

- `TABLE_NAME`

The name of the table that has the constraint.

- `COLUMN_NAME`

The name of the column that has the constraint.

If the constraint is a foreign key, then this is the column of the foreign key, not the column that the foreign key references.

- `ORDINAL_POSITION`

The column's position within the constraint, not the column's position within the table. Column positions are numbered beginning with 1.

- `POSITION_IN_UNIQUE_CONSTRAINT`

`NULL` for unique and primary-key constraints. For foreign-key constraints, this column is the ordinal position in key of the table that is being referenced.

- `REFERENCED_TABLE_SCHEMA`

The name of the schema (database) referenced by the constraint.

- `REFERENCED_TABLE_NAME`

The name of the table referenced by the constraint.

- `REFERENCED_COLUMN_NAME`

The name of the column referenced by the constraint.

Suppose that there are two tables name `t1` and `t3` that have the following definitions:

```
CREATE TABLE t1
(
```

```

s1 INT,
s2 INT,
s3 INT,
PRIMARY KEY(s3)
) ENGINE=InnoDB;
CREATE TABLE t3
(
s1 INT,
s2 INT,
s3 INT,
KEY(s1),
CONSTRAINT CO FOREIGN KEY (s2) REFERENCES t1(s3)
) ENGINE=InnoDB;

```

For those two tables, the `KEY_COLUMN_USAGE` table has two rows:

- One row with `CONSTRAINT_NAME = 'PRIMARY'`, `TABLE_NAME = 't1'`, `COLUMN_NAME = 's3'`, `ORDINAL_POSITION = 1`, `POSITION_IN_UNIQUE_CONSTRAINT = NULL`.
- One row with `CONSTRAINT_NAME = 'CO'`, `TABLE_NAME = 't3'`, `COLUMN_NAME = 's2'`, `ORDINAL_POSITION = 1`, `POSITION_IN_UNIQUE_CONSTRAINT = 1`.

## 4.12 The INFORMATION\_SCHEMA OPTIMIZER\_TRACE Table

The `OPTIMIZER_TRACE` table provides information produced by the optimizer tracing capability for traced statements. To enable tracking, use the `optimizer_trace` system variable. For details, see [MySQL Internals: Tracing the Optimizer](#).

The `OPTIMIZER_TRACE` table has these columns:

- `QUERY`

The text of the traced statement.

- `TRACE`

The trace, in `JSON` format.

- `MISSING_BYTES_BEYOND_MAX_MEM_SIZE`

Each remembered trace is a string that is extended as optimization progresses and appends data to it. The `optimizer_trace_max_mem_size` variable sets a limit on the total amount of memory used by all currently remembered traces. If this limit is reached, the current trace is not extended (and thus is incomplete), and the `MISSING_BYTES_BEYOND_MAX_MEM_SIZE` column shows the number of bytes missing from the trace.

- `INSUFFICIENT_PRIVILEGES`

If a traced query uses views or stored routines that have `SQL SECURITY` with a value of `DEFINER`, it may be that a user other than the definer is denied from seeing the trace of the query. In that case, the trace is shown as empty and `INSUFFICIENT_PRIVILEGES` has a value of 1. Otherwise, the value is 0.

## 4.13 The INFORMATION\_SCHEMA PARAMETERS Table

The `PARAMETERS` table provides information about parameters for stored routines (stored procedures and stored functions), and about return values for stored functions. The `PARAMETERS` table does not include built-in (native) functions or loadable functions. Parameter information is similar to the contents of the `param_list` column in the `mysql.proc` table.

The `PARAMETERS` table has these columns:

- `SPECIFIC_CATALOG`

The name of the catalog to which the routine containing the parameter belongs. This value is always `def`.

- `SPECIFIC_SCHEMA`

The name of the schema (database) to which the routine containing the parameter belongs.

- `SPECIFIC_NAME`

The name of the routine containing the parameter.

- `ORDINAL_POSITION`

For successive parameters of a stored procedure or function, the `ORDINAL_POSITION` values are 1, 2, 3, and so forth. For a stored function, there is also a row that applies to the function return value (as described by the `RETURNS` clause). The return value is not a true parameter, so the row that describes it has these unique characteristics:

- The `ORDINAL_POSITION` value is 0.
- The `PARAMETER_NAME` and `PARAMETER_MODE` values are `NULL` because the return value has no name and the mode does not apply.

- `PARAMETER_MODE`

The mode of the parameter. This value is one of `IN`, `OUT`, or `INOUT`. For a stored function return value, this value is `NULL`.

- `PARAMETER_NAME`

The name of the parameter. For a stored function return value, this value is `NULL`.

- `DATA_TYPE`

The parameter data type.

The `DATA_TYPE` value is the type name only with no other information. The `DTD_IDENTIFIER` value contains the type name and possibly other information such as the precision or length.

- `CHARACTER_MAXIMUM_LENGTH`

For string parameters, the maximum length in characters.

- `CHARACTER_OCTET_LENGTH`

For string parameters, the maximum length in bytes.

- `NUMERIC_PRECISION`

For numeric parameters, the numeric precision.

- `NUMERIC_SCALE`

For numeric parameters, the numeric scale.

- `DATETIME_PRECISION`

For temporal parameters, the fractional seconds precision.

- `CHARACTER_SET_NAME`

For character string parameters, the character set name.

- `COLLATION_NAME`

For character string parameters, the collation name.

- [DTD\\_IDENTIFIER](#)

The parameter data type.

The [DATA\\_TYPE](#) value is the type name only with no other information. The [DTD\\_IDENTIFIER](#) value contains the type name and possibly other information such as the precision or length.

- [ROUTINE\\_TYPE](#)

[PROCEDURE](#) for stored procedures, [FUNCTION](#) for stored functions.

## 4.14 The INFORMATION\_SCHEMA PARTITIONS Table

The [PARTITIONS](#) table provides information about table partitions. Each row in this table corresponds to an individual partition or subpartition of a partitioned table. For more information about partitioning tables, see [Partitioning](#).

The [PARTITIONS](#) table has these columns:

- [TABLE\\_CATALOG](#)

The name of the catalog to which the table belongs. This value is always `def`.

- [TABLE\\_SCHEMA](#)

The name of the database to which the table belongs.

- [TABLE\\_NAME](#)

The name of the table containing the partition.

- [PARTITION\\_NAME](#)

The name of the partition.

- [SUBPARTITION\\_NAME](#)

If the [PARTITIONS](#) table row represents a subpartition, the name of subpartition; otherwise `NULL`.

- [PARTITION\\_ORDINAL\\_POSITION](#)

All partitions are indexed in the same order as they are defined, with `1` being the number assigned to the first partition. The indexing can change as partitions are added, dropped, and reorganized; the number shown in this column reflects the current order, taking into account any indexing changes.

- [SUBPARTITION\\_ORDINAL\\_POSITION](#)

Subpartitions within a given partition are also indexed and reindexed in the same manner as partitions are indexed within a table.

- [PARTITION\\_METHOD](#)

One of the values [RANGE](#), [LIST](#), [HASH](#), [LINEAR HASH](#), [KEY](#), or [LINEAR KEY](#); that is, one of the available partitioning types as discussed in [Partitioning Types](#).

- [SUBPARTITION\\_METHOD](#)

One of the values [HASH](#), [LINEAR HASH](#), [KEY](#), or [LINEAR KEY](#); that is, one of the available subpartitioning types as discussed in [Subpartitioning](#).

- [PARTITION\\_EXPRESSION](#)

The expression for the partitioning function used in the `CREATE TABLE` or `ALTER TABLE` statement that created the table's current partitioning scheme.

For example, consider a partitioned table created in the `test` database using this statement:

```
CREATE TABLE tp (
  c1 INT,
  c2 INT,
  c3 VARCHAR(25)
)
PARTITION BY HASH(c1 + c2)
PARTITIONS 4;
```

The `PARTITION_EXPRESSION` column in a `PARTITIONS` table row for a partition from this table displays `c1 + c2`, as shown here:

```
mysql> SELECT DISTINCT PARTITION_EXPRESSION
FROM INFORMATION_SCHEMA.PARTITIONS
WHERE TABLE_NAME='tp' AND TABLE_SCHEMA='test';
+-----+
| PARTITION_EXPRESSION |
+-----+
| c1 + c2              |
+-----+
```

- `SUBPARTITION_EXPRESSION`

This works in the same fashion for the subpartitioning expression that defines the subpartitioning for a table as `PARTITION_EXPRESSION` does for the partitioning expression used to define a table's partitioning.

If the table has no subpartitions, this column is `NULL`.

- `PARTITION_DESCRIPTION`

This column is used for `RANGE` and `LIST` partitions. For a `RANGE` partition, it contains the value set in the partition's `VALUES LESS THAN` clause, which can be either an integer or `MAXVALUE`. For a `LIST` partition, this column contains the values defined in the partition's `VALUES IN` clause, which is a list of comma-separated integer values.

For partitions whose `PARTITION_METHOD` is other than `RANGE` or `LIST`, this column is always `NULL`.

- `TABLE_ROWS`

The number of table rows in the partition.

For partitioned `InnoDB` tables, the row count given in the `TABLE_ROWS` column is only an estimated value used in SQL optimization, and may not always be exact.

For `NDB` tables, you can also obtain this information using the `ndb_desc` utility.

- `AVG_ROW_LENGTH`

The average length of the rows stored in this partition or subpartition, in bytes. This is the same as `DATA_LENGTH` divided by `TABLE_ROWS`.

For `NDB` tables, you can also obtain this information using the `ndb_desc` utility.

- `DATA_LENGTH`

The total length of all rows stored in this partition or subpartition, in bytes; that is, the total number of bytes stored in the partition or subpartition.

For `NDB` tables, you can also obtain this information using the `ndb_desc` utility.

- [MAX\\_DATA\\_LENGTH](#)

The maximum number of bytes that can be stored in this partition or subpartition.

For [NDB](#) tables, you can also obtain this information using the [ndb\\_desc](#) utility.

- [INDEX\\_LENGTH](#)

The length of the index file for this partition or subpartition, in bytes.

For partitions of [NDB](#) tables, whether the tables use implicit or explicit partitioning, the [INDEX\\_LENGTH](#) column value is always 0. However, you can obtain equivalent information using the [ndb\\_desc](#) utility.

- [DATA\\_FREE](#)

The number of bytes allocated to the partition or subpartition but not used.

For [NDB](#) tables, you can also obtain this information using the [ndb\\_desc](#) utility.

- [CREATE\\_TIME](#)

The time that the partition or subpartition was created.

Prior to MySQL 5.6.25, for partitioned [InnoDB](#) tables, this column was always [NULL](#). The correct creation time is shown in MySQL 5.6.25 and later. (Bug #17299181, Bug #69990)

- [UPDATE\\_TIME](#)

The time that the partition or subpartition was last modified.

For partitioned [InnoDB](#) tables, the value is always [NULL](#).

- [CHECK\\_TIME](#)

The last time that the table to which this partition or subpartition belongs was checked.

For partitioned [InnoDB](#) tables, this column is always [NULL](#).

- [CHECKSUM](#)

The checksum value, if any; otherwise [NULL](#).

- [PARTITION\\_COMMENT](#)

The text of the comment, if the partition has one. If not, this value is empty.

The maximum length for a partition comment is defined as 1024 characters, and the display width of the [PARTITION\\_COMMENT](#) column is also 1024, characters to match this limit.

- [NODEGROUP](#)

This is the nodegroup to which the partition belongs. This is relevant only to [NDB Cluster](#) tables; otherwise, the value is always 0.

- [TABLESPACE\\_NAME](#)

The name of the tablespace to which the partition belongs. The value is always [DEFAULT](#), unless the table uses the [NDB](#) storage engine (see the *Notes* at the end of this section).

## Notes

- [PARTITIONS](#) is a nonstandard [INFORMATION\\_SCHEMA](#) table.

- A table using any storage engine other than `NDB` and which is not partitioned has one row in the `PARTITIONS` table. However, the values of the `PARTITION_NAME`, `SUBPARTITION_NAME`, `PARTITION_ORDINAL_POSITION`, `SUBPARTITION_ORDINAL_POSITION`, `PARTITION_METHOD`, `SUBPARTITION_METHOD`, `PARTITION_EXPRESSION`, `SUBPARTITION_EXPRESSION`, and `PARTITION_DESCRIPTION` columns are all `NULL`. Also, the `PARTITION_COMMENT` column in this case is blank.
- An `NDB` table which is not explicitly partitioned has one row in the `PARTITIONS` table for each data node in the `NDB` cluster. For each such row:
  - The `SUBPARTITION_NAME`, `SUBPARTITION_ORDINAL_POSITION`, `SUBPARTITION_METHOD`, `PARTITION_EXPRESSION`, `SUBPARTITION_EXPRESSION`, `CREATE_TIME`, `UPDATE_TIME`, `CHECK_TIME`, `CHECKSUM`, and `TABLESPACE_NAME` columns are all `NULL`.
  - The `PARTITION_METHOD` is always `KEY`.
  - The `NODEGROUP` column is `default`.
  - The `PARTITION_EXPRESSION` and `PARTITION_COMMENT` columns are empty.

## 4.15 The INFORMATION\_SCHEMA PLUGINS Table

The `PLUGINS` table provides information about server plugins.

The `PLUGINS` table has these columns:

- `PLUGIN_NAME`

The name used to refer to the plugin in statements such as `INSTALL PLUGIN` and `UNINSTALL PLUGIN`.

- `PLUGIN_VERSION`

The version from the plugin's general type descriptor.

- `PLUGIN_STATUS`

The plugin status, one of `ACTIVE`, `INACTIVE`, `DISABLED`, or `DELETED`.

- `PLUGIN_TYPE`

The type of plugin, such as `STORAGE ENGINE`, `INFORMATION_SCHEMA`, or `AUTHENTICATION`.

- `PLUGIN_TYPE_VERSION`

The version from the plugin's type-specific descriptor.

- `PLUGIN_LIBRARY`

The name of the plugin shared library file. This is the name used to refer to the plugin file in statements such as `INSTALL PLUGIN` and `UNINSTALL PLUGIN`. This file is located in the directory named by the `plugin_dir` system variable. If the library name is `NULL`, the plugin is compiled in and cannot be uninstalled with `UNINSTALL PLUGIN`.

- `PLUGIN_LIBRARY_VERSION`

The plugin API interface version.

- `PLUGIN_AUTHOR`

The plugin author.

- `PLUGIN_DESCRIPTION`

A short description of the plugin.

- `PLUGIN_LICENSE`

How the plugin is licensed (for example, [GPL](#)).

- `LOAD_OPTION`

How the plugin was loaded. The value is `OFF`, `ON`, `FORCE`, or `FORCE_PLUS_PERMANENT`. See [Installing and Uninstalling Plugins](#).

## Notes

- `PLUGINS` is a nonstandard `INFORMATION_SCHEMA` table.
- For plugins installed with `INSTALL PLUGIN`, the `PLUGIN_NAME` and `PLUGIN_LIBRARY` values are also registered in the `mysql.plugin` table.
- For information about plugin data structures that form the basis of the information in the `PLUGINS` table, see [The MySQL Plugin API](#).

Plugin information is also available from the `SHOW PLUGINS` statement. See [SHOW PLUGINS Statement](#). These statements are equivalent:

```
SELECT
  PLUGIN_NAME, PLUGIN_STATUS, PLUGIN_TYPE,
  PLUGIN_LIBRARY, PLUGIN_LICENSE
FROM INFORMATION_SCHEMA.PLUGINS;
SHOW PLUGINS;
```

## 4.16 The INFORMATION\_SCHEMA PROCESSLIST Table

The MySQL process list indicates the operations currently being performed by the set of threads executing within the server. The `PROCESSLIST` table is one source of process information. For a comparison of this table with other sources, see [Sources of Process Information](#).

The `PROCESSLIST` table has these columns:

- `ID`

The connection identifier. This is the same value displayed in the `Id` column of the `SHOW PROCESSLIST` statement, displayed in the `PROCESSLIST_ID` column of the Performance Schema `threads` table, and returned by the `CONNECTION_ID()` function within the thread.

- `USER`

The MySQL user who issued the statement. A value of `system user` refers to a nonclient thread spawned by the server to handle tasks internally, for example, a delayed-row handler thread or an I/O or SQL thread used on replica hosts. For `system user`, there is no host specified in the `Host` column. `unauthenticated user` refers to a thread that has become associated with a client connection but for which authentication of the client user has not yet occurred. `event_scheduler` refers to the thread that monitors scheduled events (see [Using the Event Scheduler](#)).

- `HOST`

The host name of the client issuing the statement (except for `system user`, for which there is no host). The host name for TCP/IP connections is reported in `host_name:client_port` format to make it easier to determine which client is doing what.

- `DB`

The default database for the thread, or `NULL` if none has been selected.

- `COMMAND`

The type of command the thread is executing on behalf of the client, or `Sleep` if the session is idle. For descriptions of thread commands, see [Examining Server Thread \(Process\) Information](#). The value of this column corresponds to the `COM_xxx` commands of the client/server protocol and `Com_xxx` status variables. See [Server Status Variables](#).

- `TIME`

The time in seconds that the thread has been in its current state. For a replica SQL thread, the value is the number of seconds between the timestamp of the last replicated event and the real time of the replica host. See [Replication Threads](#).

- `STATE`

An action, event, or state that indicates what the thread is doing. For descriptions of `STATE` values, see [Examining Server Thread \(Process\) Information](#).

Most states correspond to very quick operations. If a thread stays in a given state for many seconds, there might be a problem that needs to be investigated.

- `INFO`

The statement the thread is executing, or `NULL` if it is executing no statement. The statement might be the one sent to the server, or an innermost statement if the statement executes other statements. For example, if a `CALL` statement executes a stored procedure that is executing a `SELECT` statement, the `INFO` value shows the `SELECT` statement.

## Notes

- `PROCESSLIST` is a nonstandard `INFORMATION_SCHEMA` table.
- Like the output from the `SHOW PROCESSLIST` statement, the `PROCESSLIST` table provides information about all threads, even those belonging to other users, if you have the `PROCESS` privilege. Otherwise (without the `PROCESS` privilege), nonanonymous users have access to information about their own threads but not threads for other users, and anonymous users have no access to thread information.
- If an SQL statement refers to the `PROCESSLIST` table, MySQL populates the entire table once, when statement execution begins, so there is read consistency during the statement. There is no read consistency for a multi-statement transaction.

The following statements are equivalent:

```
SELECT * FROM INFORMATION_SCHEMA.PROCESSLIST
SHOW FULL PROCESSLIST
```

## 4.17 The INFORMATION\_SCHEMA PROFILING Table

The `PROFILING` table provides statement profiling information. Its contents correspond to the information produced by the `SHOW PROFILE` and `SHOW PROFILES` statements (see [SHOW PROFILE Statement](#)). The table is empty unless the `profiling` session variable is set to 1.

The `PROFILING` table has these columns:

- `QUERY_ID`

A numeric statement identifier.

- [SEQ](#)  
A sequence number indicating the display order for rows with the same [QUERY\\_ID](#) value.
- [STATE](#)  
The profiling state to which the row measurements apply.
- [DURATION](#)  
How long statement execution remained in the given state, in seconds.
- [CPU\\_USER](#), [CPU\\_SYSTEM](#)  
User and system CPU use, in seconds.
- [CONTEXT\\_VOLUNTARY](#), [CONTEXT\\_INVOLUNTARY](#)  
How many voluntary and involuntary context switches occurred.
- [BLOCK\\_OPS\\_IN](#), [BLOCK\\_OPS\\_OUT](#)  
The number of block input and output operations.
- [MESSAGES\\_SENT](#), [MESSAGES\\_RECEIVED](#)  
The number of communication messages sent and received.
- [PAGE\\_FAULTS\\_MAJOR](#), [PAGE\\_FAULTS\\_MINOR](#)  
The number of major and minor page faults.
- [SWAPS](#)  
How many swaps occurred.
- [SOURCE\\_FUNCTION](#), [SOURCE\\_FILE](#), and [SOURCE\\_LINE](#)  
Information indicating where in the source code the profiled state executes.

## Notes

- [PROFILING](#) is a nonstandard [INFORMATION\\_SCHEMA](#) table.

Profiling information is also available from the [SHOW PROFILE](#) and [SHOW PROFILES](#) statements. See [SHOW PROFILE Statement](#). For example, the following queries are equivalent:

```
SHOW PROFILE FOR QUERY 2;
SELECT STATE, FORMAT(DURATION, 6) AS DURATION
FROM INFORMATION_SCHEMA.PROFILING
WHERE QUERY_ID = 2 ORDER BY SEQ;
```

## 4.18 The INFORMATION\_SCHEMA REFERENTIAL\_CONSTRAINTS Table

The [REFERENTIAL\\_CONSTRAINTS](#) table provides information about foreign keys.

The [REFERENTIAL\\_CONSTRAINTS](#) table has these columns:

- [CONSTRAINT\\_CATALOG](#)  
The name of the catalog to which the constraint belongs. This value is always `def`.

