



- Expert Verified, Online, **Free**.

Examine this statement, which executes successfully:

```
CREATE TABLE world.city (
  ID int NOT NULL AUTO_INCREMENT,
  Name char(35) NOT NULL DEFAULT '',
  CountryCode char(3) NOT NULL DEFAULT '',
  District char(20) NOT NULL DEFAULT '',
  Population int NOT NULL DEFAULT '0',
  PRIMARY KEY (ID),
  KEY CountryCode (CountryCode)
) ENGINE=InnoDB;
```

You want to improve the performance of this query:

```
SELECT Name
FROM world.city
WHERE Population BETWEEN 1000000 AND 2000000;
```

Which change enables the query to succeed while accessing fewer rows?

- A. ALTER TABLE world.city ADD SPATIAL INDEX (Name);
- B. ALTER TABLE world.city ADD SPATIAL INDEX (Population);
- C. ALTER TABLE world.city ADD INDEX (Population);
- D. ALTER TABLE world.city ADD INDEX (Name);
- E. ALTER TABLE world.city ADD FULLTEXT INDEX (Name);
- F. ALTER TABLE world.city ADD FULLTEXT INDEX (Population);

Suggested Answer: B

Community vote distribution

C (100%)

🗨️ **LrnsTgh** 8 months, 1 week ago

Selected Answer: C

agree with C.

running command line option B generated error 1687: A SPATIAL index may only contain a geometrical type column.

upvoted 1 times

🗨️ **nkanike1** 11 months, 2 weeks ago

Answer should be C

upvoted 1 times

🗨️ **Mohindar** 1 year, 2 months ago

Selected Answer: C

Answer must be C, int is not spatial column

upvoted 1 times

🗨️ **HSong** 1 year, 3 months ago

SPATIAL INDEX creates an R-tree index. For storage engines that support nonspatial indexing of spatial columns, the engine creates a B-tree index. A B-tree index on spatial values is useful for exact-value lookups, but not for range scans.

The optimizer can use spatial indexes defined on columns that are SRID-restricted.

So answer is B.

upvoted 1 times

🗨️ **fthusa** 1 year, 5 months ago

Answer C

upvoted 1 times

🗨️ **RicoAndriol** 1 year, 6 months ago

Selected Answer: C

population is INT

upvoted 1 times

🗨️ **Nelate** 1 year, 6 months ago

Correct answer should be C . You can't add a spartial index on non-geo data types

upvoted 1 times

Which three are characteristics of a newly created role? (Choose three.)

- A. It can be dropped using the DROP ROLE statement.
- B. It is stored in the mysql.role table.
- C. It is created as a locked account.
- D. It can be renamed using the RENAME ROLE statement.
- E. It can be granted to user accounts.
- F. It can be protected with a password.

Suggested Answer: AEF

Community vote distribution

ACE (100%)

 **eena96** 7 months, 3 weeks ago

Agree, A,C,E

upvoted 1 times

 **FelipeK** 10 months, 3 weeks ago

Selected Answer: ACE

The three correct statements are A, C, and E. A newly created role in MySQL 8.0 can be dropped using the DROP ROLE statement, is created as a locked account, and can be granted to user accounts. It is not stored in a mysql.role table, but rather in the mysql.user table. It cannot be protected with a password, as it has no password when created.

upvoted 2 times

 **FelipeK** 1 year, 2 months ago

The three correct statements are A, C, and E. A newly created role in MySQL 8.0 can be dropped using the DROP ROLE statement, is created as a locked account, and can be granted to user accounts. It is not stored in a mysql.role table, but rather in the mysql.user table. It cannot be protected with a password, as it has no password when created.

upvoted 3 times

 **raju8010** 1 year, 2 months ago

hi please can anyone send me the dumps

upvoted 1 times

 **fthusa** 1 year, 5 months ago

ACE

It is created as a lock account

upvoted 1 times

 **Nelate** 1 year, 6 months ago

Correct answer is A,C,E <https://dev.mysql.com/doc/refman/8.0/en/roles.html>

upvoted 1 times

You have configured GTID-based asynchronous replication with one master and one slave.

A user accidentally updated some data on the slave.

To fix this, you stopped replication and successfully reverted the accidental changes.

Examine the current GTID information:

```

Master uuid:          aaaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa
Master gtid_executed: aaaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-10300
Master gtid_purged:   aaaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-3820

Slave uuid:          bbbbbbbbb-bbbb-bbbb-bbbb-bbbbbbbbbbbbb
Slave gtid_executed: aaaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-10167,
                    bbbbbbbbb-bbbb-bbbb-bbbb-bbbbbbbbbbbbb:1-9
Slave gtid_purged:   aaaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-2312

```

You must fix GTID sets on the slave to avoid replicating unwanted transactions in case of failover.

Which set of actions would allow the slave to continue replicating without erroneous transactions?

- A. RESET MASTER;
SET GLOBAL gtid_purged=aaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-10167;
- B. SET GLOBAL gtid_purged=aaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-2312,bbbbbbbb-bbbb-bbbb-bbbb-bbbbbbbbbbbbb:1-9;
SET GLOBAL gtid_executed=aaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-10167;
- C. RESET SLAVE;
SET GLOBAL gtid_purged=aaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-3820;
SET GLOBAL gtid_executed=aaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-10300;
- D. RESET MASTER;
SET GLOBAL gtid_purged=aaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-2312;
SET GLOBAL gtid_executed=aaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-10167;
- E. RESET SLAVE;
SET GLOBAL gtid_purged=aaaaaaaa-aaaa-aaaa-aaaa-aaaaaaaaaaaaa:1-10167;

Suggested Answer: *D*

  **marklv** 11 months, 3 weeks ago

d is right

upvoted 1 times

The data in this instance is transient; no backup or replication will be required. It is currently under performing.

The database size is static and including indexes is 19G.

Total system memory is 32G.

After profiling the system, you highlight these MySQL status and global variables:

```
Com_rollback      | 85408355 |
Com_commit        | 1242342  |
InnoDB_buffer_pool_pages_free | 163840   |
```

```
[mysqld]
buffer_pool_size=20G
innodb_flush_log_at_trx_commit=2
disable-log-bin
```

The OS metrics indicate that disk is a bottleneck.

Other variables retain their default values.

Which three changes will provide the most benefit to the instance? (Choose three.)

- A. innodb_flush_log_at_trx_commit=1
- B. buffer_pool_size=24G
- C. innodb_log_file_size=1G
- D. sync_binlog=0
- E. innodb_doublewrite=0
- F. max_connections=10000
- G. innodb_undo_directory=/dev/shm

Suggested Answer: ACF

Community vote distribution

DEG (50%)

BEG (50%)

🗨️ **FelipeK** 10 months, 3 weeks ago

Selected Answer: DEG

D. sync_binlog=0: This disables the synchronizing of the binary log to disk after every commit. This can reduce disk I/O, but it comes with a risk of losing transactions in case of a crash. Since your data is transient, this might be an acceptable risk.

E. innodb_doublewrite=0: Disabling innodb_doublewrite can reduce disk I/O, but it comes with a risk of data corruption in the event of a crash. Since your data is transient and doesn't require backup or replication, this might be an acceptable risk.