- [CONSTRAINT\\_SCHEMA](#)  
The name of the schema (database) to which the constraint belongs.
- [CONSTRAINT\\_NAME](#)  
The name of the constraint.
- [UNIQUE\\_CONSTRAINT\\_CATALOG](#)  
The name of the catalog containing the unique constraint that the constraint references. This value is always `def`.
- [UNIQUE\\_CONSTRAINT\\_SCHEMA](#)  
The name of the schema (database) containing the unique constraint that the constraint references.
- [UNIQUE\\_CONSTRAINT\\_NAME](#)  
The name of the unique constraint that the constraint references.
- [MATCH\\_OPTION](#)  
The value of the constraint `MATCH` attribute. The only valid value at this time is `NONE`.
- [UPDATE\\_RULE](#)  
The value of the constraint `ON UPDATE` attribute. The possible values are `CASCADE`, `SET NULL`, `SET DEFAULT`, `RESTRICT`, `NO ACTION`.
- [DELETE\\_RULE](#)  
The value of the constraint `ON DELETE` attribute. The possible values are `CASCADE`, `SET NULL`, `SET DEFAULT`, `RESTRICT`, `NO ACTION`.
- [TABLE\\_NAME](#)  
The name of the table. This value is the same as in the [TABLE\\_CONSTRAINTS](#) table.
- [REFERENCED\\_TABLE\\_NAME](#)  
The name of the table referenced by the constraint.

## 4.19 The INFORMATION\_SCHEMA ROUTINES Table

The [ROUTINES](#) table provides information about stored routines (stored procedures and stored functions). The [ROUTINES](#) table does not include built-in (native) functions or loadable functions.

The column named “`mysql.proc Name`” indicates the `mysql.proc` table column that corresponds to the [INFORMATION\\_SCHEMA ROUTINES](#) table column, if any.

The [ROUTINES](#) table has these columns:

- [SPECIFIC\\_NAME](#)  
The name of the routine.
- [ROUTINE\\_CATALOG](#)  
The name of the catalog to which the routine belongs. This value is always `def`.
- [ROUTINE\\_SCHEMA](#)

The name of the schema (database) to which the routine belongs.

- [ROUTINE\\_NAME](#)

The name of the routine.

- [ROUTINE\\_TYPE](#)

[PROCEDURE](#) for stored procedures, [FUNCTION](#) for stored functions.

- [DATA\\_TYPE](#)

If the routine is a stored function, the return value data type. If the routine is a stored procedure, this value is empty.

The [DATA\\_TYPE](#) value is the type name only with no other information. The [DTD\\_IDENTIFIER](#) value contains the type name and possibly other information such as the precision or length.

- [CHARACTER\\_MAXIMUM\\_LENGTH](#)

For stored function string return values, the maximum length in characters. If the routine is a stored procedure, this value is [NULL](#).

- [CHARACTER\\_OCTET\\_LENGTH](#)

For stored function string return values, the maximum length in bytes. If the routine is a stored procedure, this value is [NULL](#).

- [NUMERIC\\_PRECISION](#)

For stored function numeric return values, the numeric precision. If the routine is a stored procedure, this value is [NULL](#).

- [NUMERIC\\_SCALE](#)

For stored function numeric return values, the numeric scale. If the routine is a stored procedure, this value is [NULL](#).

- [DATETIME\\_PRECISION](#)

For stored function temporal return values, the fractional seconds precision. If the routine is a stored procedure, this value is [NULL](#).

- [CHARACTER\\_SET\\_NAME](#)

For stored function character string return values, the character set name. If the routine is a stored procedure, this value is [NULL](#).

- [COLLATION\\_NAME](#)

For stored function character string return values, the collation name. If the routine is a stored procedure, this value is [NULL](#).

- [DTD\\_IDENTIFIER](#)

If the routine is a stored function, the return value data type. If the routine is a stored procedure, this value is empty.

The [DATA\\_TYPE](#) value is the type name only with no other information. The [DTD\\_IDENTIFIER](#) value contains the type name and possibly other information such as the precision or length.

- [ROUTINE\\_BODY](#)

The language used for the routine definition. This value is always `SQL`.

- `ROUTINE_DEFINITION`

The text of the SQL statement executed by the routine.

- `EXTERNAL_NAME`

This value is always `NULL`.

- `EXTERNAL_LANGUAGE`

The language of the stored routine. MySQL calculates `EXTERNAL_LANGUAGE` thus:

- If `mysql.proc.language = 'SQL'`, `EXTERNAL_LANGUAGE` is `NULL`
- Otherwise, `EXTERNAL_LANGUAGE` is what is in `mysql.proc.language`. However, we do not have external languages yet, so it is always `NULL`.

- `PARAMETER_STYLE`

This value is always `SQL`.

- `IS_DETERMINISTIC`

`YES` or `NO`, depending on whether the routine is defined with the `DETERMINISTIC` characteristic.

- `SQL_DATA_ACCESS`

The data access characteristic for the routine. The value is one of `CONTAINS SQL`, `NO SQL`, `READS SQL DATA`, or `MODIFIES SQL DATA`.

- `SQL_PATH`

This value is always `NULL`.

- `SECURITY_TYPE`

The routine `SQL SECURITY` characteristic. The value is one of `DEFINER` or `INVOKER`.

- `CREATED`

The date and time when the routine was created. This is a `TIMESTAMP` value.

- `LAST_ALTERED`

The date and time when the routine was last modified. This is a `TIMESTAMP` value. If the routine has not been modified since its creation, this value is the same as the `CREATED` value.

- `SQL_MODE`

The SQL mode in effect when the routine was created or altered, and under which the routine executes. For the permitted values, see [Server SQL Modes](#).

- `ROUTINE_COMMENT`

The text of the comment, if the routine has one. If not, this value is empty.

- `DEFINER`

The account named in the `DEFINER` clause (often the user who created the routine), in `'user_name'@'host_name'` format.

- `CHARACTER_SET_CLIENT`

The session value of the `character_set_client` system variable when the routine was created.

- `COLLATION_CONNECTION`

The session value of the `collation_connection` system variable when the routine was created.

- `DATABASE_COLLATION`

The collation of the database with which the routine is associated.

## Notes

- To see information about a routine, you must be the user named in the routine `DEFINER` clause or have `SELECT` access to the `mysql.proc` table. If you do not have privileges for the routine itself, the value displayed for the `ROUTINE_DEFINITION` column is `NULL`.
- Information about stored function return values is also available in the `PARAMETERS` table. The return value row for a stored function can be identified as the row that has an `ORDINAL_POSITION` value of 0.

## 4.20 The INFORMATION\_SCHEMA SCHEMATA Table

A schema is a database, so the `SCHEMATA` table provides information about databases.

The `SCHEMATA` table has these columns:

- `CATALOG_NAME`

The name of the catalog to which the schema belongs. This value is always `def`.

- `SCHEMA_NAME`

The name of the schema.

- `DEFAULT_CHARACTER_SET_NAME`

The schema default character set.

- `DEFAULT_COLLATION_NAME`

The schema default collation.

- `SQL_PATH`

This value is always `NULL`.

Schema names are also available from the `SHOW DATABASES` statement. See [SHOW DATABASES Statement](#). The following statements are equivalent:

```
SELECT SCHEMA_NAME AS `Database`
FROM INFORMATION_SCHEMA.SCHEMATA
[WHERE SCHEMA_NAME LIKE 'wild']
SHOW DATABASES
[LIKE 'wild']
```

You see only those databases for which you have some kind of privilege, unless you have the global `SHOW DATABASES` privilege.

### Caution

Because a global privilege is considered a privilege for all databases, *any* global privilege enables a user to see all database names with `SHOW DATABASES` or by examining the `INFORMATION_SCHEMA.SCHEMATA` table.

## 4.21 The INFORMATION\_SCHEMA SCHEMA\_PRIVILEGES Table

The `SCHEMA_PRIVILEGES` table provides information about schema (database) privileges. It takes its values from the `mysql.db` system table.

The `SCHEMA_PRIVILEGES` table has these columns:

- `GRANTEE`

The name of the account to which the privilege is granted, in `'user_name'@'host_name'` format.

- `TABLE_CATALOG`

The name of the catalog to which the schema belongs. This value is always `def`.

- `TABLE_SCHEMA`

The name of the schema.

- `PRIVILEGE_TYPE`

The privilege granted. The value can be any privilege that can be granted at the schema level; see [GRANT Statement](#). Each row lists a single privilege, so there is one row per schema privilege held by the grantee.

- `IS_GRANTABLE`

`YES` if the user has the `GRANT OPTION` privilege, `NO` otherwise. The output does not list `GRANT OPTION` as a separate row with `PRIVILEGE_TYPE='GRANT OPTION'`.

### Notes

- `SCHEMA_PRIVILEGES` is a nonstandard `INFORMATION_SCHEMA` table.

The following statements are *not* equivalent:

```
SELECT ... FROM INFORMATION_SCHEMA.SCHEMA_PRIVILEGES
SHOW GRANTS ...
```

## 4.22 The INFORMATION\_SCHEMA STATISTICS Table

The `STATISTICS` table provides information about table indexes.

The `STATISTICS` table has these columns:

- `TABLE_CATALOG`

The name of the catalog to which the table containing the index belongs. This value is always `def`.

- `TABLE_SCHEMA`

The name of the schema (database) to which the table containing the index belongs.

- `TABLE_NAME`

The name of the table containing the index.

- `NON_UNIQUE`

0 if the index cannot contain duplicates, 1 if it can.

- `INDEX_SCHEMA`

The name of the schema (database) to which the index belongs.

- [INDEX\\_NAME](#)

The name of the index. If the index is the primary key, the name is always [PRIMARY](#).

- [SEQ\\_IN\\_INDEX](#)

The column sequence number in the index, starting with 1.

- [COLUMN\\_NAME](#)

The column name. See also the description for the [EXPRESSION](#) column.

- [COLLATION](#)

How the column is sorted in the index. This can have values [A](#) (ascending), [D](#) (descending), or [NULL](#) (not sorted).

- [CARDINALITY](#)

An estimate of the number of unique values in the index. To update this number, run [ANALYZE TABLE](#) or (for [MyISAM](#) tables) `myisamchk -a`.

[CARDINALITY](#) is counted based on statistics stored as integers, so the value is not necessarily exact even for small tables. The higher the cardinality, the greater the chance that MySQL uses the index when doing joins.

- [SUB\\_PART](#)

The index prefix. That is, the number of indexed characters if the column is only partly indexed, [NULL](#) if the entire column is indexed.

**Note**

Prefix *limits* are measured in bytes. However, prefix *lengths* for index specifications in [CREATE TABLE](#), [ALTER TABLE](#), and [CREATE INDEX](#) statements are interpreted as number of characters for nonbinary string types ([CHAR](#), [VARCHAR](#), [TEXT](#)) and number of bytes for binary string types ([BINARY](#), [VARBINARY](#), [BLOB](#)). Take this into account when specifying a prefix length for a nonbinary string column that uses a multibyte character set.

For additional information about index prefixes, see [Column Indexes](#), and [CREATE INDEX Statement](#).

- [PACKED](#)

Indicates how the key is packed. [NULL](#) if it is not.

- [NULLABLE](#)

Contains [YES](#) if the column may contain [NULL](#) values and `' '` if not.

- [INDEX\\_TYPE](#)

The index method used ([BTREE](#), [FULLTEXT](#), [HASH](#), [RTREE](#)).

- [COMMENT](#)

Information about the index not described in its own column, such as [disabled](#) if the index is disabled.

- [INDEX\\_COMMENT](#)

Any comment provided for the index with a `COMMENT` attribute when the index was created.

## Notes

- There is no standard `INFORMATION_SCHEMA` table for indexes. The MySQL column list is similar to what SQL Server 2000 returns for `sp_statistics`, except that `QUALIFIER` and `OWNER` are replaced with `CATALOG` and `SCHEMA`, respectively.

Information about table indexes is also available from the `SHOW INDEX` statement. See [SHOW INDEX Statement](#). The following statements are equivalent:

```
SELECT * FROM INFORMATION_SCHEMA.STATISTICS
  WHERE table_name = 'tbl_name'
  AND table_schema = 'db_name'
SHOW INDEX
  FROM tbl_name
  FROM db_name
```

## 4.23 The INFORMATION\_SCHEMA TABLES Table

The `TABLES` table provides information about tables in databases.

The `TABLES` table has these columns:

- `TABLE_CATALOG`

The name of the catalog to which the table belongs. This value is always `def`.

- `TABLE_SCHEMA`

The name of the schema (database) to which the table belongs.

- `TABLE_NAME`

The name of the table.

- `TABLE_TYPE`

`BASE TABLE` for a table, `VIEW` for a view, or `SYSTEM VIEW` for an `INFORMATION_SCHEMA` table.

The `TABLES` table does not list `TEMPORARY` tables.

- `ENGINE`

The storage engine for the table. See [The InnoDB Storage Engine](#), and [Alternative Storage Engines](#).

For partitioned tables, `ENGINE` shows the name of the storage engine used by all partitions.

- `VERSION`

The version number of the table's `.frm` file.

- `ROW_FORMAT`

The row-storage format (`Fixed`, `Dynamic`, `Compressed`, `Redundant`, `Compact`). For `MyISAM` tables, `Dynamic` corresponds to what `myisamchk -dvv` reports as `Packed`. `InnoDB` table format is either `Redundant` or `Compact` when using the `Antelope` file format, or `Compressed` or `Dynamic` when using the `Barracuda` file format.

- `TABLE_ROWS`

The number of rows. Some storage engines, such as `MyISAM`, store the exact count. For other storage engines, such as `InnoDB`, this value is an approximation, and may vary from the actual

value by as much as 40% to 50%. In such cases, use `SELECT COUNT(*)` to obtain an accurate count.

`TABLE_ROWS` is `NULL` for `INFORMATION_SCHEMA` tables.

For `InnoDB` tables, the row count is only a rough estimate used in SQL optimization. (This is also true if the `InnoDB` table is partitioned.)

- `AVG_ROW_LENGTH`

The average row length.

Refer to the notes at the end of this section for related information.

- `DATA_LENGTH`

For `MyISAM`, `DATA_LENGTH` is the length of the data file, in bytes.

For `InnoDB`, `DATA_LENGTH` is the approximate amount of space allocated for the clustered index, in bytes. Specifically, it is the clustered index size, in pages, multiplied by the `InnoDB` page size.

Refer to the notes at the end of this section for information regarding other storage engines.

- `MAX_DATA_LENGTH`

For `MyISAM`, `MAX_DATA_LENGTH` is maximum length of the data file. This is the total number of bytes of data that can be stored in the table, given the data pointer size used.

Unused for `InnoDB`.

Refer to the notes at the end of this section for information regarding other storage engines.

- `INDEX_LENGTH`

For `MyISAM`, `INDEX_LENGTH` is the length of the index file, in bytes.

For `InnoDB`, `INDEX_LENGTH` is the approximate amount of space allocated for non-clustered indexes, in bytes. Specifically, it is the sum of non-clustered index sizes, in pages, multiplied by the `InnoDB` page size.

Refer to the notes at the end of this section for information regarding other storage engines.

- `DATA_FREE`

The number of allocated but unused bytes.

`InnoDB` tables report the free space of the tablespace to which the table belongs. For a table located in the shared tablespace, this is the free space of the shared tablespace. If you are using multiple tablespaces and the table has its own tablespace, the free space is for only that table. Free space means the number of bytes in completely free extents minus a safety margin. Even if free space displays as 0, it may be possible to insert rows as long as new extents need not be allocated.

For `NDB Cluster`, `DATA_FREE` shows the space allocated on disk for, but not used by, a Disk Data table or fragment on disk. (In-memory data resource usage is reported by the `DATA_LENGTH` column.)

For partitioned tables, this value is only an estimate and may not be absolutely correct. A more accurate method of obtaining this information in such cases is to query the `INFORMATION_SCHEMA PARTITIONS` table, as shown in this example:

```
SELECT SUM(DATA_FREE)
FROM INFORMATION_SCHEMA.PARTITIONS
WHERE TABLE_SCHEMA = 'mydb'
```

```
AND TABLE_NAME = 'mytable';
```

For more information, see [Section 4.14, “The INFORMATION\\_SCHEMA PARTITIONS Table”](#).

- [AUTO\\_INCREMENT](#)

The next [AUTO\\_INCREMENT](#) value.

- [CREATE\\_TIME](#)

When the table was created.

Prior to MySQL 5.6.25, for partitioned [InnoDB](#) tables, the [CREATE\\_TIME](#) column shows [NULL](#). This column shows the correct table creation time for such tables in MySQL 5.6.25 and later. (Bug #17299181, Bug #69990)

- [UPDATE\\_TIME](#)

When the data file was last updated. For some storage engines, this value is [NULL](#). For example, [InnoDB](#) stores multiple tables in its [system tablespace](#) and the data file timestamp does not apply. Even with [file-per-table](#) mode with each [InnoDB](#) table in a separate `.ibd` file, [change buffering](#) can delay the write to the data file, so the file modification time is different from the time of the last insert, update, or delete. For [MyISAM](#), the data file timestamp is used; however, on Windows the timestamp is not updated by updates, so the value is inaccurate.

For partitioned [InnoDB](#) tables, [UPDATE\\_TIME](#) is always [NULL](#).

- [CHECK\\_TIME](#)

When the table was last checked. Not all storage engines update this time, in which case, the value is always [NULL](#).

For partitioned [InnoDB](#) tables, [CHECK\\_TIME](#) is always [NULL](#).

- [TABLE\\_COLLATION](#)

The table default collation. The output does not explicitly list the table default character set, but the collation name begins with the character set name.

- [CHECKSUM](#)

The live checksum value, if any.

- [CREATE\\_OPTIONS](#)

Extra options used with `CREATE TABLE`.

[CREATE\\_OPTIONS](#) shows `partitioned` for partitioned tables.

When creating a table with [strict mode](#) disabled, the storage engine's default row format is used if the specified row format is not supported. The actual row format of the table is reported in the [ROW\\_FORMAT](#) column. [CREATE\\_OPTIONS](#) shows the row format that was specified in the `CREATE TABLE` statement.

When altering the storage engine of a table, table options that are not applicable to the new storage engine are retained in the table definition to enable reverting the table with its previously defined options to the original storage engine, if necessary. The [CREATE\\_OPTIONS](#) column may show retained options.

- [TABLE\\_COMMENT](#)

The comment used when creating the table (or information as to why MySQL could not access the table information).

## Notes

- For **NDB** tables, the output of this statement shows appropriate values for the `AVG_ROW_LENGTH` and `DATA_LENGTH` columns, with the exception that **BLOB** columns are not taken into account.
- For **NDB** tables, `DATA_LENGTH` includes data stored in main memory only; the `MAX_DATA_LENGTH` and `DATA_FREE` columns apply to Disk Data.
- For **NDB Cluster Disk Data** tables, `MAX_DATA_LENGTH` shows the space allocated for the disk part of a Disk Data table or fragment. (In-memory data resource usage is reported by the `DATA_LENGTH` column.)
- For **MEMORY** tables, the `DATA_LENGTH`, `MAX_DATA_LENGTH`, and `INDEX_LENGTH` values approximate the actual amount of allocated memory. The allocation algorithm reserves memory in large amounts to reduce the number of allocation operations.
- For views, all `TABLES` columns are `NULL` except that `TABLE_NAME` indicates the view name and `TABLE_COMMENT` says `VIEW`.

Table information is also available from the `SHOW TABLE STATUS` and `SHOW TABLES` statements. See [SHOW TABLE STATUS Statement](#), and [SHOW TABLES Statement](#). The following statements are equivalent:

```
SELECT
  TABLE_NAME, ENGINE, VERSION, ROW_FORMAT, TABLE_ROWS, AVG_ROW_LENGTH,
  DATA_LENGTH, MAX_DATA_LENGTH, INDEX_LENGTH, DATA_FREE, AUTO_INCREMENT,
  CREATE_TIME, UPDATE_TIME, CHECK_TIME, TABLE_COLLATION, CHECKSUM,
  CREATE_OPTIONS, TABLE_COMMENT
FROM INFORMATION_SCHEMA.TABLES
WHERE table_schema = 'db_name'
[AND table_name LIKE 'wild']
SHOW TABLE STATUS
FROM db_name
[LIKE 'wild']
```

The following statements are equivalent:

```
SELECT
  TABLE_NAME, TABLE_TYPE
FROM INFORMATION_SCHEMA.TABLES
WHERE table_schema = 'db_name'
[AND table_name LIKE 'wild']
SHOW FULL TABLES
FROM db_name
[LIKE 'wild']
```

## 4.24 The INFORMATION\_SCHEMA TABLESPACES Table

This table is unused. Other `INFORMATION_SCHEMA` tables may provide related information:

- For **NDB**, the `INFORMATION_SCHEMA FILES` table provides tablespace-related information.
- For **InnoDB**, the `INFORMATION_SCHEMA INNODB_SYS_TABLESPACES` and `INNODB_SYS_DATAFILES` tables provide tablespace metadata.

## 4.25 The INFORMATION\_SCHEMA TABLE\_CONSTRAINTS Table

The `TABLE_CONSTRAINTS` table describes which tables have constraints.

The `TABLE_CONSTRAINTS` table has these columns:

- `CONSTRAINT_CATALOG`

The name of the catalog to which the constraint belongs. This value is always `def`.

- `CONSTRAINT_SCHEMA`

The name of the schema (database) to which the constraint belongs.

- `TABLE_SCHEMA`

The name of the schema (database) to which the table belongs.

- `TABLE_NAME`

The name of the table.

- The `CONSTRAINT_TYPE`

The type of constraint. The value can be `UNIQUE`, `PRIMARY KEY`, `FOREIGN KEY`, or `CHECK`. This is a `CHAR` (not `ENUM`) column. The `CHECK` value is not available until MySQL supports `CHECK`.

The `UNIQUE` and `PRIMARY KEY` information is about the same as what you get from the `Key_name` column in the output from `SHOW INDEX` when the `Non_unique` column is 0.

## 4.26 The INFORMATION\_SCHEMA TABLE\_PRIVILEGES Table

The `TABLE_PRIVILEGES` table provides information about table privileges. It takes its values from the `mysql.tables_priv` system table.

The `TABLE_PRIVILEGES` table has these columns:

- `GRANTEE`

The name of the account to which the privilege is granted, in `'user_name'@'host_name'` format.

- `TABLE_CATALOG`

The name of the catalog to which the table belongs. This value is always `def`.

- `TABLE_SCHEMA`

The name of the schema (database) to which the table belongs.

- `TABLE_NAME`

The name of the table.

- `PRIVILEGE_TYPE`

The privilege granted. The value can be any privilege that can be granted at the table level; see [GRANT Statement](#). Each row lists a single privilege, so there is one row per table privilege held by the grantee.

- `IS_GRANTABLE`

`YES` if the user has the `GRANT OPTION` privilege, `NO` otherwise. The output does not list `GRANT OPTION` as a separate row with `PRIVILEGE_TYPE='GRANT OPTION'`.

### Notes

- `TABLE_PRIVILEGES` is a nonstandard `INFORMATION_SCHEMA` table.

The following statements are *not* equivalent:

```
SELECT ... FROM INFORMATION_SCHEMA.TABLE_PRIVILEGES
SHOW GRANTS ...
```

## 4.27 The INFORMATION\_SCHEMA TRIGGERS Table

The `TRIGGERS` table provides information about triggers. To see information about a table's triggers, you must have the `TRIGGER` privilege for the table.

The `TRIGGERS` table has these columns:

- `TRIGGER_CATALOG`

The name of the catalog to which the trigger belongs. This value is always `def`.

- `TRIGGER_SCHEMA`

The name of the schema (database) to which the trigger belongs.

- `TRIGGER_NAME`

The name of the trigger.

- `EVENT_MANIPULATION`

The trigger event. This is the type of operation on the associated table for which the trigger activates. The value is `INSERT` (a row was inserted), `DELETE` (a row was deleted), or `UPDATE` (a row was modified).

- `EVENT_OBJECT_CATALOG`, `EVENT_OBJECT_SCHEMA`, and `EVENT_OBJECT_TABLE`

As noted in [Using Triggers](#), every trigger is associated with exactly one table. These columns indicate the catalog and schema (database) in which this table occurs, and the table name, respectively. The `EVENT_OBJECT_CATALOG` value is always `def`.

- `ACTION_ORDER`

The ordinal position of the trigger's action within the list of all similar triggers on the same table. This value is always `0` because it is not possible to have more than one trigger with the same `EVENT_MANIPULATION` and `ACTION_TIMING` on the same table.

- `ACTION_CONDITION`

This value is always `NULL`.

- `ACTION_STATEMENT`

The trigger body; that is, the statement executed when the trigger activates. This text uses UTF-8 encoding.

- `ACTION_ORIENTATION`

This value is always `ROW`.

- `ACTION_TIMING`

Whether the trigger activates before or after the triggering event. The value is `BEFORE` or `AFTER`.

- `ACTION_REFERENCE_OLD_TABLE`

This value is always `NULL`.

- `ACTION_REFERENCE_NEW_TABLE`

This value is always `NULL`.

- `ACTION_REFERENCE_OLD_ROW` and `ACTION_REFERENCE_NEW_ROW`

The old and new column identifiers, respectively. The `ACTION_REFERENCE_OLD_ROW` value is always `OLD` and the `ACTION_REFERENCE_NEW_ROW` value is always `NEW`.

- `CREATED`

This value is always `NULL`.

- `SQL_MODE`

The SQL mode in effect when the trigger was created, and under which the trigger executes. For the permitted values, see [Server SQL Modes](#).

- `DEFINER`

The account named in the `DEFINER` clause (often the user who created the trigger), in `'user_name'@'host_name'` format.

- `CHARACTER_SET_CLIENT`

The session value of the `character_set_client` system variable when the trigger was created.

- `COLLATION_CONNECTION`

The session value of the `collation_connection` system variable when the trigger was created.

- `DATABASE_COLLATION`

The collation of the database with which the trigger is associated.

## Example

The following example uses the `ins_sum` trigger defined in [Using Triggers](#):

```
mysql> SELECT * FROM INFORMATION_SCHEMA.TRIGGERS
      WHERE TRIGGER_SCHEMA='test' AND TRIGGER_NAME='ins_sum'\G
***** 1. row *****
      TRIGGER_CATALOG: def
      TRIGGER_SCHEMA: test
      TRIGGER_NAME: ins_sum
      EVENT_MANIPULATION: INSERT
      EVENT_OBJECT_CATALOG: def
      EVENT_OBJECT_SCHEMA: test
      EVENT_OBJECT_TABLE: account
      ACTION_ORDER: 0
      ACTION_CONDITION: NULL
      ACTION_STATEMENT: SET @sum = @sum + NEW.amount
      ACTION_ORIENTATION: ROW
      ACTION_TIMING: BEFORE
      ACTION_REFERENCE_OLD_TABLE: NULL
      ACTION_REFERENCE_NEW_TABLE: NULL
      ACTION_REFERENCE_OLD_ROW: OLD
      ACTION_REFERENCE_NEW_ROW: NEW
      CREATED: NULL
      SQL_MODE: NO_ENGINE_SUBSTITUTION
      DEFINER: me@localhost
      CHARACTER_SET_CLIENT: utf8
      COLLATION_CONNECTION: utf8_general_ci
      DATABASE_COLLATION: latin1_swedish_ci
```

Trigger information is also available from the `SHOW TRIGGERS` statement. See [SHOW TRIGGERS Statement](#).

## 4.28 The INFORMATION\_SCHEMA USER\_PRIVILEGES Table

The `USER_PRIVILEGES` table provides information about global privileges. It takes its values from the `mysql.user` system table.

The `USER_PRIVILEGES` table has these columns:

- `GRANTEE`

The name of the account to which the privilege is granted, in `'user_name'@'host_name'` format.

- `TABLE_CATALOG`

The name of the catalog. This value is always `def`.

- `PRIVILEGE_TYPE`

The privilege granted. The value can be any privilege that can be granted at the global level; see [GRANT Statement](#). Each row lists a single privilege, so there is one row per global privilege held by the grantee.

- `IS_GRANTABLE`

`YES` if the user has the `GRANT OPTION` privilege, `NO` otherwise. The output does not list `GRANT OPTION` as a separate row with `PRIVILEGE_TYPE='GRANT OPTION'`.

### Notes

- `USER_PRIVILEGES` is a nonstandard `INFORMATION_SCHEMA` table.

The following statements are *not* equivalent:

```
SELECT ... FROM INFORMATION_SCHEMA.USER_PRIVILEGES
SHOW GRANTS ...
```

## 4.29 The INFORMATION\_SCHEMA VIEWS Table

The `VIEWS` table provides information about views in databases. You must have the `SHOW VIEW` privilege to access this table.

The `VIEWS` table has these columns:

- `TABLE_CATALOG`

The name of the catalog to which the view belongs. This value is always `def`.

- `TABLE_SCHEMA`

The name of the schema (database) to which the view belongs.

- `TABLE_NAME`

The name of the view.

- `VIEW_DEFINITION`

The `SELECT` statement that provides the definition of the view. This column has most of what you see in the `Create Table` column that `SHOW CREATE VIEW` produces. Skip the words before `SELECT` and skip the words `WITH CHECK OPTION`. Suppose that the original statement was:

```
CREATE VIEW v AS
SELECT s2,s1 FROM t
WHERE s1 > 5
```

```
ORDER BY s1
WITH CHECK OPTION;
```

Then the view definition looks like this:

```
SELECT s2,s1 FROM t WHERE s1 > 5 ORDER BY s1
```

- [CHECK\\_OPTION](#)

The value of the [CHECK\\_OPTION](#) attribute. The value is one of [NONE](#), [CASCADE](#), or [LOCAL](#).

- [IS\\_UPDATABLE](#)

MySQL sets a flag, called the view updatability flag, at [CREATE VIEW](#) time. The flag is set to [YES](#) (true) if [UPDATE](#) and [DELETE](#) (and similar operations) are legal for the view. Otherwise, the flag is set to [NO](#) (false). The [IS\\_UPDATABLE](#) column in the [VIEWS](#) table displays the status of this flag. It means that the server always knows whether a view is updatable.

If a view is not updatable, statements such [UPDATE](#), [DELETE](#), and [INSERT](#) are illegal and are rejected. (Even if a view is updatable, it might not be possible to insert into it; for details, refer to [Updatable and Insertable Views](#).)

- [DEFINER](#)

The account of the user who created the view, in '[user\\_name](#)'@'[host\\_name](#)' format.

- [SECURITY\\_TYPE](#)

The view [SQL SECURITY](#) characteristic. The value is one of [DEFINER](#) or [INVOKER](#).

- [CHARACTER\\_SET\\_CLIENT](#)

The session value of the [character\\_set\\_client](#) system variable when the view was created.

- [COLLATION\\_CONNECTION](#)

The session value of the [collation\\_connection](#) system variable when the view was created.

## Notes

MySQL permits different [sql\\_mode](#) settings to tell the server the type of SQL syntax to support. For example, you might use the [ANSI](#) SQL mode to ensure MySQL correctly interprets the standard SQL concatenation operator, the double bar (`||`), in your queries. If you then create a view that concatenates items, you might worry that changing the [sql\\_mode](#) setting to a value different from [ANSI](#) could cause the view to become invalid. But this is not the case. No matter how you write out a view definition, MySQL always stores it the same way, in a canonical form. Here is an example that shows how the server changes a double bar concatenation operator to a [CONCAT\(\)](#) function:

```
mysql> SET sql_mode = 'ANSI';
Query OK, 0 rows affected (0.00 sec)
mysql> CREATE VIEW test.v AS SELECT 'a' || 'b' as coll;
Query OK, 0 rows affected (0.00 sec)
mysql> SELECT VIEW_DEFINITION FROM INFORMATION_SCHEMA.VIEWS
        WHERE TABLE_SCHEMA = 'test' AND TABLE_NAME = 'v';
+-----+
| VIEW_DEFINITION          |
+-----+
| select concat('a','b') AS `coll` |
+-----+
1 row in set (0.00 sec)
```

The advantage of storing a view definition in canonical form is that changes made later to the value of [sql\\_mode](#) do not affect the results from the view. However, an additional consequence is that comments prior to [SELECT](#) are stripped from the definition by the server.

---

# Chapter 5 INFORMATION\_SCHEMA InnoDB Tables

## Table of Contents

5.1 INFORMATION_SCHEMA InnoDB Table Reference .....	47
5.2 The INFORMATION_SCHEMA INNODB_BUFFER_PAGE Table .....	48
5.3 The INFORMATION_SCHEMA INNODB_BUFFER_PAGE_LRU Table .....	51
5.4 The INFORMATION_SCHEMA INNODB_BUFFER_POOL_STATS Table .....	54
5.5 The INFORMATION_SCHEMA INNODB_CMP and INNODB_CMP_RESET Tables .....	57
5.6 The INFORMATION_SCHEMA INNODB_CMPMEM and INNODB_CMPMEM_RESET Tables ..	58
5.7 The INFORMATION_SCHEMA INNODB_CMP_PER_INDEX and INNODB_CMP_PER_INDEX_RESET Tables .....	60
5.8 The INFORMATION_SCHEMA INNODB_FT_BEING_DELETED Table .....	61
5.9 The INFORMATION_SCHEMA INNODB_FT_CONFIG Table .....	62
5.10 The INFORMATION_SCHEMA INNODB_FT_DEFAULT_STOPWORD Table .....	63
5.11 The INFORMATION_SCHEMA INNODB_FT_DELETED Table .....	64
5.12 The INFORMATION_SCHEMA INNODB_FT_INDEX_CACHE Table .....	65
5.13 The INFORMATION_SCHEMA INNODB_FT_INDEX_TABLE Table .....	66
5.14 The INFORMATION_SCHEMA INNODB_LOCKS Table .....	68
5.15 The INFORMATION_SCHEMA INNODB_LOCK_WAITS Table .....	69
5.16 The INFORMATION_SCHEMA INNODB_METRICS Table .....	70
5.17 The INFORMATION_SCHEMA INNODB_SYS_COLUMNS Table .....	72
5.18 The INFORMATION_SCHEMA INNODB_SYS_DATAFILES Table .....	73
5.19 The INFORMATION_SCHEMA INNODB_SYS_FIELDS Table .....	74
5.20 The INFORMATION_SCHEMA INNODB_SYS_FOREIGN Table .....	74
5.21 The INFORMATION_SCHEMA INNODB_SYS_FOREIGN_COLS Table .....	75
5.22 The INFORMATION_SCHEMA INNODB_SYS_INDEXES Table .....	76
5.23 The INFORMATION_SCHEMA INNODB_SYS_TABLES Table .....	77
5.24 The INFORMATION_SCHEMA INNODB_SYS_TABLESPACES Table .....	78
5.25 The INFORMATION_SCHEMA INNODB_SYS_TABLESTATS View .....	79
5.26 The INFORMATION_SCHEMA INNODB_TRX Table .....	80

This section provides table definitions for [INFORMATION\\_SCHEMA InnoDB](#) tables. For related information and examples, see [InnoDB INFORMATION\\_SCHEMA Tables](#).

[INFORMATION\\_SCHEMA InnoDB](#) tables can be used to monitor ongoing [InnoDB](#) activity, to detect inefficiencies before they turn into issues, or to troubleshoot performance and capacity issues. As your database becomes bigger and busier, running up against the limits of your hardware capacity, you monitor and tune these aspects to keep the database running smoothly.

## 5.1 INFORMATION\_SCHEMA InnoDB Table Reference

The following table summarizes [INFORMATION\\_SCHEMA InnoDB](#) tables. For greater detail, see the individual table descriptions.

**Table 5.1 INFORMATION\_SCHEMA InnoDB Tables**

Name	Description
<a href="#">INNODB_BUFFER_PAGE</a>	Pages in InnoDB buffer pool
<a href="#">INNODB_BUFFER_PAGE_LRU</a>	LRU ordering of pages in InnoDB buffer pool
<a href="#">INNODB_BUFFER_POOL_STATS</a>	InnoDB buffer pool statistics
<a href="#">INNODB_CMP</a>	Status for operations related to compressed InnoDB tables

Name	Description
<a href="#">INNODB_CMP_PER_INDEX</a>	Status for operations related to compressed InnoDB tables and indexes
<a href="#">INNODB_CMP_PER_INDEX_RESET</a>	Status for operations related to compressed InnoDB tables and indexes
<a href="#">INNODB_CMP_RESET</a>	Status for operations related to compressed InnoDB tables
<a href="#">INNODB_CMPMEM</a>	Status for compressed pages within InnoDB buffer pool
<a href="#">INNODB_CMPMEM_RESET</a>	Status for compressed pages within InnoDB buffer pool
<a href="#">INNODB_FT_BEING_DELETED</a>	Snapshot of INNODB_FT_DELETED table
<a href="#">INNODB_FT_CONFIG</a>	Metadata for InnoDB table FULLTEXT index and associated processing
<a href="#">INNODB_FT_DEFAULT_STOPWORD</a>	Default list of stopwords for InnoDB FULLTEXT indexes
<a href="#">INNODB_FT_DELETED</a>	Rows deleted from InnoDB table FULLTEXT index
<a href="#">INNODB_FT_INDEX_CACHE</a>	Token information for newly inserted rows in InnoDB FULLTEXT index
<a href="#">INNODB_FT_INDEX_TABLE</a>	Inverted index information for processing text searches against InnoDB table FULLTEXT index
<a href="#">INNODB_LOCK_WAITS</a>	InnoDB transaction lock-wait information
<a href="#">INNODB_LOCKS</a>	InnoDB transaction lock information
<a href="#">INNODB_METRICS</a>	InnoDB performance information
<a href="#">INNODB_SYS_COLUMNS</a>	Columns in each InnoDB table
<a href="#">INNODB_SYS_DATAFILES</a>	Data file path information for InnoDB file-per-table and general tablespaces
<a href="#">INNODB_SYS_FIELDS</a>	Key columns of InnoDB indexes
<a href="#">INNODB_SYS_FOREIGN</a>	InnoDB foreign-key metadata
<a href="#">INNODB_SYS_FOREIGN_COLS</a>	InnoDB foreign-key column status information
<a href="#">INNODB_SYS_INDEXES</a>	InnoDB index metadata
<a href="#">INNODB_SYS_TABLES</a>	InnoDB table metadata
<a href="#">INNODB_SYS_TABLESPACES</a>	InnoDB file-per-table, general, and undo tablespace metadata
<a href="#">INNODB_SYS_TABLESTATS</a>	InnoDB table low-level status information
<a href="#">INNODB_TRX</a>	Active InnoDB transaction information

## 5.2 The INFORMATION\_SCHEMA INNODB\_BUFFER\_PAGE Table

The [INNODB\\_BUFFER\\_PAGE](#) table provides information about each [page](#) in the [InnoDB buffer pool](#).

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA Buffer Pool Tables](#).

### Warning

Querying the [INNODB\\_BUFFER\\_PAGE](#) table can affect performance. Do not query this table on a production system unless you are aware of the performance impact and have determined it to be acceptable. To avoid

impacting performance on a production system, reproduce the issue you want to investigate and query buffer pool statistics on a test instance.

The `INNODB_BUFFER_PAGE` table has these columns:

- `POOL_ID`

The buffer pool ID. This is an identifier to distinguish between multiple buffer pool instances.

- `BLOCK_ID`

The buffer pool block ID.

- `SPACE`

The tablespace ID; the same value as `INNODB_SYS_TABLES.SPACE`.

- `PAGE_NUMBER`

The page number.

- `PAGE_TYPE`

The page type. The following table shows the permitted values.

**Table 5.2 INNODB\_BUFFER\_PAGE.PAGE\_TYPE Values**

Page Type	Description
<code>ALLOCATED</code>	Freshly allocated page
<code>BLOB</code>	Uncompressed BLOB page
<code>COMPRESSED_BLOB2</code>	Subsequent comp BLOB page
<code>COMPRESSED_BLOB</code>	First compressed BLOB page
<code>EXTENT_DESCRIPTOR</code>	Extent descriptor page
<code>FILE_SPACE_HEADER</code>	File space header
<code>IBUF_BITMAP</code>	Insert buffer bitmap
<code>IBUF_FREE_LIST</code>	Insert buffer free list
<code>IBUF_INDEX</code>	Insert buffer index
<code>INDEX</code>	B-tree node
<code>INODE</code>	Index node
<code>SYSTEM</code>	System page
<code>TRX_SYSTEM</code>	Transaction system data
<code>UNDO_LOG</code>	Undo log page
<code>UNKNOWN</code>	Unknown

- `FLUSH_TYPE`

The flush type.

- `FIX_COUNT`

The number of threads using this block within the buffer pool. When zero, the block is eligible to be evicted.

- `IS_HASHED`

Whether a hash index has been built on this page.

- [NEWEST\\_MODIFICATION](#)

The Log Sequence Number of the youngest modification.

- [OLDEST\\_MODIFICATION](#)

The Log Sequence Number of the oldest modification.

- [ACCESS\\_TIME](#)

An abstract number used to judge the first access time of the page.

- [TABLE\\_NAME](#)

The name of the table the page belongs to. This column is applicable only to pages with a [PAGE\\_TYPE](#) value of [INDEX](#).

- [INDEX\\_NAME](#)

The name of the index the page belongs to. This can be the name of a clustered index or a secondary index. This column is applicable only to pages with a [PAGE\\_TYPE](#) value of [INDEX](#).

- [NUMBER\\_RECORDS](#)

The number of records within the page.

- [DATA\\_SIZE](#)

The sum of the sizes of the records. This column is applicable only to pages with a [PAGE\\_TYPE](#) value of [INDEX](#).

- [COMPRESSED\\_SIZE](#)

The compressed page size. [NULL](#) for pages that are not compressed.

- [PAGE\\_STATE](#)

The page state. The following table shows the permitted values.