G. innodb_undo_directory=/dev/shm: Moving the innodb_undo_directory to a memory-based filesystem like /dev/shm can reduce disk I/O and improve performance. However, this comes with a risk of running out of space for undo logs if the system has a high rate of updates/deletes/inserts.

upvoted 1 times

🗨️ **alemdba** 10 months, 1 week ago

why D? it log_bin has algrely disabled

upvoted 1 times

🗨️ **nkanike1** 11 months, 2 weeks ago

Agree DEG is correct

upvoted 1 times

🗨️ **alemdba** 10 months, 1 week ago

why D? it log_bin has algrely disabled

upvoted 1 times

🗨️ **marklv** 11 months, 3 weeks ago

DEG

buffer pool is already larger than data

trx makes things worse

log file doesn't necessarily make things better

upvoted 1 times

🗨️ 👤 **marklv** 11 months, 3 weeks ago
connections also doesn't do anything
upvoted 1 times

🗨️ 👤 **LrnsTgh** 8 months ago
this question is the same with number 125. should be CEG.
upvoted 2 times

🗨️ 👤 **jackymak** 7 months, 2 weeks ago
Agree. bin-log is disabled, D is wrong.
upvoted 1 times

🗨️ 👤 **alemdba** 10 months, 1 week ago
why D? it log_bin has already disabled
upvoted 1 times

🗨️ 👤 **FelipeK** 1 year, 1 month ago

Selected Answer: BEG

B. `buffer_pool_size=24G`: Increasing the buffer pool size to 24G can help reduce disk I/O by keeping more data in memory. This can improve performance, especially if disk is a bottleneck.

E. `innodb_doublewrite=0`: Disabling the doublewrite buffer can reduce disk I/O and improve performance. However, this comes at the cost of increased risk of data corruption in the event of a crash.

G. `innodb_undo_directory=/dev/shm`: Setting the undo directory to `/dev/shm` can improve performance by storing undo logs in memory instead of on disk.

upvoted 1 times

Which statement is true about InnoDB persistent index statistics?

- A. Updating index statistics is an I/O expensive operation.
- B. Index statistics are calculated from pages buffered in the buffer pool for tables with InnoDB storage engine.
- C. Setting `innodb_stats_auto_recalc=ON` causes statistics to be updated automatically when a new index is created.
- D. Execution plans based on transient index statistics improve precision when `innodb_stats_persistent_sample_pages` is increased.
- E. Increasing `innodb_stats_persistent_sample_pages` determines higher pages scanning speed, at the cost of increased memory usage.
- F. Tables are scanned and index statistics recalculated when an instance is restarted.

Suggested Answer: D

Community vote distribution

A (67%)

D (33%)

🗳️ 👤 **7206e44** 2 months ago

Selected Answer: A

- A. Correct
- D. it has to be `innodb_stats_transient_sample_pages` for this to be correct
upvoted 1 times

🗳️ 👤 **jackymak** 7 months, 2 weeks ago

Selected Answer: D

D
https://dev.mysql.com/doc/refman/8.0/en/innodb-parameters.html#sysvar_innodb_stats_persistent_sample_pages
 The number of index pages to sample when estimating cardinality and other statistics for an indexed column, such as those calculated by `ANALYZE TABLE`. [Increasing the value improves the accuracy of index statistics, which can improve the query execution plan], at the expense of increased I/O during the execution of `ANALYZE TABLE` for an InnoDB table.

Not C

https://dev.mysql.com/doc/refman/8.0/en/innodb-parameters.html#sysvar_innodb_stats_auto_recalc
 This setting applies to tables created when the `innodb_stats_persistent` option is enabled.
 upvoted 1 times

🗳️ 👤 **Zehir** 10 months, 2 weeks ago

Selected Answer: A

A.
 C is not correct because it mentions: "The `innodb_stats_auto_recalc` variable, which is enabled by default, controls whether statistics are calculated automatically when a table undergoes changes to more than 10% of its rows, not when new indexes are created."
 upvoted 1 times

🗳️ 👤 **boyhantan** 11 months, 1 week ago

C is Correct

<https://dev.mysql.com/doc/refman/8.0/en/innodb-performance-optimizer-statistics.html>
 upvoted 1 times

🗳️ 👤 **boyhantan** 11 months, 1 week ago

C is correct
 upvoted 1 times

🗳️ 👤 **nkanike1** 11 months, 2 weeks ago

D is correct according to me.
 upvoted 1 times

🗳️ 👤 **marklv** 11 months, 3 weeks ago

D is not about persistent indexed. E might be the answer but not sure

upvoted 1 times

Which two are features of MySQL Enterprise Firewall? (Choose two.)

- A. recording incoming SQL statement to facilitate the creation of a whitelist of permitted commands
- B. blocking of potential threats by configuring pre-approved whitelists
- C. modifying SQL statement dynamically with substitutions
- D. automatic locking of user accounts who break your firewall
- E. provides stateless firewall access to TCP/3306

Suggested Answer: BC

Community vote distribution

AB (100%)

  **devml** Highly Voted 1 year, 7 months ago

Selected Answer: AB

The correct answer should be A and B
upvoted 6 times

  **nkanike1** Most Recent 11 months, 2 weeks ago

The correct answer should be A and B
upvoted 2 times

  **matias886** 1 year, 3 months ago

the correct answer should be A and B
upvoted 3 times

Examine the modified output:

```
mysql> SHOW SLAVE STATUS\G
***** 1. row *****
      Slave_IO_Running: Yes
      Slave_SQL_Running: Yes
      Seconds_Behind_Master: 1612
```

Seconds_Behind_Master value is steadily growing.

What are two possible causes? (Choose two.)

- A. The master is most probably too busy to transmit data and the slave needs to wait for more data.
- B. One or more large tables do not have primary keys.
- C. This value shows only I/O latency and is not indicative of the size of the transaction queue.
- D. The master is producing a large volume of events in parallel but the slave is processing them serially.
- E. The parallel slave threads are experiencing lock contention.

Suggested Answer: CD

Community vote distribution

AD (100%)

🗨️ **MOHAMEDMOIDEEN** 6 months, 3 weeks ago

Read_Master_Log_Pos: Master Slave is connected.

upvoted 1 times

🗨️ **marklv** 9 months, 2 weeks ago

Not A. The thread is normally very fast and a master can have multiple slaves.

Tables not having primary keys can cause replication slowness.

I have seen a master execute many queries in parallel but a slave executes them serially.

upvoted 1 times

🗨️ **marklv** 9 months, 2 weeks ago

B & D, I think C may be a possible answer, but I have to think about it.

upvoted 1 times

🗨️ **nkanike1** 11 months, 2 weeks ago

Yes I agree, A & D

upvoted 2 times

🗨️ **FelipeK** 1 year, 1 month ago

Selected Answer: AD

The correct answers are:

- A. The master is most probably too busy to transmit data and the slave needs to wait for more data.
- D. The master is producing a large volume of events in parallel but the slave is processing them serially.

Explanation:

A. When the "Seconds_Behind_Master" value is steadily growing, it indicates that the slave is falling behind the master because it's waiting for data from the master. This can happen if the master is busy and cannot transmit data to the slave as quickly as the slave can process it.

D. If the master is producing a large volume of events (changes) in parallel, but the slave is processing them serially (one by one), it can lead to a situation where the slave falls behind the master. The slave might not be able to keep up with the incoming changes, causing the "Seconds_Behind_Master" value to increase.

upvoted 2 times

You must configure the MySQL command-line client to provide the highest level of trust and security when connecting to a remote MySQL Server.

Which value of `--ssl-mode` will do this?

- A. PREFERRED
- B. VERIFY_CA
- C. REQUIRED
- D. VERIFY_IDENTITY

Suggested Answer: C

Community vote distribution

D (100%)

🗨️ **FelipeK** 10 months, 3 weeks ago

Selected Answer: D

Agreed - Answer D.

upvoted 2 times

🗨️ **nkanike1** 11 months, 2 weeks ago

Agreed - Answer D.

upvoted 2 times

🗨️ **Dhanushka** 1 year, 1 month ago

Selected Answer: D

VERIFY_IDENTITY: Like VERIFY_CA, but additionally perform host name identity verification by checking the host name the client uses for connecting to the server against the identity in the certificate that the server sends to the client

upvoted 3 times

Consider this shell output and executed commands:

```
[root@oel7 ~]# ps aux | grep mysqld
mysqld 2076 3.5 24.6 1386852 372572 ? Ssl 12:01 0:01 /usr/sbin/mysqld
[root@oel7 ~]# kill -15 2076
```

Which statement is true about MySQL server shutdown?