**Table 5.3 INNODB\_BUFFER\_PAGE.PAGE\_STATE Values**

Page State	Description
<a href="#">FILE_PAGE</a>	A buffered file page
<a href="#">MEMORY</a>	Contains a main memory object
<a href="#">NOT_USED</a>	In the free list
<a href="#">NULL</a>	Clean compressed pages, compressed pages in the flush list, pages used as buffer pool watch sentinels
<a href="#">READY_FOR_USE</a>	A free page
<a href="#">REMOVE_HASH</a>	Hash index should be removed before placing in the free list

- [IO\\_FIX](#)

Whether any I/O is pending for this page: [IO\\_NONE](#) = no pending I/O, [IO\\_READ](#) = read pending, [IO\\_WRITE](#) = write pending.

- [IS\\_OLD](#)

Whether the block is in the sublist of old blocks in the LRU list.

- `FREE_PAGE_CLOCK`

The value of the `freed_page_clock` counter when the block was the last placed at the head of the LRU list. The `freed_page_clock` counter tracks the number of blocks removed from the end of the LRU list.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_BUFFER_PAGE LIMIT 1\G
***** 1. row *****
      POOL_ID: 0
      BLOCK_ID: 0
      SPACE: 97
      PAGE_NUMBER: 2473
      PAGE_TYPE: INDEX
      FLUSH_TYPE: 1
      FIX_COUNT: 0
      IS_HASHED: YES
NEWEST_MODIFICATION: 733855581
OLDEST_MODIFICATION: 0
      ACCESS_TIME: 3378385672
      TABLE_NAME: `employees`.`salaries`
      INDEX_NAME: PRIMARY
      NUMBER_RECORDS: 468
      DATA_SIZE: 14976
      COMPRESSED_SIZE: 0
      PAGE_STATE: FILE_PAGE
      IO_FIX: IO_NONE
      IS_OLD: YES
      FREE_PAGE_CLOCK: 66
```

## Notes

- This table is useful primarily for expert-level performance monitoring, or when developing performance-related extensions for MySQL.
- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA.COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.
- When tables, table rows, partitions, or indexes are deleted, associated pages remain in the buffer pool until space is required for other data. The `INNODB_BUFFER_PAGE` table reports information about these pages until they are evicted from the buffer pool. For more information about how the `InnoDB` manages buffer pool data, see [Buffer Pool](#).

## 5.3 The INFORMATION\_SCHEMA INNODB\_BUFFER\_PAGE\_LRU Table

The `INNODB_BUFFER_PAGE_LRU` table provides information about the pages in the `InnoDB` buffer pool; in particular, how they are ordered in the LRU list that determines which pages to `evict` from the buffer pool when it becomes full.

The `INNODB_BUFFER_PAGE_LRU` table has the same columns as the `INNODB_BUFFER_PAGE` table, except that the `INNODB_BUFFER_PAGE_LRU` table has `LRU_POSITION` and `COMPRESSED` columns instead of `BLOCK_ID` and `PAGE_STATE` columns.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA Buffer Pool Tables](#).

### Warning

Querying the `INNODB_BUFFER_PAGE_LRU` table can affect performance. Do not query this table on a production system unless you are aware of

the performance impact and have determined it to be acceptable. To avoid impacting performance on a production system, reproduce the issue you want to investigate and query buffer pool statistics on a test instance.

The `INNODB_BUFFER_PAGE_LRU` table has these columns:

- `POOL_ID`

The buffer pool ID. This is an identifier to distinguish between multiple buffer pool instances.

- `LRU_POSITION`

The position of the page in the LRU list.

- `SPACE`

The tablespace ID; the same value as `INNODB_SYS_TABLES.SPACE`.

- `PAGE_NUMBER`

The page number.

- `PAGE_TYPE`

The page type. The following table shows the permitted values.

**Table 5.4 INNODB\_BUFFER\_PAGE\_LRU.PAGE\_TYPE Values**

Page Type	Description
<code>ALLOCATED</code>	Freshly allocated page
<code>BLOB</code>	Uncompressed BLOB page
<code>COMPRESSED_BLOB2</code>	Subsequent comp BLOB page
<code>COMPRESSED_BLOB</code>	First compressed BLOB page
<code>EXTENT_DESCRIPTOR</code>	Extent descriptor page
<code>FILE_SPACE_HEADER</code>	File space header
<code>IBUF_BITMAP</code>	Insert buffer bitmap
<code>IBUF_FREE_LIST</code>	Insert buffer free list
<code>IBUF_INDEX</code>	Insert buffer index
<code>INDEX</code>	B-tree node
<code>INODE</code>	Index node
<code>SYSTEM</code>	System page
<code>TRX_SYSTEM</code>	Transaction system data
<code>UNDO_LOG</code>	Undo log page
<code>UNKNOWN</code>	Unknown

- `FLUSH_TYPE`

The flush type.

- `FIX_COUNT`

The number of threads using this block within the buffer pool. When zero, the block is eligible to be evicted.

- `IS_HASHED`

Whether a hash index has been built on this page.

- `NEWEST_MODIFICATION`  
The Log Sequence Number of the youngest modification.
- `OLDEST_MODIFICATION`  
The Log Sequence Number of the oldest modification.
- `ACCESS_TIME`  
An abstract number used to judge the first access time of the page.
- `TABLE_NAME`  
The name of the table the page belongs to. This column is applicable only to pages with a `PAGE_TYPE` value of `INDEX`.
- `INDEX_NAME`  
The name of the index the page belongs to. This can be the name of a clustered index or a secondary index. This column is applicable only to pages with a `PAGE_TYPE` value of `INDEX`.
- `NUMBER_RECORDS`  
The number of records within the page.
- `DATA_SIZE`  
The sum of the sizes of the records. This column is applicable only to pages with a `PAGE_TYPE` value of `INDEX`.
- `COMPRESSED_SIZE`  
The compressed page size. `NULL` for pages that are not compressed.
- `COMPRESSED`  
Whether the page is compressed.
- `IO_FIX`  
Whether any I/O is pending for this page: `IO_NONE` = no pending I/O, `IO_READ` = read pending, `IO_WRITE` = write pending.
- `IS_OLD`  
Whether the block is in the sublist of old blocks in the LRU list.
- `FREE_PAGE_CLOCK`  
The value of the `freed_page_clock` counter when the block was the last placed at the head of the LRU list. The `freed_page_clock` counter tracks the number of blocks removed from the end of the LRU list.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_BUFFER_PAGE_LRU LIMIT 1\G
***** 1. row *****
      POOL_ID: 0
    LRU_POSITION: 0
         SPACE: 97
    PAGE_NUMBER: 1984
     PAGE_TYPE: INDEX
     FLUSH_TYPE: 1
     FIX_COUNT: 0
```

```

      IS_HASHED: YES
NEWEST_MODIFICATION: 719490396
OLDEST_MODIFICATION: 0
      ACCESS_TIME: 3378383796
      TABLE_NAME: `employees`.`salaries`
      INDEX_NAME: PRIMARY
      NUMBER_RECORDS: 468
      DATA_SIZE: 14976
COMPRESSED_SIZE: 0
      COMPRESSED: NO
      IO_FIX: IO_NONE
      IS_OLD: YES
      FREE_PAGE_CLOCK: 0

```

## Notes

- This table is useful primarily for expert-level performance monitoring, or when developing performance-related extensions for MySQL.
- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.
- Querying this table can require MySQL to allocate a large block of contiguous memory, more than 64 bytes times the number of active pages in the buffer pool. This allocation could potentially cause an out-of-memory error, especially for systems with multi-gigabyte buffer pools.
- Querying this table requires MySQL to lock the data structure representing the buffer pool while traversing the LRU list, which can reduce concurrency, especially for systems with multi-gigabyte buffer pools.
- When tables, table rows, partitions, or indexes are deleted, associated pages remain in the buffer pool until space is required for other data. The [INNODB\\_BUFFER\\_PAGE\\_LRU](#) table reports information about these pages until they are evicted from the buffer pool. For more information about how the [InnoDB](#) manages buffer pool data, see [Buffer Pool](#).

## 5.4 The INFORMATION\_SCHEMA INNODB\_BUFFER\_POOL\_STATS Table

The [INNODB\\_BUFFER\\_POOL\\_STATS](#) table provides much of the same buffer pool information provided in [SHOW ENGINE INNODB STATUS](#) output. Much of the same information may also be obtained using [InnoDB](#) buffer pool [server status variables](#).

The idea of making pages in the buffer pool “young” or “not young” refers to transferring them between the [sublists](#) at the head and tail of the buffer pool data structure. Pages made “young” take longer to age out of the buffer pool, while pages made “not young” are moved much closer to the point of [eviction](#).

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA Buffer Pool Tables](#).

The [INNODB\\_BUFFER\\_POOL\\_STATS](#) table has these columns:

- [POOL\\_ID](#)

The buffer pool ID. This is an identifier to distinguish between multiple buffer pool instances.

- [POOL\\_SIZE](#)

The [InnoDB](#) buffer pool size in pages.

- [FREE\\_BUFFERS](#)

The number of free pages in the `InnoDB` buffer pool.

- `DATABASE_PAGES`

The number of pages in the `InnoDB` buffer pool containing data. This number includes both dirty and clean pages.

- `OLD_DATABASE_PAGES`

The number of pages in the `old` buffer pool sublist.

- `MODIFIED_DATABASE_PAGES`

The number of modified (dirty) database pages.

- `PENDING_DECOMPRESS`

The number of pages pending decompression.

- `PENDING_READS`

The number of pending reads.

- `PENDING_FLUSH_LRU`

The number of pages pending flush in the LRU.

- `PENDING_FLUSH_LIST`

The number of pages pending flush in the flush list.

- `PAGES_MADE_YOUNG`

The number of pages made young.

- `PAGES_NOT_MADE_YOUNG`

The number of pages not made young.

- `PAGES_MADE_YOUNG_RATE`

The number of pages made young per second (pages made young since the last printout / time elapsed).

- `PAGES_MADE_NOT_YOUNG_RATE`

The number of pages not made per second (pages not made young since the last printout / time elapsed).

- `NUMBER_PAGES_READ`

The number of pages read.

- `NUMBER_PAGES_CREATED`

The number of pages created.

- `NUMBER_PAGES_WRITTEN`

The number of pages written.

- `PAGES_READ_RATE`

The number of pages read per second (pages read since the last printout / time elapsed).

- `PAGES_CREATE_RATE`

The number of pages created per second (pages created since the last printout / time elapsed).

- `PAGES_WRITTEN_RATE`

The number of pages written per second (pages written since the last printout / time elapsed).

- `NUMBER_PAGES_GET`

The number of logical read requests.

- `HIT_RATE`

The buffer pool hit rate.

- `YOUNG_MAKE_PER_THOUSAND_GETS`

The number of pages made young per thousand gets.

- `NOT_YOUNG_MAKE_PER_THOUSAND_GETS`

The number of pages not made young per thousand gets.

- `NUMBER_PAGES_READ_AHEAD`

The number of pages read ahead.

- `NUMBER_READ_AHEAD_EVICTED`

The number of pages read into the `InnoDB` buffer pool by the read-ahead background thread that were subsequently evicted without having been accessed by queries.

- `READ_AHEAD_RATE`

The read-ahead rate per second (pages read ahead since the last printout / time elapsed).

- `READ_AHEAD_EVICTED_RATE`

The number of read-ahead pages evicted without access per second (read-ahead pages not accessed since the last printout / time elapsed).

- `LRU_IO_TOTAL`

Total LRU I/O.

- `LRU_IO_CURRENT`

LRU I/O for the current interval.

- `UNCOMPRESS_TOTAL`

The total number of pages decompressed.

- `UNCOMPRESS_CURRENT`

The number of pages decompressed in the current interval.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_BUFFER_POOL_STATS\G
***** 1. row *****
                POOL_ID: 0
                POOL_SIZE: 8192
```

```

        FREE_BUFFERS: 1
        DATABASE_PAGES: 8085
        OLD_DATABASE_PAGES: 2964
    MODIFIED_DATABASE_PAGES: 0
        PENDING_DECOMPRESS: 0
        PENDING_READS: 0
        PENDING_FLUSH_LRU: 0
        PENDING_FLUSH_LIST: 0
        PAGES_MADE_YOUNG: 22821
        PAGES_NOT_MADE_YOUNG: 3544303
        PAGES_MADE_YOUNG_RATE: 357.62602199870594
    PAGES_MADE_NOT_YOUNG_RATE: 0
        NUMBER_PAGES_READ: 2389
        NUMBER_PAGES_CREATED: 12385
        NUMBER_PAGES_WRITTEN: 13111
        PAGES_READ_RATE: 0
        PAGES_CREATE_RATE: 0
        PAGES_WRITTEN_RATE: 0
        NUMBER_PAGES_GET: 33322210
            HIT_RATE: 1000
    YOUNG_MAKE_PER_THOUSAND_GETS: 18
    NOT_YOUNG_MAKE_PER_THOUSAND_GETS: 0
        NUMBER_PAGES_READ_AHEAD: 2024
        NUMBER_READ_AHEAD_EVICTED: 0
            READ_AHEAD_RATE: 0
        READ_AHEAD_EVICTED_RATE: 0
            LRU_IO_TOTAL: 0
            LRU_IO_CURRENT: 0
            UNCOMPRESS_TOTAL: 0
            UNCOMPRESS_CURRENT: 0

```

## Notes

- This table is useful primarily for expert-level performance monitoring, or when developing performance-related extensions for MySQL.
- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.

## 5.5 The INFORMATION\_SCHEMA INNODB\_CMP and INNODB\_CMP\_RESET Tables

The `INNODB_CMP` and `INNODB_CMP_RESET` tables provide status information on operations related to [compressed InnoDB](#) tables.

The `INNODB_CMP` and `INNODB_CMP_RESET` tables have these columns:

- `PAGE_SIZE`

The compressed page size in bytes.

- `COMPRESS_OPS`

The number of times a B-tree page of size `PAGE_SIZE` has been compressed. Pages are compressed whenever an empty page is created or the space for the uncompressed modification log runs out.

- `COMPRESS_OPS_OK`

The number of times a B-tree page of size `PAGE_SIZE` has been successfully compressed. This count should never exceed `COMPRESS_OPS`.

- `COMPRESS_TIME`

The total time in seconds used for attempts to compress B-tree pages of size `PAGE_SIZE`.

- `UNCOMPRESS_OPS`

The number of times a B-tree page of size `PAGE_SIZE` has been uncompressed. B-tree pages are uncompressed whenever compression fails or at first access when the uncompressed page does not exist in the buffer pool.

- `UNCOMPRESS_TIME`

The total time in seconds used for uncompressing B-tree pages of the size `PAGE_SIZE`.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_CMP\G
***** 1. row *****
    page_size: 1024
    compress_ops: 0
    compress_ops_ok: 0
    compress_time: 0
    uncompress_ops: 0
    uncompress_time: 0
***** 2. row *****
    page_size: 2048
    compress_ops: 0
    compress_ops_ok: 0
    compress_time: 0
    uncompress_ops: 0
    uncompress_time: 0
***** 3. row *****
    page_size: 4096
    compress_ops: 0
    compress_ops_ok: 0
    compress_time: 0
    uncompress_ops: 0
    uncompress_time: 0
***** 4. row *****
    page_size: 8192
    compress_ops: 86955
    compress_ops_ok: 81182
    compress_time: 27
    uncompress_ops: 26828
    uncompress_time: 5
***** 5. row *****
    page_size: 16384
    compress_ops: 0
    compress_ops_ok: 0
    compress_time: 0
    uncompress_ops: 0
    uncompress_time: 0
```

## Notes

- Use these tables to measure the effectiveness of `InnoDB` table `compression` in your database.
- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA.COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.
- For usage information, see [Monitoring InnoDB Table Compression at Runtime](#) and [Using the Compression Information Schema Tables](#). For general information about `InnoDB` table compression, see [InnoDB Table Compression](#).

## 5.6 The INFORMATION\_SCHEMA INNODB\_CMPMEM and INNODB\_CMPMEM\_RESET Tables

The `INNODB_CMPMEM` and `INNODB_CMPMEM_RESET` tables provide status information on compressed pages within the InnoDB buffer pool.

The `INNODB_CMPMEM` and `INNODB_CMPMEM_RESET` tables have these columns:

- `PAGE_SIZE`

The block size in bytes. Each record of this table describes blocks of this size.

- `BUFFER_POOL_INSTANCE`

A unique identifier for the buffer pool instance.

- `PAGES_USED`

The number of blocks of size `PAGE_SIZE` that are currently in use.

- `PAGES_FREE`

The number of blocks of size `PAGE_SIZE` that are currently available for allocation. This column shows the external fragmentation in the memory pool. Ideally, these numbers should be at most 1.

- `RELOCATION_OPS`

The number of times a block of size `PAGE_SIZE` has been relocated. The buddy system can relocate the allocated “buddy neighbor” of a freed block when it tries to form a bigger freed block. Reading from the `INNODB_CMPMEM_RESET` table resets this count.

- `RELOCATION_TIME`

The total time in microseconds used for relocating blocks of size `PAGE_SIZE`. Reading from the table `INNODB_CMPMEM_RESET` resets this count.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_CMPMEM\G
***** 1. row *****
      page_size: 1024
buffer_pool_instance: 0
      pages_used: 0
      pages_free: 0
      relocation_ops: 0
      relocation_time: 0
***** 2. row *****
      page_size: 2048
buffer_pool_instance: 0
      pages_used: 0
      pages_free: 0
      relocation_ops: 0
      relocation_time: 0
***** 3. row *****
      page_size: 4096
buffer_pool_instance: 0
      pages_used: 0
      pages_free: 0
      relocation_ops: 0
      relocation_time: 0
***** 4. row *****
      page_size: 8192
buffer_pool_instance: 0
      pages_used: 7673
      pages_free: 15
      relocation_ops: 4638
      relocation_time: 0
***** 5. row *****
      page_size: 16384
```

```
buffer_pool_instance: 0
  pages_used: 0
  pages_free: 0
  relocation_ops: 0
  relocation_time: 0
```

## Notes

- Use these tables to measure the effectiveness of [InnoDB](#) table [compression](#) in your database.
- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.
- For usage information, see [Monitoring InnoDB Table Compression at Runtime](#) and [Using the Compression Information Schema Tables](#). For general information about [InnoDB](#) table compression, see [InnoDB Table Compression](#).

## 5.7 The INFORMATION\_SCHEMA INNODB\_CMP\_PER\_INDEX and INNODB\_CMP\_PER\_INDEX\_RESET Tables

The [INNODB\\_CMP\\_PER\\_INDEX](#) and [INNODB\\_CMP\\_PER\\_INDEX\\_RESET](#) tables provide status information on operations related to [compressed](#) [InnoDB](#) tables and indexes, with separate statistics for each combination of database, table, and index, to help you evaluate the performance and usefulness of compression for specific tables.

For a compressed [InnoDB](#) table, both the table data and all the [secondary indexes](#) are compressed. In this context, the table data is treated as just another index, one that happens to contain all the columns: the [clustered index](#).

The [INNODB\\_CMP\\_PER\\_INDEX](#) and [INNODB\\_CMP\\_PER\\_INDEX\\_RESET](#) tables have these columns:

- [DATABASE\\_NAME](#)

The schema (database) containing the applicable table.

- [TABLE\\_NAME](#)

The table to monitor for compression statistics.

- [INDEX\\_NAME](#)

The index to monitor for compression statistics.

- [COMPRESS\\_OPS](#)

The number of compression operations attempted. [Pages](#) are compressed whenever an empty page is created or the space for the uncompressed modification log runs out.

- [COMPRESS\\_OPS\\_OK](#)

The number of successful compression operations. Subtract from the [COMPRESS\\_OPS](#) value to get the number of [compression failures](#). Divide by the [COMPRESS\\_OPS](#) value to get the percentage of compression failures.

- [COMPRESS\\_TIME](#)

The total time in seconds used for compressing data in this index.

- [UNCOMPRESS\\_OPS](#)

The number of uncompression operations performed. Compressed [InnoDB](#) pages are uncompressed whenever compression [fails](#), or the first time a compressed page is accessed in the [buffer pool](#) and the uncompressed page does not exist.

- [UNCOMPRESS\\_TIME](#)

The total time in seconds used for uncompressing data in this index.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_CMP_PER_INDEX\G
***** 1. row *****
  database_name: employees
    table_name: salaries
    index_name: PRIMARY
    compress_ops: 0
compress_ops_ok: 0
  compress_time: 0
uncompress_ops: 23451
uncompress_time: 4
***** 2. row *****
  database_name: employees
    table_name: salaries
    index_name: emp_no
    compress_ops: 0
compress_ops_ok: 0
  compress_time: 0
uncompress_ops: 1597
uncompress_time: 0
```

## Notes

- Use these tables to measure the effectiveness of [InnoDB](#) table [compression](#) for specific tables, indexes, or both.
- You must have the [PROCESS](#) privilege to query these tables.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of these tables, including data types and default values.
- Because collecting separate measurements for every index imposes substantial performance overhead, [INNODB\\_CMP\\_PER\\_INDEX](#) and [INNODB\\_CMP\\_PER\\_INDEX\\_RESET](#) statistics are not gathered by default. You must enable the [innodb\\_cmp\\_per\\_index\\_enabled](#) system variable before performing the operations on compressed tables that you want to monitor.
- For usage information, see [Monitoring InnoDB Table Compression at Runtime](#) and [Using the Compression Information Schema Tables](#). For general information about [InnoDB](#) table compression, see [InnoDB Table Compression](#).