- A. kill -15 should be avoided. Use other methods such as mysqladmin shutdown or systemctl stop mysqld.
- B. kill -15 and kill -9 are effectively the same forced shutdown that risk committed transactions not written to disk.
- C. kill -15 carries out a normal shutdown process, such as mysqladmin shutdown.
- D. mysqld_safe prohibits commands that would harm the operation of the server. An error would be returned by the kill command.

Suggested Answer: C

Community vote distribution

C (100%)

 **marklv** 11 months, 3 weeks ago

c is corect

upvoted 1 times

 **FelipeK** 1 year, 2 months ago

Selected Answer: C

The correct answer is C kill -15 carries out a normal shutdown process, such as mysqladmin shutdown.

It sends a SIGTERM signal to the MySQL server to perform a graceful shutdown.

This means that the server will complete any ongoing transactions, flush caches, and close connections before shutting down.

upvoted 1 times

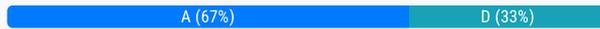
You wish to protect your MySQL database against SQL injection attacks.

Which method would fail to do this?

- A. installing and configuring the Connection Control plugin
- B. avoiding concatenation of SQL statements and user-supplied values in an application
- C. using stored procedures for any database access
- D. using PREPARED STATEMENTS

Suggested Answer: C

Community vote distribution



7206e44 2 months ago

Selected Answer: A

A. Can be used for DDoS or limit number of connections, but not necessarily sql injection
upvoted 1 times

rjaf 3 months, 3 weeks ago

Selected Answer: A

Answer is A. Connection control plugin cannot protect against sql injection.
upvoted 1 times

GalambG 8 months, 2 weeks ago

Selected Answer: D

The parameterized queries or prepared statements, which help prevent SQL injection vulnerabilities. Answer: D
upvoted 1 times

eena96 7 months, 3 weeks ago

But, the question asks which method would fail to to protect SQL injection attack
upvoted 1 times

FelipeK 10 months, 3 weeks ago

Selected Answer: A

A. installing and configuring the Connection Control plugin. This plugin is used to limit the number of connection attempts from a host, but it does not directly prevent SQL injection attacks.
upvoted 2 times

nkanike1 11 months, 2 weeks ago

Answer is A. Connection control plugin cannot protect against sql injection.
upvoted 1 times

marklv 11 months, 3 weeks ago

which one does not work, A. Connection Control plugin has nothing to do with sql injection. The rest work.
upvoted 1 times

Dhanushka 1 year, 1 month ago

Selected Answer: D

A prepared statement is a parameterized and reusable SQL query which forces the developer to write the SQL command and the user-provided data separately. The SQL command is executed safely, preventing SQL Injection vulnerabilities
upvoted 1 times

You have just installed MySQL on Oracle Linux and adjusted your `/etc/my.cnf` parameters to suit your installation.

Examine the output:

```
# systemctl start mysqld
Job for mysqld.service failed because the control process exited with error code. See "systemctl status mysqld.service" and
"journalctl -xe" for details.

# systemctl status mysqld.service
mysqld.service - MySQL Server
Loaded: loaded (/usr/lib/systemd/system/mysqld.service; enabled; vendor preset: disabled)
Active: failed (Result: exit-code) since Thu 2019-12-12 07:54:53 ACDT; 33s ago
Docs: man:mysqld(8)
      http://dev.mysql.com/doc/refman/en/using-systemd.html
Process: 2732 ExecStart=/usr/sbin/mysqld $MYSQLD_OPTS (code=exited, status=1/FAILURE)
Process: 2705 ExecStartPre=/usr/bin/mysqld_pre_systemd (code=exited, status=0/SUCCESS)
Main PID: 2732 (code=exited, status=1/FAILURE)
Status: "Server startup in progress"

Dec 12 07:54:49 oel7 systemd[1]: Starting MySQL Server...
Dec 12 07:54:53 oel7 systemd[1]: mysqld.service: main process exited, code=exited, status=1/FAILURE
Dec 12 07:54:53 oel7 systemd[1]: Failed to start MySQL Server.
Dec 12 07:54:53 oel7 systemd[1]: Unit mysqld.service entered failed state.
Dec 12 07:54:53 oel7 systemd[1]: mysqld.service failed.
```

What statement is true about the start attempt?

- A. MySQL server was not started due to a problem while executing process 2732.
- B. MySQL server continued to start up even though another process existed.
- C. systemd found the mysqld service disabled and failed to start it.
- D. systemd waited for 30 seconds before timing out and start up failed.
- E. systemd attempted to start mysqld, found another systemd mysqld process running, and shut it down.

Suggested Answer: E

 **nkanike1** 11 months, 2 weeks ago

something happened at 2732, so A
upvoted 1 times

 **marklv** 11 months, 3 weeks ago

something happened at 2732, so A.
upvoted 1 times

 **Steve26NL** 1 year, 1 month ago

Correct answer should be D

Service has been correctly enabled. No indication that systemd found another systemd mysql proces and shut it down.

upvoted 1 times

Examine these entries from the general query log:

Time	Id	Command	Argument
2019-12-17T00:36:23.389450Z	24	Connect	root@localhost on mydb using SSL/TLS
2019-12-17T00:36:23.389754Z	24	Query	select @@version_comment limit 1
2019-12-17T00:36:23.929519Z	25	Connect	root@localhost on mydb using SSL/TLS
2019-12-17T00:36:23.929846Z	25	Query	select @@version_comment limit 1
2019-12-17T00:36:27.633082Z	24	Query	START TRANSACTION
2019-12-17T00:36:30.321657Z	24	Query	UPDATE t1 SET val = 1 WHERE ID = 130
2019-12-17T00:36:32.417433Z	25	Query	START TRANSACTION
2019-12-17T00:36:33.617642Z	25	Query	UPDATE t2 SET val = 5 WHERE ID = 3805
2019-12-17T00:36:36.049458Z	25	Query	UPDATE t1 SET val = 10 WHERE ID = 130
2019-12-17T00:36:38.513674Z	24	Query	UPDATE t2 SET val = 42 WHERE ID = 3805

All UPDATE statements reference existing rows.

Which describes the outcome of the sequence of statements?

- A. Connection 24 experiences a lock wait timeout.
- B. Connection 25 experiences a lock wait timeout.
- C. A deadlock occurs immediately.
- D. All statements execute without error.
- E. A deadlock occurs after innodb_lock_wait_timeout seconds.

Suggested Answer: E

Community vote distribution

C (50%)

D (50%)

🗳️ 👤 **FelipeK** 10 months, 3 weeks ago

Selected Answer: C

```
root@localhost[world]> update t2 set val=42 where id =3805;
ERROR 1213 (40001): Deadlock found when trying to get lock; try restarting transaction
```

```
+-----+-----+
| Variable_name | Value |
+-----+-----+
| innodb_lock_wait_timeout | 900 |
+-----+-----+
upvoted 1 times
```

🗳️ 👤 **marklv** 11 months, 3 weeks ago

D is right
upvoted 1 times

🗳️ 👤 **marklv** 9 months, 1 week ago

correction, 24 and 25 hit the same id, an started a transaction. B and C are correct except b is NOT correct if 24 commits before timeout. So C happens. And C happens because the updates happen BEFORE 24 commits.
upvoted 1 times

🗳️ 👤 **marklv** 2 months ago

I see I didn't say the answer. No commit is ever executed, so the transactions are in limbo. So E is the correct answer. Eventually a timeout will occur. Auto commit does not matter when you begin a transaction. Both connections hit the same row, and are not committed, so there are lock.
upvoted 1 times

🗳️ 👤 **Dhanushka** 1 year, 2 months ago

Selected Answer: D

Auto-commit not been disabled hence it should be able to execute all
upvoted 1 times

Examine this command, which executes successfully:

```
$ mysqlrouter --bootstrap user@hostname:port --directory=directory_path
```

Which activity is performed?