## 5.8 The INFORMATION\_SCHEMA INNODB\_FT\_BEING\_DELETED Table

The [INNODB\\_FT\\_BEING\\_DELETED](#) table provides a snapshot of the [INNODB\\_FT\\_DELETED](#) table; it is used only during an [OPTIMIZE TABLE](#) maintenance operation. When [OPTIMIZE TABLE](#) is run, the [INNODB\\_FT\\_BEING\\_DELETED](#) table is emptied, and [DOC\\_ID](#) values are removed from the [INNODB\\_FT\\_DELETED](#) table. Because the contents of [INNODB\\_FT\\_BEING\\_DELETED](#) typically have a short lifetime, this table has limited utility for monitoring or debugging. For information about running [OPTIMIZE TABLE](#) on tables with [FULLTEXT](#) indexes, see [Fine-Tuning MySQL Full-Text Search](#).

This table is empty initially. Before querying it, set the value of the [innodb\\_ft\\_aux\\_table](#) system variable to the name (including the database name) of the table that contains the [FULLTEXT](#)

index; for example `test/articles`. The output appears similar to the example provided for the `INNODB_FT_DELETED` table.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA FULLTEXT Index Tables](#).

The `INNODB_FT_BEING_DELETED` table has these columns:

- `DOC_ID`

The document ID of the row that is in the process of being deleted. This value might reflect the value of an ID column that you defined for the underlying table, or it can be a sequence value generated by `InnoDB` when the table contains no suitable column. This value is used when you do text searches, to skip rows in the `INNODB_FT_INDEX_TABLE` table before data for deleted rows is physically removed from the `FULLTEXT` index by an `OPTIMIZE TABLE` statement. For more information, see [Optimizing InnoDB Full-Text Indexes](#).

## Notes

- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.
- For more information about `InnoDB FULLTEXT` search, see [InnoDB Full-Text Indexes](#), and [Full-Text Search Functions](#).

## 5.9 The INFORMATION\_SCHEMA INNODB\_FT\_CONFIG Table

The `INNODB_FT_CONFIG` table provides metadata about the `FULLTEXT` index and associated processing for an `InnoDB` table.

This table is empty initially. Before querying it, set the value of the `innodb_ft_aux_table` system variable to the name (including the database name) of the table that contains the `FULLTEXT` index; for example `test/articles`.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA FULLTEXT Index Tables](#).

The `INNODB_FT_CONFIG` table has these columns:

- `KEY`

The name designating an item of metadata for an `InnoDB` table containing a `FULLTEXT` index.

The values for this column might change, depending on the needs for performance tuning and debugging for `InnoDB` full-text processing. The key names and their meanings include:

- `optimize_checkpoint_limit`: The number of seconds after which an `OPTIMIZE TABLE` run stops.
- `synced_doc_id`: The next `DOC_ID` to be issued.
- `stopword_table_name`: The `database/table` name for a user-defined stopword table. The `VALUE` column is empty if there is no user-defined stopword table.
- `use_stopword`: Indicates whether a stopword table is used, which is defined when the `FULLTEXT` index is created.
- `VALUE`

The value associated with the corresponding [KEY](#) column, reflecting some limit or current value for an aspect of a [FULLTEXT](#) index for an [InnoDB](#) table.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_FT_CONFIG;
+-----+-----+
| KEY          | VALUE          |
+-----+-----+
| optimize_checkpoint_limit | 180            |
| synced_doc_id | 0              |
| stopword_table_name | test/my_stopwords |
| use_stopword  | 1              |
+-----+-----+
```

## Notes

- This table is intended only for internal configuration. It is not intended for statistical information purposes.
- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.
- For more information about [InnoDB FULLTEXT](#) search, see [InnoDB Full-Text Indexes](#), and [Full-Text Search Functions](#).

## 5.10 The INFORMATION\_SCHEMA INNODB\_FT\_DEFAULT\_STOPWORD Table

The [INNODB\\_FT\\_DEFAULT\\_STOPWORD](#) table holds a list of [stopwords](#) that are used by default when creating a [FULLTEXT](#) index on [InnoDB](#) tables. For information about the default [InnoDB](#) stopword list and how to define your own stopword lists, see [Full-Text Stopwords](#).

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA FULLTEXT Index Tables](#).

The [INNODB\\_FT\\_DEFAULT\\_STOPWORD](#) table has these columns:

- [value](#)

A word that is used by default as a stopword for [FULLTEXT](#) indexes on [InnoDB](#) tables. This is not used if you override the default stopword processing with either the [innodb\\_ft\\_server\\_stopword\\_table](#) or the [innodb\\_ft\\_user\\_stopword\\_table](#) system variable.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_FT_DEFAULT_STOPWORD;
+-----+
| value |
+-----+
| a     |
| about |
| an    |
| are   |
| as    |
| at    |
| be    |
| by    |
| com   |
| de    |
+-----+
```

```

| en
| for
| from
| how
| i
| in
| is
| it
| la
| of
| on
| or
| that
| the
| this
| to
| was
| what
| when
| where
| who
| will
| with
| und
| the
| www
+-----+
36 rows in set (0.00 sec)

```

## Notes

- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.
- For more information about [InnoDB FULLTEXT](#) search, see [InnoDB Full-Text Indexes](#), and [Full-Text Search Functions](#).

## 5.11 The INFORMATION\_SCHEMA INNODB\_FT\_DELETED Table

The [INNODB\\_FT\\_DELETED](#) table stores rows that are deleted from the [FULLTEXT](#) index for an [InnoDB](#) table. To avoid expensive index reorganization during DML operations for an [InnoDB FULLTEXT](#) index, the information about newly deleted words is stored separately, filtered out of search results when you do a text search, and removed from the main search index only when you issue an [OPTIMIZE TABLE](#) statement for the [InnoDB](#) table. For more information, see [Optimizing InnoDB Full-Text Indexes](#).

This table is empty initially. Before querying it, set the value of the [innodb\\_ft\\_aux\\_table](#) system variable to the name (including the database name) of the table that contains the [FULLTEXT](#) index; for example [test/articles](#).

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA FULLTEXT Index Tables](#).

The [INNODB\\_FT\\_DELETED](#) table has these columns:

- [DOC\\_ID](#)

The document ID of the newly deleted row. This value might reflect the value of an ID column that you defined for the underlying table, or it can be a sequence value generated by [InnoDB](#) when the table contains no suitable column. This value is used when you do text searches, to skip rows in the [INNODB\\_FT\\_INDEX\\_TABLE](#) table before data for deleted rows is physically removed from the [FULLTEXT](#) index by an [OPTIMIZE TABLE](#) statement. For more information, see [Optimizing InnoDB Full-Text Indexes](#).

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_FT_DELETED;
+-----+
| DOC_ID |
+-----+
|      6 |
|      7 |
|      8 |
+-----+
```

## Notes

- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.
- For more information about [InnoDB FULLTEXT](#) search, see [InnoDB Full-Text Indexes](#), and [Full-Text Search Functions](#).

## 5.12 The INFORMATION\_SCHEMA INNODB\_FT\_INDEX\_CACHE Table

The [INNODB\\_FT\\_INDEX\\_CACHE](#) table provides token information about newly inserted rows in a [FULLTEXT](#) index. To avoid expensive index reorganization during DML operations, the information about newly indexed words is stored separately, and combined with the main search index only when [OPTIMIZE TABLE](#) is run, when the server is shut down, or when the cache size exceeds a limit defined by the [innodb\\_ft\\_cache\\_size](#) or [innodb\\_ft\\_total\\_cache\\_size](#) system variable.

This table is empty initially. Before querying it, set the value of the [innodb\\_ft\\_aux\\_table](#) system variable to the name (including the database name) of the table that contains the [FULLTEXT](#) index; for example `test/articles`.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA FULLTEXT Index Tables](#).

The [INNODB\\_FT\\_INDEX\\_CACHE](#) table has these columns:

- [WORD](#)  
A word extracted from the text of a newly inserted row.
- [FIRST\\_DOC\\_ID](#)  
The first document ID in which this word appears in the [FULLTEXT](#) index.
- [LAST\\_DOC\\_ID](#)  
The last document ID in which this word appears in the [FULLTEXT](#) index.
- [DOC\\_COUNT](#)  
The number of rows in which this word appears in the [FULLTEXT](#) index. The same word can occur several times within the cache table, once for each combination of [DOC\\_ID](#) and [POSITION](#) values.
- [DOC\\_ID](#)  
The document ID of the newly inserted row. This value might reflect the value of an ID column that you defined for the underlying table, or it can be a sequence value generated by [InnoDB](#) when the table contains no suitable column.

- [POSITION](#)

The position of this particular instance of the word within the relevant document identified by the [DOC\\_ID](#) value. The value does not represent an absolute position; it is an offset added to the [POSITION](#) of the previous instance of that word.

## Notes

- This table is empty initially. Before querying it, set the value of the [innodb\\_ft\\_aux\\_table](#) system variable to the name (including the database name) of the table that contains the [FULLTEXT](#) index; for example `test/articles`. The following example demonstrates how to use the [innodb\\_ft\\_aux\\_table](#) system variable to show information about a [FULLTEXT](#) index for a specified table.

```
mysql> USE test;
mysql> CREATE TABLE articles (
    id INT UNSIGNED AUTO_INCREMENT NOT NULL PRIMARY KEY,
    title VARCHAR(200),
    body TEXT,
    FULLTEXT (title,body)
) ENGINE=InnoDB;
mysql> INSERT INTO articles (title,body) VALUES
('MySQL Tutorial','DBMS stands for DataBase ...'),
('How To Use MySQL Well','After you went through a ...'),
('Optimizing MySQL','In this tutorial we show ...'),
('1001 MySQL Tricks','1. Never run mysqld as root. 2. ...'),
('MySQL vs. YourSQL','In the following database comparison ...'),
('MySQL Security','When configured properly, MySQL ...');
mysql> SET GLOBAL innodb_ft_aux_table = 'test/articles';
mysql> SELECT WORD, DOC_COUNT, DOC_ID, POSITION
FROM INFORMATION_SCHEMA.INNODB_FT_INDEX_CACHE LIMIT 5;
```

WORD	DOC_COUNT	DOC_ID	POSITION
1001	1	4	0
after	1	2	22
comparison	1	5	44
configured	1	6	20
database	2	1	31

- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.
- For more information about [InnoDB FULLTEXT](#) search, see [InnoDB Full-Text Indexes](#), and [Full-Text Search Functions](#).

## 5.13 The INFORMATION\_SCHEMA INNODB\_FT\_INDEX\_TABLE Table

The [INNODB\\_FT\\_INDEX\\_TABLE](#) table provides information about the inverted index used to process text searches against the [FULLTEXT](#) index of an [InnoDB](#) table.

This table is empty initially. Before querying it, set the value of the [innodb\\_ft\\_aux\\_table](#) system variable to the name (including the database name) of the table that contains the [FULLTEXT](#) index; for example `test/articles`.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA FULLTEXT Index Tables](#).

The [INNODB\\_FT\\_INDEX\\_TABLE](#) table has these columns:

- [WORD](#)

A word extracted from the text of the columns that are part of a `FULLTEXT`.

- `FIRST_DOC_ID`

The first document ID in which this word appears in the `FULLTEXT` index.

- `LAST_DOC_ID`

The last document ID in which this word appears in the `FULLTEXT` index.

- `DOC_COUNT`

The number of rows in which this word appears in the `FULLTEXT` index. The same word can occur several times within the cache table, once for each combination of `DOC_ID` and `POSITION` values.

- `DOC_ID`

The document ID of the row containing the word. This value might reflect the value of an ID column that you defined for the underlying table, or it can be a sequence value generated by `InnoDB` when the table contains no suitable column.

- `POSITION`

The position of this particular instance of the word within the relevant document identified by the `DOC_ID` value.

## Notes

- This table is empty initially. Before querying it, set the value of the `innodb_ft_aux_table` system variable to the name (including the database name) of the table that contains the `FULLTEXT` index; for example `test/articles`. The following example demonstrates how to use the `innodb_ft_aux_table` system variable to show information about a `FULLTEXT` index for a specified table. Before information for newly inserted rows appears in `INNODB_FT_INDEX_TABLE`, the `FULLTEXT` index cache must be flushed to disk. This is accomplished by running an `OPTIMIZE TABLE` operation on the indexed table with the `innodb_optimize_fulltext_only` system variable enabled. (The example disables that variable again at the end because it is intended to be enabled only temporarily.)

```
mysql> USE test;
mysql> CREATE TABLE articles (
    id INT UNSIGNED AUTO_INCREMENT NOT NULL PRIMARY KEY,
    title VARCHAR(200),
    body TEXT,
    FULLTEXT (title,body)
) ENGINE=InnoDB;
mysql> INSERT INTO articles (title,body) VALUES
('MySQL Tutorial','DBMS stands for DataBase ...'),
('How To Use MySQL Well','After you went through a ...'),
('Optimizing MySQL','In this tutorial we show ...'),
('1001 MySQL Tricks','1. Never run mysqld as root. 2. ...'),
('MySQL vs. YourSQL','In the following database comparison ...'),
('MySQL Security','When configured properly, MySQL ...');
mysql> SET GLOBAL innodb_optimize_fulltext_only=ON;
mysql> OPTIMIZE TABLE articles;
+-----+-----+-----+-----+
| Table          | Op      | Msg_type | Msg_text |
+-----+-----+-----+-----+
| test.articles | optimize | status   | OK       |
+-----+-----+-----+-----+
mysql> SET GLOBAL innodb_ft_aux_table = 'test/articles';
mysql> SELECT WORD, DOC_COUNT, DOC_ID, POSITION
FROM INFORMATION_SCHEMA.INNODB_FT_INDEX_TABLE LIMIT 5;
+-----+-----+-----+-----+
| WORD          | DOC_COUNT | DOC_ID | POSITION |
+-----+-----+-----+-----+
```

```

| 1001 | 1 | 4 | 0 |
| after | 1 | 2 | 22 |
| comparison | 1 | 5 | 44 |
| configured | 1 | 6 | 20 |
| database | 2 | 1 | 31 |
+-----+-----+-----+-----+
mysql> SET GLOBAL innodb_optimize_fulltext_only=OFF;

```

- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.
- For more information about [InnoDB FULLTEXT](#) search, see [InnoDB Full-Text Indexes](#), and [Full-Text Search Functions](#).

## 5.14 The INFORMATION\_SCHEMA INNODB\_LOCKS Table

The [INNODB\\_LOCKS](#) table provides information about each lock that an [InnoDB](#) transaction has requested but not yet acquired, and each lock that a transaction holds that is blocking another transaction.

The [INNODB\\_LOCKS](#) table has these columns:

- [LOCK\\_ID](#)

A unique lock ID number, internal to [InnoDB](#). Treat it as an opaque string. Although [LOCK\\_ID](#) currently contains [TRX\\_ID](#), the format of the data in [LOCK\\_ID](#) is subject to change at any time. Do not write applications that parse the [LOCK\\_ID](#) value.

- [LOCK\\_TRX\\_ID](#)

The ID of the transaction holding the lock. To obtain details about the transaction, join this column with the [TRX\\_ID](#) column of the [INNODB\\_TRX](#) table.

- [LOCK\\_MODE](#)

How the lock is requested. Permitted lock mode descriptors are [S](#), [X](#), [IS](#), [IX](#), [GAP](#), [AUTO\\_INC](#), and [UNKNOWN](#). Lock mode descriptors may be used in combination to identify particular lock modes. For information about [InnoDB](#) lock modes, see [InnoDB Locking](#).

- [LOCK\\_TYPE](#)

The type of lock. Permitted values are [RECORD](#) for a row-level lock, [TABLE](#) for a table-level lock.

- [LOCK\\_TABLE](#)

The name of the table that has been locked or contains locked records.

- [LOCK\\_INDEX](#)

The name of the index, if [LOCK\\_TYPE](#) is [RECORD](#); otherwise [NULL](#).

- [LOCK\\_SPACE](#)

The tablespace ID of the locked record, if [LOCK\\_TYPE](#) is [RECORD](#); otherwise [NULL](#).

- [LOCK\\_PAGE](#)

The page number of the locked record, if [LOCK\\_TYPE](#) is [RECORD](#); otherwise [NULL](#).

- [LOCK\\_REC](#)

The heap number of the locked record within the page, if [LOCK\\_TYPE](#) is [RECORD](#); otherwise [NULL](#).

- [LOCK\\_DATA](#)

The data associated with the lock, if any. A value is shown if the [LOCK\\_TYPE](#) is [RECORD](#), otherwise the value is [NULL](#). Primary key values of the locked record are shown for a lock placed on the primary key index. Secondary index values of the locked record are shown for a lock placed on a unique secondary index. Secondary index values are shown with primary key values appended if the secondary index is not unique. If there is no primary key, [LOCK\\_DATA](#) shows either the key values of a selected unique index or the unique [InnoDB](#) internal row ID number, according to the rules governing [InnoDB](#) clustered index use (see [Clustered and Secondary Indexes](#)). [LOCK\\_DATA](#) reports “supremum pseudo-record” for a lock taken on a supremum pseudo-record. If the page containing the locked record is not in the buffer pool because it was written to disk while the lock was held, [InnoDB](#) does not fetch the page from disk. Instead, [LOCK\\_DATA](#) reports [NULL](#).

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_LOCKS\G
***** 1. row *****
    lock_id: 3723:72:3:2
lock_trx_id: 3723
  lock_mode: X
  lock_type: RECORD
lock_table: `mysql`.`t`
lock_index: PRIMARY
lock_space: 72
  lock_page: 3
   lock_rec: 2
  lock_data: 1, 9
***** 2. row *****
    lock_id: 3722:72:3:2
lock_trx_id: 3722
  lock_mode: S
  lock_type: RECORD
lock_table: `mysql`.`t`
lock_index: PRIMARY
lock_space: 72
  lock_page: 3
   lock_rec: 2
  lock_data: 1, 9
```

## Notes

- Use this table to help diagnose performance problems that occur during times of heavy concurrent load. Its contents are updated as described in [Persistence and Consistency of InnoDB Transaction and Locking Information](#).
- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.
- For usage information, see [Using InnoDB Transaction and Locking Information](#).

## 5.15 The INFORMATION\_SCHEMA INNODB\_LOCK\_WAITS Table

The [INNODB\\_LOCK\\_WAITS](#) table contains one or more rows for each blocked [InnoDB](#) transaction, indicating the lock it has requested and any locks that are blocking that request.

The [INNODB\\_LOCK\\_WAITS](#) table has these columns:

- [REQUESTING\\_TRX\\_ID](#)

The ID of the requesting (blocked) transaction.

- [REQUESTED\\_LOCK\\_ID](#)

The ID of the lock for which a transaction is waiting. To obtain details about the lock, join this column with the `LOCK_ID` column of the `INNODB_LOCKS` table.

- `BLOCKING_TRX_ID`

The ID of the blocking transaction.

- `BLOCKING_LOCK_ID`

The ID of a lock held by a transaction blocking another transaction from proceeding. To obtain details about the lock, join this column with the `LOCK_ID` column of the `INNODB_LOCKS` table.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_LOCK_WAITS\G
***** 1. row *****
requesting_trx_id: 3396
requested_lock_id: 3396:91:3:2
  blocking_trx_id: 3395
  blocking_lock_id: 3395:91:3:2
```

## Notes

- Use this table to help diagnose performance problems that occur during times of heavy concurrent load. Its contents are updated as described in [Persistence and Consistency of InnoDB Transaction and Locking Information](#).
- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA.COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.
- For usage information, see [Using InnoDB Transaction and Locking Information](#).

## 5.16 The INFORMATION\_SCHEMA INNODB\_METRICS Table

The `INNODB_METRICS` table provides a wide variety of `InnoDB` performance information, complementing the specific focus areas of the Performance Schema tables for `InnoDB`. With simple queries, you can check the overall health of the system. With more detailed queries, you can diagnose issues such as performance bottlenecks, resource shortages, and application issues.

Each monitor represents a point within the `InnoDB` source code that is instrumented to gather counter information. Each counter can be started, stopped, and reset. You can also perform these actions for a group of counters using their common module name.

By default, relatively little data is collected. To start, stop, and reset counters, set one of the system variables `innodb_monitor_enable`, `innodb_monitor_disable`, `innodb_monitor_reset`, or `innodb_monitor_reset_all`, using the name of the counter, the name of the module, a wildcard match for such a name using the “%” character, or the special keyword `all`.

For usage information, see [InnoDB INFORMATION\\_SCHEMA Metrics Table](#).

The `INNODB_METRICS` table has these columns:

- `NAME`  
A unique name for the counter.
- `SUBSYSTEM`  
The aspect of `InnoDB` that the metric applies to.
- `COUNT`

The value since the counter was enabled.

- `MAX_COUNT`

The maximum value since the counter was enabled.

- `MIN_COUNT`

The minimum value since the counter was enabled.

- `AVG_COUNT`

The average value since the counter was enabled.

- `COUNT_RESET`

The counter value since it was last reset. (The `_RESET` columns act like the lap counter on a stopwatch: you can measure the activity during some time interval, while the cumulative figures are still available in `COUNT`, `MAX_COUNT`, and so on.)

- `MAX_COUNT_RESET`

The maximum counter value since it was last reset.

- `MIN_COUNT_RESET`

The minimum counter value since it was last reset.

- `AVG_COUNT_RESET`

The average counter value since it was last reset.

- `TIME_ENABLED`

The timestamp of the last start.

- `TIME_DISABLED`

The timestamp of the last stop.

- `TIME_ELAPSED`

The elapsed time in seconds since the counter started.

- `TIME_RESET`

The timestamp of the last reset.

- `STATUS`

Whether the counter is still running (`enabled`) or stopped (`disabled`).

- `TYPE`

Whether the item is a cumulative counter, or measures the current value of some resource.

- `COMMENT`

The counter description.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_METRICS WHERE NAME='dml_inserts'\G
***** 1. row *****
```

```

        NAME: dml_inserts
        SUBSYSTEM: dml
        COUNT: 3
        MAX_COUNT: 3
        MIN_COUNT: NULL
        AVG_COUNT: 0.046153846153846156
        COUNT_RESET: 3
        MAX_COUNT_RESET: 3
        MIN_COUNT_RESET: NULL
        AVG_COUNT_RESET: NULL
        TIME_ENABLED: 2014-12-04 14:18:28
        TIME_DISABLED: NULL
        TIME_ELAPSED: 65
        TIME_RESET: NULL
        STATUS: enabled
        TYPE: status_counter
        COMMENT: Number of rows inserted

```

## Notes

- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.