- A. MySQL Router is configured based on the information in files in `directory_path`.
- B. MySQL Router configures itself based on the information retrieved from the InnoDB cluster metadata server.
- C. MySQL Router is restarted.
- D. MySQL Router configures all the cluster nodes based on the information retrieved from the InnoDB cluster metadata server.

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ **marklv** 11 months, 3 weeks ago

B is it

upvoted 1 times

🗨️ **FelipeK** 1 year, 1 month ago

Selected Answer: B

The given command `$ mysqlrouter --bootstrap user@hostname:port --directory=directory_path` is using the MySQL Router's `--bootstrap` option. The purpose of the `--bootstrap` option is to configure MySQL Router to connect to an InnoDB cluster metadata server and retrieve configuration information from it.

upvoted 2 times

You encountered an insufficient privilege error in the middle of a long transaction.
The database administrator is informed and immediately grants the required privilege:
`GRANT UPDATE ON world.city TO 'user1';`
How can you proceed with your transaction with the least interruption?

- A. Roll back the transaction and start the transaction again in the same session.
- B. Re-execute the failed statement in your transaction.
- C. Change the default database and re-execute the failed statement in your transaction.
- D. Close the connection, reconnect, and start the transaction again.

Suggested Answer: B

Community vote distribution

B (100%)

 **FelipeK** 1 year, 1 month ago

Selected Answer: B

In this situation, the database administrator has granted the required privilege (UPDATE on world.city) to 'user1', which was missing initially. To proceed with the transaction with the least interruption, you can re-execute the failed statement in your transaction. Since the privilege issue has been resolved by the administrator's action, re-executing the statement should work without any further issues. This approach minimizes interruption and allows you to continue your transaction without having to roll back the entire transaction or close and reconnect the connection.

upvoted 2 times

Examine these statements, which execute successfully:

```
CREATE ROLE r_world_rd;
GRANT SELECT ON world.* TO r_world_rd;
CREATE USER john IDENTIFIED BY 'P@ssw0rd';
GRANT r_world_rd TO john;
```

Examine these statements issued by user John:

```
mysql> SHOW GRANTS;
+-----+
| Grants for john@% |
+-----+
| GRANT USAGE ON *.* TO 'john'@'%' |
| GRANT 'r_world_rd'@'%' TO 'john'@'%' |
+-----+
2 rows in set (0.01 sec)

mysql> SELECT * FROM world.city;
ERROR 1142 (42000): SELECT command denied to user 'john'@'localhost' for table 'city'
```

What is the reason for the error?

- A. The statement was blocked by MySQL Firewall.
- B. John has not activated the role.
- C. John needs to reconnect to the database.
- D. The DBA needs to execute FLUSH PRIVILEGES.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **FelipeK** 11 months ago

Selected Answer: B

The correct is B

upvoted 2 times

🗳️ 👤 **nkanike1** 11 months, 2 weeks ago

Answer B

upvoted 2 times

🗳️ 👤 **dieymir** 1 year, 3 months ago

answer is correct

upvoted 1 times

🗳️ 👤 **fthusa** 1 year, 5 months ago

Answer is not there city is no in the world database that's why this error getting

upvoted 1 times

Binary log events for the 'mydb1' schema must be copied to a different schema name 'mydb2'. Which command will do this?

- A. `mysqlbinlog --rewrite-db='mydb1->mydb2' | mysql`
- B. `mysqlbinlog --datebase=mydb1 --database=mydb2 | mysql`
- C. `mysqlbinlog --rewrite-db='mydb1' --rewrite-db='mydb2' | mysql`
- D. `mysqlbinlog --read-from-remote-server --raw | sed 's/mydb1/mydb2/g' | mysql`

Suggested Answer: D

Community vote distribution

A (100%)

🗨️ **nkanike1** 11 months, 2 weeks ago

Answer A

upvoted 1 times

🗨️ **FelipeK** 1 year, 1 month ago

Selected Answer: A

The correct command to copy binary log events for the 'mydb1' schema to a different schema name 'mydb2' is `mysqlbinlog --rewrite-db='mydb1->mydb2' | mysql1`. So, the correct answer is A. This command updates a database with a different name than the original

upvoted 1 times

🗨️ **Dhanushka** 1 year, 2 months ago

Selected Answer: A

It does not tell read from remote server

upvoted 2 times

Examine this MySQL Shell command:

```
dba.rebootClusterFromCompleteOutage()
```

Which two statements are true? (Choose two.)

- A. It reconfigures InnoDB Cluster if the cluster was stopped.
- B. It performs InnoDB Cluster instances rolling restart.
- C. It only starts all InnoDB Cluster instances.
- D. It is not mandatory that all instances are running and reachable before running the command.
- E. It stops and restarts all InnoDB Cluster instances and initializes the metadata.
- F. It only stops and restarts all InnoDB Cluster instances.
- G. It picks the minimum number of instances necessary to rebuild the quorum and reconfigures InnoDB Cluster.

Suggested Answer: *BD*

Community vote distribution

AG (100%)

🗨️ 👤 **FelipeK** 8 months ago

Selected Answer: AG

A. It reconfigures InnoDB Cluster if the cluster was stopped. This function is specifically designed to recover a cluster after a complete outage, which means that group replication has stopped on all member instances.

G. It picks the minimum number of instances necessary to rebuild the quorum and reconfigures InnoDB Cluster. The function retrieves the cluster metadata and topology from the current instance, checks which instances of the cluster are currently reachable, and recovers the cluster based on the metadata stored in that instance. If any member is currently unreachable, the function fails, but this check can be bypassed with the force option.

upvoted 1 times

🗨️ 👤 **marklv** 2 months ago

I don't know if it reconfigures anything.

upvoted 1 times

🗨️ 👤 **marklv** 11 months, 3 weeks ago

C and D. It doesn't restart. Only a majority are needed to start the cluster. It just starts nodes.

upvoted 1 times

Which two statements are true about MySQL server multi-source replication? (Choose two.)

- A. It is not compatible with auto-positioning.
- B. It needs to be re-instanced after a crash to maintain consistency.
- C. It uses only time-based replication conflict resolution.
- D. It relies on relay_log_recovery for resilient operations.
- E. It does not attempt to detect or resolve replication conflicts.
- F. It must use GTID replication.

Suggested Answer: CF

Community vote distribution

DE (75%)

EF (25%)

🗨️ 👤 **jackymak** 7 months, 1 week ago

Selected Answer: DE

For D, <https://dev.mysql.com/doc/refman/8.3/en/replication-options-reference.html>

relay_log_recovery: Whether automatic recovery of relay log files from source at startup is enabled; must be enabled for crash-safe replica.

Not F, <https://dev.mysql.com/doc/refman/8.3/en/replication-multi-source-configuration.html>

Sources in a multi-source replication topology can be configured to use either GTID-based replication, or binary log position-based replication.

upvoted 1 times

🗨️ 👤 **FelipeK** 8 months ago

Selected Answer: EF

E. Multi-source replication in MySQL does not attempt to detect or resolve replication conflicts. Conflicts need to be handled by the application or by using other mechanisms outside the scope of multi-source replication.