## 5.17 The INFORMATION\_SCHEMA INNODB\_SYS\_COLUMNS Table

The [INNODB\\_SYS\\_COLUMNS](#) table provides metadata about [InnoDB](#) table columns, equivalent to the information from the [SYS\\_COLUMNS](#) table in the [InnoDB](#) data dictionary.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA System Tables](#).

The [INNODB\\_SYS\\_COLUMNS](#) table has these columns:

- [TABLE\\_ID](#)

An identifier representing the table associated with the column; the same value as [INNODB\\_SYS\\_TABLES.TABLE\\_ID](#).

- [NAME](#)

The name of the column. These names can be uppercase or lowercase depending on the [lower\\_case\\_table\\_names](#) setting. There are no special system-reserved names for columns.

- [POS](#)

The ordinal position of the column within the table, starting from 0 and incrementing sequentially. When a column is dropped, the remaining columns are reordered so that the sequence has no gaps.

- [MTYPE](#)

Stands for “main type”. A numeric identifier for the column type. 1 = [VARCHAR](#), 2 = [CHAR](#), 3 = [FIXBINARY](#), 4 = [BINARY](#), 5 = [BLOB](#), 6 = [INT](#), 7 = [SYS\\_CHILD](#), 8 = [SYS](#), 9 = [FLOAT](#), 10 = [DOUBLE](#), 11 = [DECIMAL](#), 12 = [VARMYSQL](#), 13 = [MYSQL](#).

- [PRTYPE](#)

The [InnoDB](#) “precise type”, a binary value with bits representing MySQL data type, character set code, and nullability.

- [LEN](#)

The column length, for example 4 for [INT](#) and 8 for [BIGINT](#). For character columns in multibyte character sets, this length value is the maximum length in bytes needed to represent a definition such as [VARCHAR\(N\)](#); that is, it might be  $2*N$ ,  $3*N$ , and so on depending on the character encoding.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_SYS_COLUMNS where TABLE_ID = 71\G
***** 1. row *****
TABLE_ID: 71
  NAME: col1
  POS: 0
  MTYPE: 6
  PRTYPE: 1027
  LEN: 4
***** 2. row *****
TABLE_ID: 71
  NAME: col2
  POS: 1
  MTYPE: 2
  PRTYPE: 524542
  LEN: 10
***** 3. row *****
TABLE_ID: 71
  NAME: col3
  POS: 2
  MTYPE: 1
  PRTYPE: 524303
  LEN: 10
```

## Notes

- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.

## 5.18 The INFORMATION\_SCHEMA INNODB\_SYS\_DATAFILES Table

The [INNODB\\_SYS\\_DATAFILES](#) table provides data file path information for [InnoDB](#) tablespaces, equivalent to the information in the [SYS\\_DATAFILES](#) table in the [InnoDB](#) data dictionary.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA System Tables](#).

The [INNODB\\_SYS\\_DATAFILES](#) table has these columns:

- [SPACE](#)

The tablespace ID.

- [PATH](#)

The tablespace data file path. If a [file-per-table](#) tablespace is created in a location outside the MySQL data directory, the path value is a fully qualified directory path. Otherwise, the path is relative to the data directory.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_SYS_DATAFILES WHERE SPACE = 57\G
***** 1. row *****
SPACE: 57
PATH: ./test/t1.ibd
```

## Notes

- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.

## 5.19 The INFORMATION\_SCHEMA INNODB\_SYS\_FIELDS Table

The `INNODB_SYS_FIELDS` table provides metadata about the key columns (fields) of `InnoDB` indexes, equivalent to the information from the `SYS_FIELDS` table in the `InnoDB` data dictionary.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA System Tables](#).

The `INNODB_SYS_FIELDS` table has these columns:

- `INDEX_ID`

An identifier for the index associated with this key field; the same value as `INNODB_SYS_INDEXES.INDEX_ID`.

- `NAME`

The name of the original column from the table; the same value as `INNODB_SYS_COLUMNS.NAME`.

- `POS`

The ordinal position of the key field within the index, starting from 0 and incrementing sequentially. When a column is dropped, the remaining columns are reordered so that the sequence has no gaps.

### Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_SYS_FIELDS WHERE INDEX_ID = 117\G
***** 1. row *****
INDEX_ID: 117
  NAME: coll
   POS: 0
```

### Notes

- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA_COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.

## 5.20 The INFORMATION\_SCHEMA INNODB\_SYS\_FOREIGN Table

The `INNODB_SYS_FOREIGN` table provides metadata about `InnoDB` foreign keys, equivalent to the information from the `SYS_FOREIGN` table in the `InnoDB` data dictionary.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA System Tables](#).

The `INNODB_SYS_FOREIGN` table has these columns:

- `ID`

The name (not a numeric value) of the foreign key index, preceded by the schema (database) name (for example, `test/products_fk`).

- `FOR_NAME`

The name of the `child table` in this foreign key relationship.

- `REF_NAME`

The name of the `parent table` in this foreign key relationship.

- `N_COLS`

The number of columns in the foreign key index.

- `TYPE`

A collection of bit flags with information about the foreign key column, ORed together. 0 = `ON DELETE/UPDATE RESTRICT`, 1 = `ON DELETE CASCADE`, 2 = `ON DELETE SET NULL`, 4 = `ON UPDATE CASCADE`, 8 = `ON UPDATE SET NULL`, 16 = `ON DELETE NO ACTION`, 32 = `ON UPDATE NO ACTION`.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_SYS_FOREIGN\G
***** 1. row *****
      ID: test/fk1
FOR_NAME: test/child
REF_NAME: test/parent
  N_COLS: 1
    TYPE: 1
```

## Notes

- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA.COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.

## 5.21 The INFORMATION\_SCHEMA INNODB\_SYS\_FOREIGN\_COLS Table

The `INNODB_SYS_FOREIGN_COLS` table provides status information about the columns of InnoDB foreign keys, equivalent to the information from the `SYS_FOREIGN_COLS` table in the InnoDB data dictionary.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA System Tables](#).

The `INNODB_SYS_FOREIGN_COLS` table has these columns:

- `ID`

The foreign key index associated with this index key field, using the same value as `INNODB_SYS_FOREIGN.ID`.

- `FOR_COL_NAME`

The name of the associated column in the child table.

- `REF_COL_NAME`

The name of the associated column in the parent table.

- `POS`

The ordinal position of this key field within the foreign key index, starting from 0.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_SYS_FOREIGN_COLS WHERE ID = 'test/fk1'\G
***** 1. row *****
      ID: test/fk1
FOR_COL_NAME: parent_id
REF_COL_NAME: id
      POS: 0
```

## Notes

- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA.COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.

## 5.22 The `INFORMATION_SCHEMA.INNODB_SYS_INDEXES` Table

The `INNODB_SYS_INDEXES` table provides metadata about InnoDB indexes, equivalent to the information in the internal `SYS_INDEXES` table in the InnoDB data dictionary.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA System Tables](#).

The `INNODB_SYS_INDEXES` table has these columns:

- `INDEX_ID`

An identifier for the index. Index identifiers are unique across all the databases in an instance.

- `NAME`

The name of the index. Most indexes created implicitly by InnoDB have consistent names but the index names are not necessarily unique. Examples: `PRIMARY` for a primary key index, `GEN_CLUST_INDEX` for the index representing a primary key when one is not specified, and `ID_IND`, `FOR_IND`, and `REF_IND` for foreign key constraints.

- `TABLE_ID`

An identifier representing the table associated with the index; the same value as `INNODB_SYS_TABLES.TABLE_ID`.

- `TYPE`

A numeric value derived from bit-level information that identifies the index type. 0 = nonunique secondary index; 1 = automatically generated clustered index (`GEN_CLUST_INDEX`); 2 = unique nonclustered index; 3 = clustered index; 32 = full-text index

- `N_FIELDS`

The number of columns in the index key. For `GEN_CLUST_INDEX` indexes, this value is 0 because the index is created using an artificial value rather than a real table column.

- `PAGE_NO`

The root page number of the index B-tree. For full-text indexes, the `PAGE_NO` column is unused and set to -1 (`FIL_NULL`) because the full-text index is laid out in several B-trees (auxiliary tables).

- `SPACE`

An identifier for the tablespace where the index resides. 0 means the InnoDB system tablespace. Any other number represents a table created with a separate `.ibd` file in `file-per-table` mode. This identifier stays the same after a `TRUNCATE TABLE` statement. Because all indexes for a table reside in the same tablespace as the table, this value is not necessarily unique.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_SYS_INDEXES WHERE TABLE_ID = 74\G
***** 1. row *****
INDEX_ID: 116
NAME: GEN_CLUST_INDEX
TABLE_ID: 74
```

```

TYPE: 1
N_FIELDS: 0
PAGE_NO: 3
SPACE: 60
***** 2. row *****
INDEX_ID: 117
NAME: i1
TABLE_ID: 74
TYPE: 0
N_FIELDS: 1
PAGE_NO: 4
SPACE: 60

```

## Notes

- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.

## 5.23 The `INFORMATION_SCHEMA INNODB_SYS_TABLES` Table

The `INNODB_SYS_TABLES` table provides metadata about `InnoDB` tables, equivalent to the information from the `SYS_TABLES` table in the `InnoDB` data dictionary.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA System Tables](#).

The `INNODB_SYS_TABLES` table has these columns:

- `TABLE_ID`  
An identifier for the `InnoDB` table. This value is unique across all databases in the instance.
- `NAME`  
The name of the table, preceded by the schema (database) name where appropriate (for example, `test/t1`). Names of databases and user tables are in the same case as they were originally defined, possibly influenced by the `lower_case_table_names` setting.
- `FLAG`  
A numeric value that represents bit-level information about table format and storage characteristics.
- `N_COLS`  
The number of columns in the table. The number reported includes three hidden columns that are created by `InnoDB` (`DB_ROW_ID`, `DB_TRX_ID`, and `DB_ROLL_PTR`).
- `SPACE`  
An identifier for the tablespace where the table resides. 0 means the `InnoDB system tablespace`. Any other number represents a table created in `file-per-table` mode with a separate `.ibd` file. This identifier stays the same after a `TRUNCATE TABLE` statement. Other than the zero value, this identifier is unique for tables across all the databases in the instance.
- `FILE_FORMAT`  
The table's file format (`Antelope` or `Barracuda`).
- `ROW_FORMAT`  
The table's row format (`Compact`, `Redundant`, `Dynamic`, or `Compressed`).
- `ZIP_PAGE_SIZE`

The zip page size. Applies only to tables with a row format of [Compressed](#).

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_SYS_TABLES WHERE TABLE_ID = 74\G
***** 1. row *****
TABLE_ID: 74
NAME: test/t1
FLAG: 1
N_COLS: 6
SPACE: 60
FILE_FORMAT: Antelope
ROW_FORMAT: Compact
ZIP_PAGE_SIZE: 0
```

## Notes

- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.

## 5.24 The INFORMATION\_SCHEMA INNODB\_SYS\_TABLESPACES Table

The [INNODB\\_SYS\\_TABLESPACES](#) table provides metadata about [InnoDB](#) tablespaces, equivalent to the information in the [SYS\\_TABLESPACES](#) table in the [InnoDB](#) data dictionary.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA System Tables](#).

The [INNODB\\_SYS\\_TABLESPACES](#) table has these columns:

- [SPACE](#)

The tablespace ID.

- [NAME](#)

The schema (database) and table name.

- [FLAG](#)

A numeric value that represents bit-level information about tablespace format and storage characteristics.

- [FILE\\_FORMAT](#)

The tablespace file format (for example, [Antelope](#) or [Barracuda](#)). The data in this field is interpreted from the tablespace flags information that resides in the [.ibd file](#). For more information about [InnoDB](#) file formats, see [InnoDB File-Format Management](#).

- [ROW\\_FORMAT](#)

The tablespace row format ([Compact](#) or [Redundant](#), [Dynamic](#), or [Compressed](#)). The data in this column is interpreted from the tablespace flags information that resides in the [.ibd file](#).

- [PAGE\\_SIZE](#)

The tablespace page size. The data in this column is interpreted from the tablespace flags information that resides in the [.ibd file](#).

- [ZIP\\_PAGE\\_SIZE](#)

The tablespace zip page size. The data in this column is interpreted from the tablespace flags information that resides in the `.ibd` file.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_SYS_TABLESPACES WHERE SPACE = 57\G
***** 1. row *****
      SPACE: 57
      NAME: test/t1
      FLAG: 0
      FILE_FORMAT: Antelope
      ROW_FORMAT: Compact or Redundant
      PAGE_SIZE: 16384
      ZIP_PAGE_SIZE: 0
```

## Notes

- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA.COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.
- Because tablespace flags are always zero for all Antelope file formats (unlike table flags), there is no way to determine from this flag integer if the tablespace row format is Redundant or Compact. As a result, the possible values for the `ROW_FORMAT` field are “Compact or Redundant”, “Compressed”, or “Dynamic.”

## 5.25 The INFORMATION\_SCHEMA INNODB\_SYS\_TABLESTATS View

The `INNODB_SYS_TABLESTATS` table provides a view of low-level status information about `InnoDB` tables. This data is used by the MySQL optimizer to calculate which index to use when querying an `InnoDB` table. This information is derived from in-memory data structures rather than data stored on disk. There is no corresponding internal `InnoDB` system table.

`InnoDB` tables are represented in this view if they have been opened since the last server restart and have not aged out of the table cache. Tables for which persistent stats are available are always represented in this view.

Table statistics are updated only for `DELETE` or `UPDATE` operations that modify indexed columns. Statistics are not updated by operations that modify only nonindexed columns.

`ANALYZE TABLE` clears table statistics and sets the `STATS_INITIALIZED` column to `Uninitialized`. Statistics are collected again the next time the table is accessed.

For related usage information and examples, see [InnoDB INFORMATION\\_SCHEMA System Tables](#).

The `INNODB_SYS_TABLESTATS` table has these columns:

- `TABLE_ID`

An identifier representing the table for which statistics are available; the same value as `INNODB_SYS_TABLES.TABLE_ID`.

- `NAME`

The name of the table; the same value as `INNODB_SYS_TABLES.NAME`.

- `STATS_INITIALIZED`

The value is `Initialized` if the statistics are already collected, `Uninitialized` if not.

- [NUM\\_ROWS](#)

The current estimated number of rows in the table. Updated after each DML operation. The value could be imprecise if uncommitted transactions are inserting into or deleting from the table.

- [CLUST\\_INDEX\\_SIZE](#)

The number of pages on disk that store the clustered index, which holds the [InnoDB](#) table data in primary key order. This value might be null if no statistics are collected yet for the table.

- [OTHER\\_INDEX\\_SIZE](#)

The number of pages on disk that store all secondary indexes for the table. This value might be null if no statistics are collected yet for the table.

- [MODIFIED\\_COUNTER](#)

The number of rows modified by DML operations, such as [INSERT](#), [UPDATE](#), [DELETE](#), and also cascade operations from foreign keys. This column is reset each time table statistics are recalculated

- [AUTOINC](#)

The next number to be issued for any auto-increment-based operation. The rate at which the [AUTOINC](#) value changes depends on how many times auto-increment numbers have been requested and how many numbers are granted per request.

- [REF\\_COUNT](#)

When this counter reaches zero, the table metadata can be evicted from the table cache.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_SYS_TABLESTATS where TABLE_ID = 71\G
***** 1. row *****
      TABLE_ID: 71
          NAME: test/t1
STATS_INITIALIZED: Initialized
      NUM_ROWS: 1
CLUST_INDEX_SIZE: 1
OTHER_INDEX_SIZE: 0
MODIFIED_COUNTER: 1
          AUTOINC: 0
      REF_COUNT: 1
```

## Notes

- This table is useful primarily for expert-level performance monitoring, or when developing performance-related extensions for MySQL.
- You must have the [PROCESS](#) privilege to query this table.
- Use the [INFORMATION\\_SCHEMA COLUMNS](#) table or the [SHOW COLUMNS](#) statement to view additional information about the columns of this table, including data types and default values.

## 5.26 The INFORMATION\_SCHEMA INNODB\_TRX Table

The [INNODB\\_TRX](#) table provides information about every transaction currently executing inside [InnoDB](#), including whether the transaction is waiting for a lock, when the transaction started, and the SQL statement the transaction is executing, if any.

For usage information, see [Using InnoDB Transaction and Locking Information](#).

The [INNODB\\_TRX](#) table has these columns:

- [TRX\\_ID](#)

A unique transaction ID number, internal to [InnoDB](#). (Starting in MySQL 5.6, these IDs are not created for transactions that are read only and nonlocking. For details, see [Optimizing InnoDB Read-Only Transactions](#).)

- [TRX\\_WEIGHT](#)

The weight of a transaction, reflecting (but not necessarily the exact count of) the number of rows altered and the number of rows locked by the transaction. To resolve a deadlock, [InnoDB](#) selects the transaction with the smallest weight as the “victim” to roll back. Transactions that have changed nontransactional tables are considered heavier than others, regardless of the number of altered and locked rows.

- [TRX\\_STATE](#)

The transaction execution state. Permitted values are [RUNNING](#), [LOCK WAIT](#), [ROLLING BACK](#), and [COMMITTING](#).

- [TRX\\_STARTED](#)

The transaction start time.

- [TRX\\_REQUESTED\\_LOCK\\_ID](#)

The ID of the lock the transaction is currently waiting for, if [TRX\\_STATE](#) is [LOCK WAIT](#); otherwise [NULL](#). To obtain details about the lock, join this column with the [LOCK\\_ID](#) column of the [INNODB\\_LOCKS](#) table.

- [TRX\\_WAIT\\_STARTED](#)

The time when the transaction started waiting on the lock, if [TRX\\_STATE](#) is [LOCK WAIT](#); otherwise [NULL](#).

- [TRX\\_MYSQL\\_THREAD\\_ID](#)

The MySQL thread ID. To obtain details about the thread, join this column with the [ID](#) column of the [INFORMATION\\_SCHEMA PROCESSLIST](#) table, but see [Persistence and Consistency of InnoDB Transaction and Locking Information](#).

- [TRX\\_QUERY](#)

The SQL statement that is being executed by the transaction.

- [TRX\\_OPERATION\\_STATE](#)

The transaction's current operation, if any; otherwise [NULL](#).

- [TRX\\_TABLES\\_IN\\_USE](#)

The number of [InnoDB](#) tables used while processing the current SQL statement of this transaction.

- [TRX\\_TABLES\\_LOCKED](#)

The number of [InnoDB](#) tables that the current SQL statement has row locks on. (Because these are row locks, not table locks, the tables can usually still be read from and written to by multiple transactions, despite some rows being locked.)

- [TRX\\_LOCK\\_STRUCTS](#)

The number of locks reserved by the transaction.

- [TRX\\_LOCK\\_MEMORY\\_BYTES](#)

The total size taken up by the lock structures of this transaction in memory.

- [TRX\\_ROWS\\_LOCKED](#)

The approximate number of rows locked by this transaction. The value might include delete-marked rows that are physically present but not visible to the transaction.

- [TRX\\_ROWS\\_MODIFIED](#)

The number of modified and inserted rows in this transaction.

- [TRX\\_CONCURRENCY\\_TICKETS](#)

A value indicating how much work the current transaction can do before being swapped out, as specified by the `innodb_concurrency_tickets` system variable.

- [TRX\\_ISOLATION\\_LEVEL](#)

The isolation level of the current transaction.

- [TRX\\_UNIQUE\\_CHECKS](#)

Whether unique checks are turned on or off for the current transaction. For example, they might be turned off during a bulk data load.

- [TRX\\_FOREIGN\\_KEY\\_CHECKS](#)

Whether foreign key checks are turned on or off for the current transaction. For example, they might be turned off during a bulk data load.

- [TRX\\_LAST\\_FOREIGN\\_KEY\\_ERROR](#)

The detailed error message for the last foreign key error, if any; otherwise `NULL`.

- [TRX\\_ADAPTIVE\\_HASH\\_LATCHED](#)

Whether the adaptive hash index is locked by the current transaction. (Only a single transaction at a time can modify the adaptive hash index.)

- [TRX\\_ADAPTIVE\\_HASH\\_TIMEOUT](#)

Whether to relinquish the search latch immediately for the adaptive hash index, or reserve it across calls from MySQL. When there is no adaptive hash index contention, this value remains zero and statements reserve the latch until they finish. During times of contention, it counts down to zero, and statements release the latch immediately after each row lookup.

- [TRX\\_IS\\_READ\\_ONLY](#)

A value of 1 indicates the transaction is read only.

- [TRX\\_AUTOCOMMIT\\_NON\\_LOCKING](#)

A value of 1 indicates the transaction is a `SELECT` statement that does not use the `FOR UPDATE` or `LOCK IN SHARED MODE` clauses, and is executing with `autocommit` enabled so that the transaction contains only this one statement. When this column and `TRX_IS_READ_ONLY` are both 1, `InnoDB` optimizes the transaction to reduce the overhead associated with transactions that change table data.

## Example