F. Multi-source replication in MySQL 8.0 requires the use of Global Transaction Identifiers (GTID) for replication. GTID helps in uniquely identifying transactions across different servers, which is essential for multi-source replication to work effectively.

upvoted 1 times

🗨️ 👤 **FelipeK** 10 months, 3 weeks ago

Selected Answer: DE

The correct answers are:

D. It relies on relay_log_recovery for resilient operations. E. It does not attempt to detect or resolve replication conflicts.

In a multi-source replication setup, MySQL relies on the relay_log_recovery option to enable or disable automatic relay log recovery just after server startup, which helps in resilient operations. Also, MySQL does not have a built-in mechanism to detect or resolve conflicts if the same data is modified by transactions from different sources. It's the responsibility of the database administrators to ensure that the data modified by each source is distinct to avoid conflicts.

upvoted 2 times

🗨️ 👤 **marklv** 11 months, 3 weeks ago

DE are it

upvoted 2 times

Which command enables rule-based MySQL Auditing capabilities?

- A. shell> mysql < audit_log_filter_linux_install.sql
- B. shell> mysqld --initialize --log-raw=audit.log
- C. mysql> INSTALL PLUGIN audit_log;
- D. mysql> INSTALL COMPONENT audit_log;

Suggested Answer: A

  **marklv** 11 months, 3 weeks ago

A is it

upvoted 1 times

Examine this SQL statement:

```
mysql> GRANT r_read@localhost TO mark WITH ADMIN OPTION;
```

Which two are true? (Choose two.)

- A. Mark can grant the privileges assigned to the r_read@localhost role to another user.
- B. ADMIN OPTION causes the role to be activated by default.
- C. Mark can grant the r_read@localhost role to another user.
- D. Mark can revoke the r_read@localhost role from another role.
- E. ADMIN OPTION allows Mark to drop the role.
- F. Mark must connect from localhost to activate the r_read@localhost role.

Suggested Answer: AE

Community vote distribution

CD (100%)

🗨️ **nkanike1** 11 months, 2 weeks ago

C and D

upvoted 1 times

🗨️ **xexeu** 1 year, 4 months ago

C and D

upvoted 1 times

🗨️ **fthusa** 1 year, 5 months ago

What about C , answer C also correct

upvoted 1 times

🗨️ **FelipeK** 1 year, 5 months ago

Selected Answer: CD

<https://dev.mysql.com/doc/refman/8.0/en/grant.html>

upvoted 2 times

Which four are types of information stored in the MySQL data dictionary? (Choose four.)

- A. performance metrics
- B. table definitions
- C. access control lists
- D. view definitions
- E. server runtime configuration
- F. server configuration rollback
- G. stored procedure definitions
- H. InnoDB buffer pool LRU management data

Suggested Answer: *BDFG*

Community vote distribution

BCDG (100%)

 **FelipeK** 11 months ago

Selected Answer: BCDG

B-Table Definitions: The INFORMATION_SCHEMA.TABLES provides details about tables in the database.

D-View Definitions: Information about views is also stored in the data dictionary.

G-Stored Procedure Definitions: Information about stored procedures is stored in the data dictionary.

C-Access Control Lists: While not explicitly mentioned, access control lists would be part of the metadata about database objects stored in the data dictionary.

upvoted 2 times

 **nkanike1** 11 months, 2 weeks ago

I feel BCDG

upvoted 1 times

 **marklv** 11 months, 3 weeks ago

at least BDG

upvoted 1 times

You have an InnoDB Cluster configured with three servers.

Examine this command, which executes successfully:

```
mysqldump -uroot -p -d mydatabase > mydatabase_backup.sql
```

Due to data loss, the cluster is initialized and a restore is attempted resulting in this error:

```
ERROR 13176 (HY000) at line 23: Cannot update GTID_PURGED with the Group Replication plugin running
```

Which two actions, either one of which, can fix this error and allow a successful restore of the cluster? (Choose two.)

- A. Stop all instances except the primary read/write master instance and run the restore.
- B. Remove the @@GLOBAL.gtid_purged statement from the dump file.
- C. Create the backup by using the --set-gtid-purged=OFF option.
- D. Remove the group replication plugin from each instance before restoring.
- E. Remove the @@GLOBAL.gtid_executed statement from the dump file.
- F. Restore using the --set-gtid-purged=OFF option.

Suggested Answer: *BC*

  **marklv** 11 months, 3 weeks ago

I think B and C

upvoted 1 times

Which statement is true about MySQL Enterprise Transparent Data Encryption (TDE)?

- A. MySQL TDE uses an appropriate keyring plugin to store the keys in a centralized location.
- B. Both MyISAM and InnoDB tables can be encrypted by setting the keyring_engine = ALL variable in the MySQL configuration file.
- C. Lost tablespace encryption keys can be regenerated only if the master database key is known or present in the Key Vault specification.
- D. TDE can encrypt InnoDB and MyISAM tables only when the tables are stored in the SYSTEM tablespace.

Suggested Answer: C

Community vote distribution

A (100%)

🗨️ **FelipeK** 10 months, 3 weeks ago

Selected Answer: A

InnoDB uses a two-tier encryption key architecture, consisting of a master encryption key and tablespace keys. When a tablespace is encrypted, a tablespace key is encrypted and stored in the tablespace header. When an application or authenticated user wants to access encrypted tablespace data, InnoDB uses a master encryption key to decrypt the tablespace key. The decrypted version of a tablespace key never changes, but the master encryption key can be changed as required. This action is referred to as master key rotation.

upvoted 2 times

🗨️ **Zehir** 10 months, 3 weeks ago

A.

[https://www.mysql.com/products/enterprise/tde.html#:~:text=MySQL%20Enterprise%20Transparent%20Data%20Encryption%20\(TDE\)%20protects%20your](https://www.mysql.com/products/enterprise/tde.html#:~:text=MySQL%20Enterprise%20Transparent%20Data%20Encryption%20(TDE)%20protects%20your)

upvoted 1 times

🗨️ **marklv** 11 months, 3 weeks ago

can't encrypt MyISAM, so BD are out.,

A and C look right.

upvoted 2 times

You are upgrading a MySQL instance to the latest 8.0 version.

Examine this output:

```
mysql> SHOW GLOBAL VARIABLES LIKE '%dir';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| basedir       | /usr  |
| ..           | ..    |
| datadir       | /var/lib/mysql |
| innodb_data_home_dir | /innodb_data |
| innodb_log_group_home_dir | ./ |
| innodb_temp_tablespaces_dir | ./#innodb_temp/ |
| innodb_tmpdir |      |
| ..           | ..    |
| plugin_dir    | /usr/lib/plugin |
| ..           | ..    |
| tmpdir        | /tmp:/var/tmp |
+-----+-----+
```

You plan to add this parameter to the configuration:

```
innodb_directories='/innodb_extras'
```

Which statement is true?

- A. It defines all innodb tablespace options relative to a starting parent directory.
- B. It is not necessary because innodb_data_home_dir is already defined.
- C. It allows scanning of other locations to discover more innodb tablespaces.
- D. It moves all innodb tablespaces to the /innodb_extras directory to enable a new innodb_data_home_dir to be defined.
- E. It adds more temporary workspace in addition to the innodb_tmpdir location.

Suggested Answer: B

Community vote distribution

C (100%)

 **dudeDaG** 1 year, 3 months ago

Even though the correct answer is "C", so the correct/approved/recommended answer for Exam is B?

upvoted 1 times

 **xexeu** 1 year, 4 months ago

Answer C

upvoted 1 times

 **FelipeK** 1 year, 5 months ago

Selected Answer: C

<https://dev.mysql.com/doc/refman/8.0/en/innodb-moving-data-files-offline.html>

upvoted 1 times

 **FelipeK** 10 months, 3 weeks ago

The correct statement is:

- C. It allows scanning of other locations to discover more innodb tablespaces.

The innodb_directories option in MySQL 8.0 allows the server to scan additional directories to discover more InnoDB tablespaces. This is particularly useful when you have tablespaces located outside the main data directory. So, by adding '/innodb_extras' to innodb_directories, MySQL will be able to discover any tablespaces located in this directory.

upvoted 1 times

You wish to store the username and password for a client connection to MySQL server in a file on a local file system. Which is the best way to encrypt the file?

- A. Use the AES_ENCRYPT() MySQL function on the option file.
- B. Use mysql_secure_installation to encrypt stored login credentials.
- C. Use a text editor to create a new defaults file and encrypt it from Linux prompt.
- D. Use mysql_config_editor to create an encrypted file.

Suggested Answer: D

Community vote distribution

D (100%)

🗨️ 👤 **FelipeK** 11 months ago

Selected Answer: D

The best way to encrypt the file is to use mysql_config_editor to create an encrypted file. This utility encrypts the .mylogin.cnf file which contains the login credentials. The other options do not provide the necessary encryption for a file stored on a local file system. So, the correct answer is D. Use mysql_config_editor to create an encrypted file.

upvoted 1 times

🗨️ 👤 **marklv** 11 months, 3 weeks ago

D is easiest, but it isn't perfect

upvoted 1 times

You are backing up raw InnoDB files by using mysqlbackup.

Which two groups of files will be backed up during a full backup? (Choose two.)

- A. ibbackup files
- B. *.CSM files
- C. *.sdi files
- D. *.ibd files
- E. ib_logfile* files

Suggested Answer: DE

Community vote distribution

DE (100%)

🗨️ 👤 **FelipeK** 11 months ago

Selected Answer: DE

InnoDB data files (ibdata files and .ibd files)

InnoDB redo log files (#ib_redo N files in MySQL 8.0.30 and higher or ib_logfile files in earlier releases)

upvoted 3 times

🗨️ 👤 **marklv** 11 months, 3 weeks ago

AB. CSV tables create CSM files, innodb stores data in ibd files

upvoted 1 times

🗨️ 👤 **marklv** 9 months, 1 week ago

Sorry its says its only backing up innodb.

<https://dev.mysql.com/doc/mysql-enterprise-backup/4.1/en/meb-files-backed-up-innodb.html>

D and E. ibd and ib_logfile

upvoted 3 times

You made some table definition changes to a schema in your MySQL Server.

Which two statements reflect how MySQL Server handles the table definition changes? (Choose two.)

- A. MySQL writes SDI to the binary log for distributed backups.
- B. MySQL keeps InnoDB metadata changes in .sdi files in datadir.
- C. The metadata is serialized in JSON format in Serialized Dictionary Information (SDI).
- D. MySQL Server stores a copy of the serialized data in the InnoDB user tablespace.
- E. MySQL implicitly executes FLUSH TABLES and stores a snapshot backup of the metadata.

Suggested Answer: CD

Community vote distribution

CD (100%)

🗨️ 👤 **FelipeK** 11 months ago

Selected Answer: CD

C. The metadata is serialized in JSON format in Serialized Dictionary Information (SDI): When you make changes to the table schema (e.g., adding or modifying columns), MySQL serializes the metadata changes into a JSON format known as Serialized Dictionary Information (SDI). This information is stored in-memory and in the binary log for replication purposes.

D. MySQL Server stores a copy of the serialized data in the InnoDB user tablespace: InnoDB tables store the serialized SDI in the InnoDB system tablespace (not in the user tablespace). This allows MySQL to recover the table structure even after a server restart.

<https://dev.mysql.com/doc/refman/8.0/en/serialized-dictionary-information.html>

upvoted 3 times

🗨️ 👤 **marklv** 11 months, 1 week ago

For DDL changes. a serialized copy is kept in SDI in json format, a copy is stored in tablespace for innodb tables

upvoted 1 times

Which two are characteristics of snapshot-based backups? (Choose two.)

- A. Snapshot-based backups greatly reduce time during which the database and applications are unavailable.
- B. There is no need for InnoDB tables to perform its own recovery when restoring from the snapshot backup.
- C. The frozen file system can be cloned to another virtual machine immediately into active service.
- D. A separate physical copy must be made before releasing the snapshot backup.
- E. Snapshot backups can be used only in virtual machines.

Suggested Answer: *DE*

 **FelipeK** 7 months, 2 weeks ago

I think is A and C.

Snapshot-based backups greatly reduce time during which the database and applications are unavailable. This is because snapshots provide logical copies of the file system at a given point in time, without requiring a physical copy of the entire file system.

The frozen file system can be cloned to another virtual machine immediately into active service. This is implied by the nature of snapshot backups, as they capture the state of the file system at a specific point in time.

upvoted 2 times

 **xjllll** 9 months ago

A and D

upvoted 2 times

 **jackymak** 5 months, 3 weeks ago

agree D

upvoted 2 times

 **marklv** 2 months ago

Not D, if the snapshot is writable, you can run mysqld off of it, no need to copy.

upvoted 1 times

 **marklv** 11 months, 1 week ago

B and C I think.

upvoted 1 times

You must run multiple instances of MySQL Server on a single host.

Which three methods are supported? (Choose three.)

- A. Run MySQL Server docker containers.
- B. Use systemd with different settings for each instance.
- C. Use system tools to lock each instance to its own CPU.
- D. Start mysqld or mysqld_safe using different option files for each instance.
- E. Run mysqld with --datadir defined for each instance.
- F. Use resource groups to lock different instances on separate CPUs.

Suggested Answer: BDE

Community vote distribution

ABD (50%)

ADE (25%)

BDE (25%)

🗳️ 👤 **xjlll** 9 months ago

Selected Answer: ABD

--datadir with --initialize

upvoted 2 times

🗳️ 👤 **FelipeK** 10 months, 3 weeks ago

Selected Answer: ADE

A. Run MySQL Server docker containers: Docker containers are a popular way to isolate and run multiple instances of MySQL on the same host.

D. Start mysqld or mysqld_safe using different option files for each instance: You can start multiple instances of MySQL by specifying different option files for each instance, which allows you to configure them separately.

E. Run mysqld with --datadir defined for each instance: You can specify a different data directory for each MySQL instance using the --datadir option.

upvoted 1 times

🗳️ 👤 **FelipeK** 11 months ago

Selected Answer: BDE

The correct options are BDE

upvoted 1 times

🗳️ 👤 **marklv** 11 months, 1 week ago

ABD

A ocker

B systemd and specify different config option files

D mutiple things need to be changed, so different configurations

upvoted 4 times

🗳️ 👤 **marklv** 9 months, 1 week ago

BTW, not E because you just specified the datadir, what about the port, the logfiles, etc.

upvoted 2 times

There are five MySQL instances configured with a working group replication.

Examine the output of the group members:

```
mysql> SELECT MEMBER_ID, MEMBER_STATE FROM performance_schema.replication_group_members;
+-----+-----+
| MEMBER_ID | MEMBER_STATE |
+-----+-----+
| 1999b9fb-4aaf-11e6-bb54-28b2bd168d07 | UNREACHABLE |
| 199b2df7-4aaf-11e6-bb16-28b2bd168d07 | ONLINE      |
| 199bb88e-4aaf-11e6-babe-28b2bd168d07 | ONLINE      |
| 19ab72fc-4aaf-11e6-bb51-28b2bd168d07 | UNREACHABLE |
| 19b33846-4aaf-11e6-ba81-28b2bd168d07 | UNREACHABLE |
+-----+-----+
```

Which two statements are true about network partitioning in the cluster? (Choose two.)

- A. The cluster will shut down to preserve data consistency.
- B. The cluster has built-in high availability and updates group_replication_ip_whitelist to remove the unreachable nodes.
- C. The group replication will buffer the transactions on the online nodes until the unreachable nodes return online.
- D. There could be both a 2 node and 3 node group replication still running, so shutting down group replication and diagnosing the issue is recommended.
- E. A manual intervention to force group members to be only the working two instances is required.