```
mysql> SELECT * FROM INFORMATION_SCHEMA.INNODB_TRX\G
***** 1. row *****
```

```
      trx_id: 3298
      trx_state: RUNNING
      trx_started: 2014-11-19 13:54:39
trx_requested_lock_id: NULL
      trx_wait_started: NULL
      trx_weight: 316436
      trx_mysql_thread_id: 2
      trx_query: DELETE FROM employees.salaries WHERE salary > 65000
      trx_operation_state: updating or deleting
      trx_tables_in_use: 1
      trx_tables_locked: 1
      trx_lock_structs: 1621
      trx_lock_memory_bytes: 243240
      trx_rows_locked: 759343
      trx_rows_modified: 314815
trx_concurrency_tickets: 0
      trx_isolation_level: REPEATABLE READ
      trx_unique_checks: 1
      trx_foreign_key_checks: 1
trx_last_foreign_key_error: NULL
      trx_adaptive_hash_latched: 0
      trx_adaptive_hash_timeout: 10000
      trx_is_read_only: 0
      trx_autocommit_non_locking: 0
```

## Notes

- Use this table to help diagnose performance problems that occur during times of heavy concurrent load. Its contents are updated as described in [Persistence and Consistency of InnoDB Transaction and Locking Information](#).
- You must have the `PROCESS` privilege to query this table.
- Use the `INFORMATION_SCHEMA.COLUMNS` table or the `SHOW COLUMNS` statement to view additional information about the columns of this table, including data types and default values.



---

# Chapter 6 INFORMATION\_SCHEMA NDB Cluster Tables

## Table of Contents

6.1 The INFORMATION_SCHEMA FILES Table .....	85
6.2 The INFORMATION_SCHEMA ndb_transid_mysql_connection_map Table .....	90

The following sections provide information about `INFORMATION_SCHEMA` tables which are specific to NDB Cluster. (The `FILES` table is available in standard MySQL 5.6 but is not used there.) The `ndb_transid_mysql_connection_map` table is implemented as an `INFORMATION_SCHEMA` plugin available only in NDB Cluster binaries or source, and does not exist in MySQL Server 5.6.

Additional statistical and other data about NDB Cluster transactions, operations, threads, blocks, and other aspects of performance can be obtained from the tables in the `ndbinfo` database. For information about these tables, see [ndbinfo: The NDB Cluster Information Database](#).

## 6.1 The INFORMATION\_SCHEMA FILES Table

The `FILES` table provides information about the files in which MySQL `NDB` Disk Data tables are stored.

The `FILES` table has these columns:

- `FILE_ID`  
A file identifier. `FILE_ID` column values are auto-generated.
- `FILE_NAME`  
The name of an `UNDO` log file created by `CREATE LOGFILE GROUP` or `ALTER LOGFILE GROUP`, or of a data file created by `CREATE TABLESPACE` or `ALTER TABLESPACE`.
- `FILE_TYPE`  
One of the values `UNDO LOG`, `DATAFILE`, or `TABLESPACE`.
- `TABLESPACE_NAME`  
The name of the tablespace with which the file is associated.
- `TABLE_CATALOG`  
This value is always empty.
- `TABLE_SCHEMA`  
This value is always `NULL`.
- `TABLE_NAME`  
The name of the Disk Data table with which the file is associated, if any.
- `LOGFILE_GROUP_NAME`  
The name of the log file group to which the log file or data file belongs.
- `LOGFILE_GROUP_NUMBER`  
For an `UNDO` log file, the auto-generated ID number of the log file group to which the log file belongs.
- `ENGINE`  
For an NDB Cluster Disk Data log file or data file, this value always `NDB` or `NDBCLUSTER`.

- `FULLTEXT_KEYS`

For an NDB Cluster Disk Data log file or data file, this value is always empty.

- `DELETED_ROWS`

This value is always `NULL`.

- `UPDATE_COUNT`

This value is always `NULL`.

- `FREE_EXTENTS`

The number of extents which have not yet been used by the file.

- `TOTAL_EXTENTS`

The total number of extents allocated to the file.

- `EXTENT_SIZE`

The size of an extent for the file in bytes.

- `INITIAL_SIZE`

The size of the file in bytes. This is the same value that was used in the `INITIAL_SIZE` clause of the `CREATE LOGFILE GROUP`, `ALTER LOGFILE GROUP`, `CREATE TABLESPACE`, or `ALTER TABLESPACE` statement used to create the file.

- `MAXIMUM_SIZE`

For NDB Cluster Disk Data files, this value is always the same as the `INITIAL_SIZE` value.

- `AUTOEXTEND_SIZE`

For NDB Cluster Disk Data files, this value is always empty.

- `CREATION_TIME`

The date and time when the file was created.

- `LAST_UPDATE_TIME`

The date and time when the file was last modified.

- `LAST_ACCESS_TIME`

The date and time when the file was last accessed by the server.

- `RECOVER_TIME`

For NDB Cluster Disk Data files, this value is always `0`.

- `TRANSACTION_COUNTER`

For NDB Cluster Disk Data files, this value is always `0`.

- `VERSION`

For NDB Cluster Disk Data files, this value is always `NULL`.

- `ROW_FORMAT`

For NDB Cluster Disk Data files, this value is always `NULL`.

- [TABLE\\_ROWS](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [AVG\\_ROW\\_LENGTH](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [DATA\\_LENGTH](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [MAX\\_DATA\\_LENGTH](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [INDEX\\_LENGTH](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [DATA\\_FREE](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [CREATE\\_TIME](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [UPDATE\\_TIME](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [CHECK\\_TIME](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [CHECKSUM](#)  
For NDB Cluster Disk Data files, this value is always [NULL](#).
- [STATUS](#)  
For NDB Cluster Disk Data files, this value is always [NORMAL](#).
- [EXTRA](#)  
For NDB Cluster Disk Data files, the [EXTRA](#) column shows which data node the file belongs to (each data node having its own copy), as well as the size of its undo buffer. Suppose that you use this statement on an NDB Cluster with four data nodes:

```
CREATE LOGFILE GROUP mygroup
  ADD UNDOFILE 'new_undo.dat'
  INITIAL_SIZE 2G
  ENGINE NDB;
```

After running the [CREATE LOGFILE GROUP](#) statement successfully, you should see a result similar to the one shown here for this query against the [FILES](#) table:

```
mysql> SELECT LOGFILE_GROUP_NAME, FILE_TYPE, EXTRA
  FROM INFORMATION_SCHEMA.FILES
  WHERE FILE_NAME = 'new_undo.dat';
```

LOGFILE_GROUP_NAME	FILE_TYPE	EXTRA
mygroup	UNDO LOG	CLUSTER_NODE=5;UNDO_BUFFER_SIZE=8388608
mygroup	UNDO LOG	CLUSTER_NODE=6;UNDO_BUFFER_SIZE=8388608

mygroup	UNDO LOG	CLUSTER_NODE=7;UNDO_BUFFER_SIZE=8388608
mygroup	UNDO LOG	CLUSTER_NODE=8;UNDO_BUFFER_SIZE=8388608

## Notes

- The `FILES` table is a nonstandard `INFORMATION_SCHEMA` table.

## NDB Notes

- This table provides information about Disk Data *files* only; you cannot use it for determining disk space allocation or availability for individual `NDB` tables. However, it is possible to see how much space is allocated for each `NDB` table having data stored on disk—as well as how much remains available for storage of data on disk for that table—using `ndb_desc`. For more information, see [ndb\\_desc — Describe NDB Tables](#).
- The `CREATION_TIME`, `LAST_UPDATE_TIME`, and `LAST_ACCESSED` values are as reported by the operating system, and are not supplied by the `NDB` storage engine. Where no value is provided by the operating system, these columns display `0000-00-00 00:00:00`.
- The difference between the `TOTAL_EXTENTS` and `FREE_EXTENTS` columns is the number of extents currently in use by the file:

```
SELECT TOTAL_EXTENTS - FREE_EXTENTS AS extents_used
FROM INFORMATION_SCHEMA.FILES
WHERE FILE_NAME = 'myfile.dat';
```

To approximate the amount of disk space in use by the file, multiply that difference by the value of the `EXTENT_SIZE` column, which gives the size of an extent for the file in bytes:

```
SELECT (TOTAL_EXTENTS - FREE_EXTENTS) * EXTENT_SIZE AS bytes_used
FROM INFORMATION_SCHEMA.FILES
WHERE FILE_NAME = 'myfile.dat';
```

Similarly, you can estimate the amount of space that remains available in a given file by multiplying `FREE_EXTENTS` by `EXTENT_SIZE`:

```
SELECT FREE_EXTENTS * EXTENT_SIZE AS bytes_free
FROM INFORMATION_SCHEMA.FILES
WHERE FILE_NAME = 'myfile.dat';
```

### Important

The byte values produced by the preceding queries are approximations only, and their precision is inversely proportional to the value of `EXTENT_SIZE`. That is, the larger `EXTENT_SIZE` becomes, the less accurate the approximations are.

It is also important to remember that once an extent is used, it cannot be freed again without dropping the data file of which it is a part. This means that deletes from a Disk Data table do *not* release disk space.

The extent size can be set in a `CREATE TABLESPACE` statement. For more information, see [CREATE TABLESPACE Statement](#).

- An additional row is present in the `FILES` table following the creation of a logfile group. This row has `NULL` for the value of the `FILE_NAME` column. For this row, the value of the `FILE_ID` column is always `0`, that of the `FILE_TYPE` column is always `UNDO LOG`, and that of the `STATUS` column is always `NORMAL`. The value of the `ENGINE` column is always `NDBCLUSTER`.

The `FREE_EXTENTS` column in this row shows the total number of free extents available to all undo files belonging to a given log file group whose name and number are shown in the `LOGFILE_GROUP_NAME` and `LOGFILE_GROUP_NUMBER` columns, respectively.

Suppose there are no existing log file groups on your NDB Cluster, and you create one using the following statement:

```
mysql> CREATE LOGFILE GROUP lg1
      ADD UNDOFILE 'undofile.dat'
      INITIAL_SIZE = 16M
      UNDO_BUFFER_SIZE = 1M
      ENGINE = NDB;
```

You can now see this `NULL` row when you query the `FILES` table:

```
mysql> SELECT DISTINCT
      FILE_NAME AS File,
      FREE_EXTENTS AS Free,
      TOTAL_EXTENTS AS Total,
      EXTENT_SIZE AS Size,
      INITIAL_SIZE AS Initial
      FROM INFORMATION_SCHEMA.FILES;
```

File	Free	Total	Size	Initial
undofile.dat	NULL	4194304	4	16777216
NULL	4184068	NULL	4	NULL

The total number of free extents available for undo logging is always somewhat less than the sum of the `TOTAL_EXTENTS` column values for all undo files in the log file group due to overhead required for maintaining the undo files. This can be seen by adding a second undo file to the log file group, then repeating the previous query against the `FILES` table:

```
mysql> ALTER LOGFILE GROUP lg1
      ADD UNDOFILE 'undofile02.dat'
      INITIAL_SIZE = 4M
      ENGINE = NDB;
```

```
mysql> SELECT DISTINCT
      FILE_NAME AS File,
      FREE_EXTENTS AS Free,
      TOTAL_EXTENTS AS Total,
      EXTENT_SIZE AS Size,
      INITIAL_SIZE AS Initial
      FROM INFORMATION_SCHEMA.FILES;
```

File	Free	Total	Size	Initial
undofile.dat	NULL	4194304	4	16777216
undofile02.dat	NULL	1048576	4	4194304
NULL	5223944	NULL	4	NULL

The amount of free space in bytes which is available for undo logging by Disk Data tables using this log file group can be approximated by multiplying the number of free extents by the initial size:

```
mysql> SELECT
      FREE_EXTENTS AS 'Free Extents',
      FREE_EXTENTS * EXTENT_SIZE AS 'Free Bytes'
      FROM INFORMATION_SCHEMA.FILES
      WHERE LOGFILE_GROUP_NAME = 'lg1'
      AND FILE_NAME IS NULL;
```

Free Extents	Free Bytes
5223944	20895776

If you create an NDB Cluster Disk Data table and then insert some rows into it, you can see approximately how much space remains for undo logging afterward, for example:

```
mysql> CREATE TABLESPACE ts1
```

```

ADD DATAFILE 'data1.dat'
USE LOGFILE GROUP lg1
INITIAL_SIZE 512M
ENGINE = NDB;
mysql> CREATE TABLE dd (
  c1 INT NOT NULL PRIMARY KEY,
  c2 INT,
  c3 DATE
)
  TABLESPACE ts1 STORAGE DISK
  ENGINE = NDB;
mysql> INSERT INTO dd VALUES
  (NULL, 1234567890, '2007-02-02'),
  (NULL, 1126789005, '2007-02-03'),
  (NULL, 1357924680, '2007-02-04'),
  (NULL, 1642097531, '2007-02-05');
mysql> SELECT
  FREE_EXTENTS AS 'Free Extents',
  FREE_EXTENTS * EXTENT_SIZE AS 'Free Bytes'
FROM INFORMATION_SCHEMA.FILES
WHERE LOGFILE_GROUP_NAME = 'lg1'
  AND FILE_NAME IS NULL;
+-----+-----+
| Free Extents | Free Bytes |
+-----+-----+
|          5207565 |      20830260 |
+-----+-----+

```

- An additional row is present in the `FILES` table for any NDB Cluster tablespace, whether or not any data files are associated with the tablespace. This row has `NULL` for the value of the `FILE_NAME` column. For this row, the value of the `FILE_ID` column is always 0, that of the `FILE_TYPE` column is always `TABLESPACE`, and that of the `STATUS` column is always `NORMAL`. The value of the `ENGINE` column is always `NDBCLUSTER`.
- For additional information, and examples of creating and dropping NDB Cluster Disk Data objects, see [NDB Cluster Disk Data Tables](#).

## 6.2 The INFORMATION\_SCHEMA ndb\_transid\_mysql\_connection\_map Table

The `ndb_transid_mysql_connection_map` table provides a mapping between NDB transactions, NDB transaction coordinators, and MySQL Servers attached to an NDB Cluster as API nodes. This information is used when populating the `server_operations` and `server_transactions` tables of the `ndbinfo` NDB Cluster information database.

The `ndb_transid_mysql_connection_map` table has these columns:

- `mysql_connection_id`  
The MySQL server connection ID.
- `node_id`  
The transaction coordinator node ID.
- `ndb_transid`  
The NDB transaction ID.

### Notes

The `mysql_connection_id` value is the same as the connection or session ID shown in the output of `SHOW PROCESSLIST`.

There are no `SHOW` statements associated with this table.

This is a nonstandard table, specific to NDB Cluster. It is implemented as an `INFORMATION_SCHEMA` plugin; you can verify that it is supported by checking the output of `SHOW PLUGINS`. If `ndb_transid_mysql_connection_map` support is enabled, the output from this statement includes a plugin having this name, of type `INFORMATION_SCHEMA`, and having status `ACTIVE`, as shown here (using emphasized text):

```
mysql> SHOW PLUGINS;
```

Name	Status	Type	Library	License
binlog	ACTIVE	STORAGE ENGINE	NULL	GPL
mysql_native_password	ACTIVE	AUTHENTICATION	NULL	GPL
mysql_old_password	ACTIVE	AUTHENTICATION	NULL	GPL
CSV	ACTIVE	STORAGE ENGINE	NULL	GPL
MEMORY	ACTIVE	STORAGE ENGINE	NULL	GPL
MRG_MYISAM	ACTIVE	STORAGE ENGINE	NULL	GPL
MyISAM	ACTIVE	STORAGE ENGINE	NULL	GPL
PERFORMANCE_SCHEMA	ACTIVE	STORAGE ENGINE	NULL	GPL
BLACKHOLE	ACTIVE	STORAGE ENGINE	NULL	GPL
ARCHIVE	ACTIVE	STORAGE ENGINE	NULL	GPL
ndbcluster	ACTIVE	STORAGE ENGINE	NULL	GPL
ndbinfo	ACTIVE	STORAGE ENGINE	NULL	GPL
<i>ndb_transid_mysql_connection_map</i>	<i>ACTIVE</i>	<i>INFORMATION_SCHEMA</i>	<i>NULL</i>	<i>GPL</i>
InnoDB	ACTIVE	STORAGE ENGINE	NULL	GPL
INNODB_TRX	ACTIVE	INFORMATION_SCHEMA	NULL	GPL
INNODB_LOCKS	ACTIVE	INFORMATION_SCHEMA	NULL	GPL
INNODB_LOCK_WAITS	ACTIVE	INFORMATION_SCHEMA	NULL	GPL
INNODB_CMP	ACTIVE	INFORMATION_SCHEMA	NULL	GPL
INNODB_CMP_RESET	ACTIVE	INFORMATION_SCHEMA	NULL	GPL
INNODB_CMPMEM	ACTIVE	INFORMATION_SCHEMA	NULL	GPL
INNODB_CMPMEM_RESET	ACTIVE	INFORMATION_SCHEMA	NULL	GPL
partition	ACTIVE	STORAGE ENGINE	NULL	GPL