Suggested Answer: AB

Community vote distribution

DE (100%)

 **FelipeK** 11 months ago

Selected Answer: DE

D. There could be both a 2 node and 3 node group replication still running, so shutting down group replication and diagnosing the issue is recommended. In a situation where there is a network partition, the remaining nodes are not able to tell if the unreachable servers have crashed or whether a network partition has isolated these nodes alone and therefore the group cannot be reconfigured automatically¹. This could potentially lead to both a 2 node and 3 node group replication still running.

E. A manual intervention to force group members to be only the working two instances is required. When a majority of group members is lost, the group is unable to progress and blocks because it cannot secure majority or quorum¹. In this case, manual intervention may be required to reconfigure the group with only the working instances.

<https://dev.mysql.com/doc/refman/8.0/en/group-replication-network-partitioning.html>

upvoted 2 times

 **xjlll** 9 months ago

Why not C

upvoted 1 times

 **marklv** 11 months, 1 week ago

D, the 3 unreachable nodes could be together in a group

E, If you want to up nodes in a cluster, you need to remove the other 3

upvoted 1 times

Examine this statement and output:

```
mysql> SELECT ROW_NUMBER() OVER() AS QN,
        query, exec_count, avg_latency, lock_latency
        FROM sys.statement_analysis
        ORDER BY exec_count;
```

QN	query	exec_count	avg_latency	lock_latency
1	SELECT SUM ('k') FROM `mysch ...` () - INTERVAL ? SQL_TSI_HOUR	381268	31.44 ms	1.01 m
2	SELECT `id`, `val`, `a`, `b` ... `updated` WHERE `created` < ?	150317	358.34 us	30.06 s
3	SELECT `emp_no`, `val`, `cre ...` ated' + INTERVAL ? SQL_TSI_DAY	600	523.32 ms	120.24 ms
4	SELECT `a`, `b`, `c` FROM `m ...` ? AND ? OR `k` BETWEEN ? AND ?	200	10.32 s	40.19 ms
5	SELECT `a`, `b` FROM `myschem ...` G (`emp_no`) WHERE `val` = ?	1	21.03 s	274.00 us

You must try to reduce query execution time.

Which two queries should you focus on? (Choose two.)

- A. QN = 3
- B. QN = 5
- C. QN = 1
- D. QN = 4
- E. QN = 2

Suggested Answer: CE

 **marklv** 11 months ago

- E. high lock time -- query 2
- A. Long execution time, somewhat locking. -- query 3

query 1 looks fine though there are many executions.
upvoted 1 times

 **alemdba** 10 months, 1 week ago

i think should focus on time -- ms s m us
upvoted 1 times

 **alemdba** 10 months, 1 week ago

CE i think.
E has 358 us letency
C has 1 min letency
A has only ms lenetcy
upvoted 1 times

```
mysql> SHOW FULL PROCESSLIST;
+----+-----+-----+-----+
| Id | User          | Host          | ...  |
+----+-----+-----+-----+
| 4  | event_scheduler | localhost    | ...  |
| 9  | root          | localhost:51502 | ...  |
| 10 | root          | localhost:51670 | ...  |
```

Examine this query:

```
SELECT SUM(m.CURRENT_NUMBER_OF_BYTES_USED) AS TOTAL
FROM performance_schema.memory_summary_by_thread_by_event_name m
INNER JOIN performance_schema.threads t
ON m.THREAD_ID = t.THREAD_ID
WHERE t.PROCESSLIST_ID = 10;
```

What information does this query provide?

- A. total memory used across all connections associated with the user on connection number 10
- B. total memory used by the first 10 connections
- C. total memory used by thread number 10
- D. total memory used across all connections associated with the user on thread number 10
- E. total memory used by connection number 10
- F. total memory used by the first 10 threads

Suggested Answer: E

  **marklv** 9 months, 1 week ago

E. counts memory where processid = 10
upvoted 1 times

On examination, your MySQL installation datadir has become recursively world read/write/executable. What are two major concerns of running an installation with incorrect file privileges? (Choose two.)

- A. Users could overwrite configuration files.
- B. Data files could be deleted.
- C. SQL injections could be used to insert bad data into the database.
- D. MySQL binaries could be damaged, deleted, or altered.
- E. Extra startup time would be required for the MySQL server to reset the privileges.

Suggested Answer: *BE*

Community vote distribution



🗳️ 👤 **xjlll** 9 months ago

Selected Answer: AB

A and B

upvoted 1 times

🗳️ 👤 **marklv** 11 months ago

- A. Configuration could be changed, but my.cnf is most likely not in it, but others.
- B. Data could be altered.

D IS NOT CORRECT.

Data injection is done over a connection, binaries ARE not affected since they are not in the datadir, startup doesn't change permissions.

upvoted 2 times

🗳️ 👤 **LrnsTgh** 8 months, 1 week ago

<https://dev.mysql.com/doc/refman/8.0/en/data-directory.html>

maybe mysqld-auto.cnf. it is in the subdirectory of datadir. me, A & B

upvoted 1 times

🗳️ 👤 **FelipeK** 11 months ago

Selected Answer: AD

A. Users could overwrite configuration files: If unauthorized users have write access to MySQL configuration files, they could modify them, potentially leading to misconfigurations that may compromise the security or stability of the database server.

D. MySQL binaries could be damaged, deleted, or altered: Incorrect file privileges could allow unauthorized users to tamper with MySQL binary files, potentially damaging, deleting, or altering them. This could render the MySQL server inoperable or introduce malicious code.

upvoted 1 times

🗳️ 👤 **marklv** 9 months, 1 week ago

binaries are not in the datadir, do D is not it

upvoted 1 times

User account `baduser@hostname` on your MySQL instance has been compromised.

Which two commands stop any new connections using the compromised account? (Choose two.)

- A. `ALTER USER baduser@hostname PASSWORD DISABLED;`
- B. `ALTER USER baduser@hostname MAX_USER_CONNECTIONS 0;`
- C. `ALTER USER baduser@hostname ACCOUNT LOCK;`
- D. `ALTER USER baduser@hostname IDENTIFIED WITH mysql_no_login;`
- E. `ALTER USER baduser@hostname DEFAULT ROLE NONE;`

Suggested Answer: *CD*

Community vote distribution

CD (100%)

 **marklv** 11 months ago

A is no such thing.

B and E do no stop connecting.

C and D are it.

upvoted 1 times

 **FelipeK** 11 months ago

Selected Answer: CD

C. `ALTER USER baduser@hostname ACCOUNT LOCK;`

This command locks the user account, preventing any new connections.

D. `ALTER USER baduser@hostname IDENTIFIED WITH mysql_no_login;`

This command changes the user's authentication method to one that doesn't allow login, effectively blocking new connections.

upvoted 1 times

An existing asynchronous replication setup is running MySQL 8.
Which two steps are a part of implementing GTID replication? (Choose two.)