```
22 rows in set (0.00 sec)
```

The plugin is enabled by default. You can disable it (or force the server not to run unless the plugin starts) by starting the server with the `--ndb-transid-mysql-connection-map` option. If the plugin is disabled, the status is shown by `SHOW PLUGINS` as `DISABLED`. The plugin cannot be enabled or disabled at runtime.

Although the names of this table and its columns are displayed using lowercase, you can use uppercase or lowercase when referring to them in SQL statements.

For this table to be created, the MySQL Server must be a binary supplied with the NDB Cluster distribution, or one built from the NDB Cluster sources with `NDB` storage engine support enabled. It is not available in the standard MySQL 5.6 Server.



---

# Chapter 7 INFORMATION\_SCHEMA Thread Pool Tables

## Table of Contents

7.1 INFORMATION_SCHEMA Thread Pool Table Reference .....	93
7.2 The INFORMATION_SCHEMA TP_THREAD_GROUP_STATE Table .....	93
7.3 The INFORMATION_SCHEMA TP_THREAD_GROUP_STATS Table .....	95
7.4 The INFORMATION_SCHEMA TP_THREAD_STATE Table .....	96

The following sections describe the [INFORMATION\\_SCHEMA](#) tables associated with the thread pool plugin (see [MySQL Enterprise Thread Pool](#)). They provide information about thread pool operation:

- [TP\\_THREAD\\_GROUP\\_STATE](#): Information about thread pool thread group states
- [TP\\_THREAD\\_GROUP\\_STATS](#): Thread group statistics
- [TP\\_THREAD\\_STATE](#): Information about thread pool thread states

Rows in these tables represent snapshots in time. In the case of [TP\\_THREAD\\_STATE](#), all rows for a thread group comprise a snapshot in time. Thus, the MySQL server holds the mutex of the thread group while producing the snapshot. But it does not hold mutexes on all thread groups at the same time, to prevent a statement against [TP\\_THREAD\\_STATE](#) from blocking the entire MySQL server.

The thread pool [INFORMATION\\_SCHEMA](#) tables are implemented by individual plugins and the decision whether to load one can be made independently of the others (see [Thread Pool Installation](#)). However, the content of all the tables depends on the thread pool plugin being enabled. If a table plugin is enabled but the thread pool plugin is not, the table becomes visible and can be accessed but is empty.

## 7.1 INFORMATION\_SCHEMA Thread Pool Table Reference

The following table summarizes [INFORMATION\\_SCHEMA](#) thread pool tables. For greater detail, see the individual table descriptions.

**Table 7.1 INFORMATION\_SCHEMA Thread Pool Tables**

Name	Description
<a href="#">TP_THREAD_GROUP_STATE</a>	Thread pool thread group states
<a href="#">TP_THREAD_GROUP_STATS</a>	Thread pool thread group statistics
<a href="#">TP_THREAD_STATE</a>	Thread pool thread information

## 7.2 The INFORMATION\_SCHEMA TP\_THREAD\_GROUP\_STATE Table

The [TP\\_THREAD\\_GROUP\\_STATE](#) table has one row per thread group in the thread pool. Each row provides information about the current state of a group.

The [TP\\_THREAD\\_GROUP\\_STATE](#) table has these columns:

- [TP\\_GROUP\\_ID](#)

The thread group ID. This is a unique key within the table.

- [CONSUMER\\_THREADS](#)

The number of consumer threads. There is at most one thread ready to start executing if the active threads become stalled or blocked.

- [RESERVE\\_THREADS](#)

The number of threads in the reserved state. This means that they are not started until there is a need to wake a new thread and there is no consumer thread. This is where most threads end up when the thread group has created more threads than needed for normal operation. Often a thread group needs additional threads for a short while and then does not need them again for a while. In this case, they go into the reserved state and remain until needed again. They take up some extra memory resources, but no extra computing resources.

- [CONNECT\\_THREAD\\_COUNT](#)

The number of threads that are processing or waiting to process connection initialization and authentication. There can be a maximum of four connection threads per thread group; these threads expire after a period of inactivity.

This column was added in MySQL 5.6.36.

- [CONNECTION\\_COUNT](#)

The number of connections using this thread group.

- [QUEUED\\_QUERIES](#)

The number of statements waiting in the high-priority queue.

- [QUEUED\\_TRANSACTIONS](#)

The number of statements waiting in the low-priority queue. These are the initial statements for transactions that have not started, so they also represent queued transactions.

- [STALL\\_LIMIT](#)

The value of the `thread_pool_stall_limit` system variable for the thread group. This is the same value for all thread groups.

- [PRIO\\_KICKUP\\_TIMER](#)

The value of the `thread_pool_prio_kickup_timer` system variable for the thread group. This is the same value for all thread groups.

- [ALGORITHM](#)

The value of the `thread_pool_algorithm` system variable for the thread group. This is the same value for all thread groups.

- [THREAD\\_COUNT](#)

The number of threads started in the thread pool as part of this thread group.

- [ACTIVE\\_THREAD\\_COUNT](#)

The number of threads active in executing statements.

- [STALLED\\_THREAD\\_COUNT](#)

The number of stalled statements in the thread group. A stalled statement could be executing, but from a thread pool perspective it is stalled and making no progress. A long-running statement quickly ends up in this category.

- [WAITING\\_THREAD\\_NUMBER](#)

If there is a thread handling the polling of statements in the thread group, this specifies the thread number within this thread group. It is possible that this thread could be executing a statement.

- [OLDEST\\_QUEUED](#)

How long in milliseconds the oldest queued statement has been waiting for execution.

- [MAX\\_THREAD\\_IDS\\_IN\\_GROUP](#)

The maximum thread ID of the threads in the group. This is the same as `MAX(TP_THREAD_NUMBER)` for the threads when selected from the `TP_THREAD_STATE` table. That is, these two queries are equivalent:

```
SELECT TP_GROUP_ID, MAX_THREAD_IDS_IN_GROUP
FROM TP_THREAD_GROUP_STATE;
SELECT TP_GROUP_ID, MAX(TP_THREAD_NUMBER)
FROM TP_THREAD_STATE GROUP BY TP_GROUP_ID;
```

## 7.3 The INFORMATION\_SCHEMA TP\_THREAD\_GROUP\_STATS Table

The `TP_THREAD_GROUP_STATS` table reports statistics per thread group. There is one row per group.

The `TP_THREAD_GROUP_STATS` table has these columns:

- [TP\\_GROUP\\_ID](#)

The thread group ID. This is a unique key within the table.

- [CONNECTIONS\\_STARTED](#)

The number of connections started.

- [CONNECTIONS\\_CLOSED](#)

The number of connections closed.

- [QUERIES\\_EXECUTED](#)

The number of statements executed. This number is incremented when a statement starts executing, not when it finishes.

- [QUERIES\\_QUEUED](#)

The number of statements received that were queued for execution. This does not count statements that the thread group was able to begin executing immediately without queuing, which can happen under the conditions described in [Thread Pool Operation](#).

- [THREADS\\_STARTED](#)

The number of threads started.

- [PRIO\\_KICKUPS](#)

The number of statements that have been moved from low-priority queue to high-priority queue based on the value of the `thread_pool_prio_kickup_timer` system variable. If this number increases quickly, consider increasing the value of that variable. A quickly increasing counter means that the priority system is not keeping transactions from starting too early. For `InnoDB`, this most likely means deteriorating performance due to too many concurrent transactions..

- [STALLED\\_QUERIES\\_EXECUTED](#)

The number of statements that have become defined as stalled due to executing for longer than the value of the `thread_pool_stall_limit` system variable.

- [BECOME\\_CONSUMER\\_THREAD](#)

The number of times thread have been assigned the consumer thread role.

- [BECOME\\_RESERVE\\_THREAD](#)

The number of times threads have been assigned the reserve thread role.

- [BECOME\\_WAITING\\_THREAD](#)

The number of times threads have been assigned the waiter thread role. When statements are queued, this happens very often, even in normal operation, so rapid increases in this value are normal in the case of a highly loaded system where statements are queued up.

- [WAKE\\_THREAD\\_STALL\\_CHECKER](#)

The number of times the stall check thread decided to wake or create a thread to possibly handle some statements or take care of the waiter thread role.

- [SLEEP\\_WAITS](#)

The number of [THD\\_WAIT\\_SLEEP](#) waits. These occur when threads go to sleep; for example, by calling the [SLEEP\(\)](#) function.

- [DISK\\_IO\\_WAITS](#)

The number of [THD\\_WAIT\\_DISKIO](#) waits. These occur when threads perform disk I/O that is likely to not hit the file system cache. Such waits occur when the buffer pool reads and writes data to disk, not for normal reads from and writes to files.

- [ROW\\_LOCK\\_WAITS](#)

The number of [THD\\_WAIT\\_ROW\\_LOCK](#) waits for release of a row lock by another transaction.

- [GLOBAL\\_LOCK\\_WAITS](#)

The number of [THD\\_WAIT\\_GLOBAL\\_LOCK](#) waits for a global lock to be released.

- [META\\_DATA\\_LOCK\\_WAITS](#)

The number of [THD\\_WAIT\\_META\\_DATA\\_LOCK](#) waits for a metadata lock to be released.

- [TABLE\\_LOCK\\_WAITS](#)

The number of [THD\\_WAIT\\_TABLE\\_LOCK](#) waits for a table to be unlocked that the statement needs to access.

- [USER\\_LOCK\\_WAITS](#)

The number of [THD\\_WAIT\\_USER\\_LOCK](#) waits for a special lock constructed by the user thread.

- [BINLOG\\_WAITS](#)

The number of [THD\\_WAIT\\_BINLOG\\_WAITS](#) waits for the binary log to become free.

- [GROUP\\_COMMIT\\_WAITS](#)

The number of [THD\\_WAIT\\_GROUP\\_COMMIT](#) waits. These occur when a group commit must wait for the other parties to complete their part of a transaction.

- [FSYNC\\_WAITS](#)

The number of [THD\\_WAIT\\_SYNC](#) waits for a file sync operation.

## 7.4 The INFORMATION\_SCHEMA TP\_THREAD\_STATE Table

The `TP_THREAD_STATE` table has one row per thread created by the thread pool to handle connections.

The `TP_THREAD_STATE` table has these columns:

- `TP_GROUP_ID`

The thread group ID.

- `TP_THREAD_NUMBER`

The ID of the thread within its thread group. `TP_GROUP_ID` and `TP_THREAD_NUMBER` together provide a unique key within the table.

- `PROCESS_COUNT`

The 10ms interval in which the statement that uses this thread is currently executing. 0 means no statement is executing, 1 means it is in the first 10ms, and so forth.

- `WAIT_TYPE`

The type of wait for the thread. `NULL` means the thread is not blocked. Otherwise, the thread is blocked by a call to `thd_wait_begin()` and the value specifies the type of wait. The `xxx_WAIT` columns of the `TP_THREAD_GROUP_STATS` table accumulate counts for each wait type.

The `WAIT_TYPE` value is a string that describes the type of wait, as shown in the following table.

**Table 7.2 TP\_THREAD\_STATE Table WAIT\_TYPE Values**

Wait Type	Meaning
<code>THD_WAIT_SLEEP</code>	Waiting for sleep
<code>THD_WAIT_DISKIO</code>	Waiting for Disk IO
<code>THD_WAIT_ROW_LOCK</code>	Waiting for row lock
<code>THD_WAIT_GLOBAL_LOCK</code>	Waiting for global lock
<code>THD_WAIT_META_DATA_LOCK</code>	Waiting for metadata lock
<code>THD_WAIT_TABLE_LOCK</code>	Waiting for table lock
<code>THD_WAIT_USER_LOCK</code>	Waiting for user lock
<code>THD_WAIT_BINLOG</code>	Waiting for binlog
<code>THD_WAIT_GROUP_COMMIT</code>	Waiting for group commit
<code>THD_WAIT_SYNC</code>	Waiting for fsync



---

# Chapter 8 INFORMATION\_SCHEMA Connection-Control Tables

## Table of Contents

8.1 INFORMATION_SCHEMA Connection-Control Table Reference .....	99
8.2 The INFORMATION_SCHEMA CONNECTION_CONTROL_FAILED_LOGIN_ATTEMPTS Table .....	99

The following sections describe the [INFORMATION\\_SCHEMA](#) tables associated with the [CONNECTION\\_CONTROL](#) plugin.

## 8.1 INFORMATION\_SCHEMA Connection-Control Table Reference

The following table summarizes [INFORMATION\\_SCHEMA](#) connection-control tables. For greater detail, see the individual table descriptions.

**Table 8.1 INFORMATION\_SCHEMA Connection-Control Tables**

Name	Description	Introduced
<a href="#">CONNECTION_CONTROL_FAILED_LOGIN_ATTEMPTS</a>	Current number of consecutive failed connection attempts per account	5.6.35

## 8.2 The INFORMATION\_SCHEMA CONNECTION\_CONTROL\_FAILED\_LOGIN\_ATTEMPTS Table

This table provides information about the current number of consecutive failed connection attempts per account (user/host combination). The table was added in MySQL 5.6.35.

[CONNECTION\\_CONTROL\\_FAILED\\_LOGIN\\_ATTEMPTS](#) has these columns:

- [USERHOST](#)

The user/host combination indicating an account that has failed connection attempts, in `'user_name'@'host_name'` format.

- [FAILED\\_ATTEMPTS](#)

The current number of consecutive failed connection attempts for the [USERHOST](#) value. This counts all failed attempts, regardless of whether they were delayed. The number of attempts for which the server added a delay to its response is the difference between the [FAILED\\_ATTEMPTS](#) value and the `connection_control_failed_connections_threshold` system variable value.

### Notes

- The [CONNECTION\\_CONTROL\\_FAILED\\_LOGIN\\_ATTEMPTS](#) plugin must be activated for this table to be available, and the [CONNECTION\\_CONTROL](#) plugin must be activated or the table contents are always empty. See [The Connection-Control Plugins](#).
- The table contains rows only for accounts that have had one or more consecutive failed connection attempts without a subsequent successful attempt. When an account connects successfully, its failed-connection count is reset to zero and the server removes any row corresponding to the account.

- Assigning a value to the `connection_control_failed_connections_threshold` system variable at runtime resets all accumulated failed-connection counters to zero, which causes the table to become empty.

---

# Chapter 9 INFORMATION\_SCHEMA MySQL Enterprise Firewall Tables

## Table of Contents

9.1 INFORMATION_SCHEMA Firewall Table Reference .....	101
9.2 The INFORMATION_SCHEMA MYSQL_FIREWALL_USERS Table .....	101
9.3 The INFORMATION_SCHEMA MYSQL_FIREWALL_WHITELIST Table .....	101

The following sections describe the [INFORMATION\\_SCHEMA](#) tables associated with MySQL Enterprise Firewall (see [MySQL Enterprise Firewall](#)). They provide views into the firewall in-memory data cache. These tables are available only if the appropriate firewall plugins are enabled.

## 9.1 INFORMATION\_SCHEMA Firewall Table Reference

The following table summarizes [INFORMATION\\_SCHEMA](#) firewall tables. For greater detail, see the individual table descriptions.

**Table 9.1 INFORMATION\_SCHEMA Firewall Tables**

Name	Description
<a href="#">MYSQL_FIREWALL_USERS</a>	Firewall in-memory data for account profiles
<a href="#">MYSQL_FIREWALL_WHITELIST</a>	Firewall in-memory data for account profile allowlists

## 9.2 The INFORMATION\_SCHEMA MYSQL\_FIREWALL\_USERS Table

The [MYSQL\\_FIREWALL\\_USERS](#) table provides a view into the in-memory data cache for MySQL Enterprise Firewall. It lists names and operational modes of registered firewall account profiles. It is used in conjunction with the `mysql.firewall_users` system table that provides persistent storage of firewall data; see [MySQL Enterprise Firewall Tables](#).

The [MYSQL\\_FIREWALL\\_USERS](#) table has these columns:

- [USERHOST](#)

The account profile name. Each account name has the format `user_name@host_name`.

- [MODE](#)

The current operational mode for the profile. Permitted mode values are `OFF`, `DETECTING` (as of MySQL 5.6.26), `PROTECTING`, `RECORDING`, and `RESET`. For details about their meanings, see [Firewall Concepts](#).

## 9.3 The INFORMATION\_SCHEMA MYSQL\_FIREWALL\_WHITELIST Table

The [MYSQL\\_FIREWALL\\_WHITELIST](#) table provides a view into the in-memory data cache for MySQL Enterprise Firewall. It lists allowlist rules of registered firewall account profiles. It is used in conjunction with the `mysql.firewall_whitelist` system table that provides persistent storage of firewall data; see [MySQL Enterprise Firewall Tables](#).

The `MYSQL_FIREWALL_WHITELIST` table has these columns:

- `USERHOST`

The account profile name. Each account name has the format `user_name@host_name`.

- `RULE`

A normalized statement indicating an acceptable statement pattern for the profile. A profile allowlist is the union of its rules.

---

## Chapter 10 Extensions to SHOW Statements

Some extensions to `SHOW` statements accompany the implementation of `INFORMATION_SCHEMA`:

- `SHOW` can be used to get information about the structure of `INFORMATION_SCHEMA` itself.
- Several `SHOW` statements accept a `WHERE` clause that provides more flexibility in specifying which rows to display.

`INFORMATION_SCHEMA` is an information database, so its name is included in the output from `SHOW DATABASES`. Similarly, `SHOW TABLES` can be used with `INFORMATION_SCHEMA` to obtain a list of its tables:

```
mysql> SHOW TABLES FROM INFORMATION_SCHEMA;
+-----+
| Tables_in_INFORMATION_SCHEMA |
+-----+
| CHARACTER_SETS               |
| COLLATIONS                   |
| COLLATION_CHARACTER_SET_APPLICABILITY |
| COLUMNS                     |
| COLUMN_PRIVILEGES            |
| ENGINES                      |
| EVENTS                       |
| FILES                        |
| GLOBAL_STATUS                |
| GLOBAL_VARIABLES            |
| KEY_COLUMN_USAGE             |
| PARTITIONS                   |
| PLUGINS                      |
| PROCESSLIST                  |
| REFERENTIAL_CONSTRAINTS     |
| ROUTINES                     |
| SCHEMATA                     |
| SCHEMA_PRIVILEGES           |
| SESSION_STATUS              |
| SESSION_VARIABLES           |
| STATISTICS                   |
| TABLES                      |
| TABLE_CONSTRAINTS          |
| TABLE_PRIVILEGES           |
| TRIGGERS                     |
| USER_PRIVILEGES             |
| VIEWS                        |
+-----+
```

`SHOW COLUMNS` and `DESCRIBE` can display information about the columns in individual `INFORMATION_SCHEMA` tables.

`SHOW` statements that accept a `LIKE` clause to limit the rows displayed also permit a `WHERE` clause that specifies more general conditions that selected rows must satisfy:

```
SHOW CHARACTER SET
SHOW COLLATION
SHOW COLUMNS
SHOW DATABASES
SHOW FUNCTION STATUS
SHOW INDEX
SHOW OPEN TABLES
SHOW PROCEDURE STATUS
SHOW STATUS
SHOW TABLE STATUS
SHOW TABLES
SHOW TRIGGERS
SHOW VARIABLES
```

The `WHERE` clause, if present, is evaluated against the column names displayed by the `SHOW` statement. For example, the `SHOW CHARACTER SET` statement produces these output columns:

```
mysql> SHOW CHARACTER SET;
```

Charset	Description	Default collation	Maxlen
big5	Big5 Traditional Chinese	big5_chinese_ci	2
dec8	DEC West European	dec8_swedish_ci	1
cp850	DOS West European	cp850_general_ci	1
hp8	HP West European	hp8_english_ci	1
koi8r	KOI8-R Relcom Russian	koi8r_general_ci	1
latin1	cp1252 West European	latin1_swedish_ci	1
latin2	ISO 8859-2 Central European	latin2_general_ci	1
...			

To use a `WHERE` clause with `SHOW CHARACTER SET`, you would refer to those column names. As an example, the following statement displays information about character sets for which the default collation contains the string `'japanese'`:

```
mysql> SHOW CHARACTER SET WHERE `Default collation` LIKE '%japanese%';
```

Charset	Description	Default collation	Maxlen
ujis	EUC-JP Japanese	ujis_japanese_ci	3
sjis	Shift-JIS Japanese	sjis_japanese_ci	2
cp932	SJIS for Windows Japanese	cp932_japanese_ci	2
eucjpms	UJIS for Windows Japanese	eucjpms_japanese_ci	3

This statement displays the multibyte character sets:

```
mysql> SHOW CHARACTER SET WHERE Maxlen > 1;
```

Charset	Description	Default collation	Maxlen
big5	Big5 Traditional Chinese	big5_chinese_ci	2
ujis	EUC-JP Japanese	ujis_japanese_ci	3
sjis	Shift-JIS Japanese	sjis_japanese_ci	2
euckr	EUC-KR Korean	euckr_korean_ci	2
gb2312	GB2312 Simplified Chinese	gb2312_chinese_ci	2
gbk	GBK Simplified Chinese	gbk_chinese_ci	2
utf8	UTF-8 Unicode	utf8_general_ci	3
ucs2	UCS-2 Unicode	ucs2_general_ci	2
cp932	SJIS for Windows Japanese	cp932_japanese_ci	2
eucjpms	UJIS for Windows Japanese	eucjpms_japanese_ci	3

---

# Chapter 11 MySQL 5.6 FAQ: INFORMATION\_SCHEMA

## Questions

- **11.1:** Where can I find documentation for the MySQL [INFORMATION\\_SCHEMA](#) database?
- **11.2:** Is there a discussion forum for [INFORMATION\\_SCHEMA](#)?
- **11.3:** Where can I find the ANSI SQL 2003 specification for [INFORMATION\\_SCHEMA](#)?
- **11.4:** What is the difference between the Oracle Data Dictionary and MySQL [INFORMATION\\_SCHEMA](#)?
- **11.5:** Can I add to or otherwise modify the tables found in the [INFORMATION\\_SCHEMA](#) database?

## Questions and Answers

### **11.1: Where can I find documentation for the MySQL [INFORMATION\\_SCHEMA](#) database?**

See [Chapter 1, INFORMATION\\_SCHEMA Tables](#)

### **11.2: Is there a discussion forum for [INFORMATION\\_SCHEMA](#)?**

See <https://forums.mysql.com/list.php?101>.

### **11.3: Where can I find the ANSI SQL 2003 specification for [INFORMATION\\_SCHEMA](#)?**

Unfortunately, the official specifications are not freely available. (ANSI makes them available for purchase.) However, there are books available, such as *SQL-99 Complete, Really* by Peter Gultzan and Trudy Pelzer, that provide a comprehensive overview of the standard, including [INFORMATION\\_SCHEMA](#).

### **11.4: What is the difference between the Oracle Data Dictionary and MySQL [INFORMATION\\_SCHEMA](#)?**

Both Oracle and MySQL provide metadata in tables. However, Oracle and MySQL use different table names and column names. The MySQL implementation is more similar to those found in DB2 and SQL Server, which also support [INFORMATION\\_SCHEMA](#) as defined in the SQL standard.

### **11.5: Can I add to or otherwise modify the tables found in the [INFORMATION\\_SCHEMA](#) database?**

No. Since applications may rely on a certain standard structure, this should not be modified. For this reason, *we cannot support bugs or other issues which result from modifying [INFORMATION\\_SCHEMA](#) tables or data.*