- A. Enable GTID by executing this on the master and the slave:
`SET GLOBAL GTID_ENABLED=on;`
- B. Execute this on the slave to enable GTID:
`START SLAVE IO_THREAD WITH GTID;`
- C. Restart MySQL (master and slave) with these options enabled:
`--gtid_mode=ON`
`--log-bin`
`--log-slave-updates`
`--enforce-gtid-consistency`
- D. Execute this on the slave to enable GTID:
`RESET SLAVE; START SLAVE GTID_NEXT=AUTOMATIC;`
- E. On the slave, alter the MySQL master connection setting with:
`ALTER channel CHANGE MASTER TO MASTER_AUTO_POSITION = 1;`
- F. On the slave, alter the MySQL master connection setting with:
`CHANGE MASTER TO MASTER_AUTO_POSITION = 1;`

Suggested Answer: CF

Community vote distribution

CF (100%)

 **marklv** 9 months, 1 week ago

C and maybe D

<https://dev.mysql.com/doc/refman/8.0/en/replication-gtids-failover.html>

F doesn't start replication. D does. Not sure.

upvoted 1 times

 **FelipeK** 11 months ago

Selected Answer: CF

C. Restart MySQL (master and slave) with these options enabled: `--gtid_mode=ON --log-bin --log-slave-updates --enforce-gtid-consistency`

F. On the slave, alter the MySQL master connection setting with: `CHANGE MASTER TO MASTER_AUTO_POSITION = 1`

upvoted 1 times

Which four connection methods can MySQL clients specify with the --protocol option when connecting to a MySQL server? (Choose four.)

- A. TCP
- B. SOCKET
- C. PIPE
- D. DIRECT
- E. IPv6
- F. FILE
- G. IPv4
- H. MEMORY

Suggested Answer: *ABCH*

  **marklv** 11 months ago

* TCP, socket

* For windows, pipe, memory

ABCH

upvoted 1 times

Examine this command and output:

```
mysql> SELECT * FROM data_locks LIMIT 1\G
***** 1. row *****
ENGINE: INNODB
ENGINE_LOCK_ID: 1200:146
ENGINE_TRANSACTION_ID: 1200
THREAD_ID: 45
EVENT_ID: 11
OBJECT_SCHEMA: mydb
OBJECT_NAME: mytable1
PARTITION_NAME: NULL
SUBPARTITION_NAME: NULL
INDEX_NAME: NULL
OBJECT_INSTANCE_BEGIN: 118793337250203
LOCK_TYPE: RECORD
LOCK_MODE: X
LOCK_STATUS: GRANTED
LOCK_DATA: 1922,1922
```

Which two statements are true? (Choose two.)

- A. The lock is at the metadata object level.
- B. The lock is a shared lock.
- C. The lock is an intentional lock.
- D. The lock is at the table object level.
- E. The lock is a row-level lock.
- F. The lock is an exclusive lock.

Suggested Answer: *DF*

Community vote distribution

EF (100%)

 **marklv** 11 months ago

EF

X is exclusive

RECORD is row lock

upvoted 1 times

 **xexeu** 1 year, 4 months ago

Answer EF

upvoted 1 times

 **FelipeK** 1 year, 5 months ago

Selected Answer: EF

<https://dev.mysql.com/doc/refman/8.0/en/innodb-locking.html>

<https://dev.mysql.com/doc/refman/8.0/en/performance-schema-data-locks-table.html>

upvoted 3 times

Which three statements are true about MySQL replication? (Choose three.)

- A. Replication can use only TCP/IP connections.
- B. Any instance can have multiple slaves, but it can have only one master.
- C. Each instance in a replication topology must have a unique server ID.
- D. Binary logs contain only transactions originating from a single MySQL instance.
- E. Each slave must have its own MySQL user for replication.
- F. Binary logging must be enabled on the master in order to replicate to other instances.
- G. A replication user must have the SELECT privilege for all tables that need to be replicated.

Suggested Answer: CFG

Community vote distribution

CFG (100%)

 **marklv** 11 months ago

ACF

G IS NOT RIGHT. The slave does not need to be able to select data. It just needs to be able to get to the binlog and download commands.

upvoted 1 times

 **marklv** 9 months, 1 week ago

<https://dev.mysql.com/doc/refman/8.0/en/change-replication-source-to.html>

"Note

Replication cannot use Unix socket files. You must be able to connect to the replication source server using TCP/IP."

For replication, you just need to be able read the binlogs. You do not need access to the tables on the master, so G is not it. A is.

upvoted 2 times

 **FelipeK** 11 months ago

Selected Answer: CFG

C. Each instance in a replication topology must have a unique server ID. This is important to ensure that each instance can be uniquely identified in the replication process.

F. Binary logging must be enabled on the master in order to replicate to other instances. Binary logging is essential for the master server to record changes that need to be replicated to the slave servers.

G. A replication user must have the SELECT privilege for all tables that need to be replicated. The replication user on the slave server needs the SELECT privilege on the tables it's replicating to retrieve the data changes from the master.

upvoted 1 times

Which two statements are true about the `mysql_config_editor` program? (Choose two.)

- A. It manages the configuration of the MySQL Firewall feature.
- B. It manages the configuration of client programs.
- C. It can move `datadir` to a new location.
- D. It manages the configuration of user privileges for accessing the server.
- E. It will use `[client]` options by default unless you provide `--login-path`.
- F. It can be used to create and edit SSL certificates and log locations.
- G. It provides an interface to change `my.cnf` files.

Suggested Answer: *BE*

Community vote distribution

BE (100%)

 **FelipeK** 11 months ago

Selected Answer: BE

B. It manages the configuration of client programs. This includes storing credentials and other client-specific configuration settings securely.

E. It will use `[client]` options by default unless you provide `--login-path`. When you use `mysql_config_editor` to store credentials, it stores them under a `[client]` login path by default.

upvoted 1 times

Examine this command, which executes successfully:

```
shell> mysqldump --master-data=2 --single-transaction --result-file=dump.sql mydb
```

Which two statements are true? (Choose two.)

- A. It executes flush tables with read lock.
- B. It enforces consistent backups for all storage engines.
- C. The backup created is a consistent data dump.
- D. This option uses the READ COMMITTED transaction isolation mode.
- E. It is a cold backup.

Suggested Answer: AD

Community vote distribution



🗉 **marklv** 11 months ago

They are all wrong.

It would be B and C if only innodb storage engine is used. MyISAM tables are not consistently backed up

https://dev.mysql.com/doc/refman/8.0/en/mysqldump.html#option_mysqldump_single-transaction

"When using this option, you should keep in mind that only InnoDB tables are dumped in a consistent state. For example, any MyISAM or MEMORY tables dumped while using this option may still change state."

upvoted 1 times

🗉 **FelipeK** 11 months ago

Selected Answer: BC

B. It enforces consistent backups for all storage engines.

This statement is true. The --single-transaction option enforces a consistent backup for all storage engines by using a single transaction to capture the data, making sure that the data dump is consistent across different storage engines.

C. The backup created is a consistent data dump.

This statement is true. The --single-transaction option ensures a consistent data dump by capturing data from a single transaction, preventing inconsistencies due to concurrent write operations.

upvoted 1 times

🗉 **Dhanushka** 1 year, 1 month ago

Selected Answer: AC

--single-transaction

This option sets the transaction isolation mode to REPEATABLE READ and sends a START TRANSACTION SQL statement to the server before dumping data. It is useful only with transactional tables such as InnoDB, because then it dumps the consistent state of the database at the time when START TRANSACTION was issued without blocking any applications.

upvoted 1 times

🗉 **Dhanushka** 1 year, 1 month ago

--single-transaction

This option sets the transaction isolation mode to REPEATABLE READ and sends a START TRANSACTION SQL statement to the server before dumping data. It is useful only with transactional tables such as InnoDB, because then it dumps the consistent state of the database at the time when START TRANSACTION was issued without blocking any applications.

upvoted 1 times

Examine this query and output:

```
mysql> EXPLAIN ANALYZE
SELECT city.CountryCode, country.Name AS Country_Name,
       city.Name, city.District, city.Population
FROM world.city
     INNER JOIN world.country ON country.Code = city.CountryCode
WHERE country.Continent = 'Asia'
     AND city.Population > 1000000
ORDER BY city.Population DESC\G

***** 1. row *****
EXPLAIN:
-> Sort: <temporary>.Population DESC (actual time=8.306..8.431 rows=125 loops=1)
   -> Stream results (actual time=0.145..8.033 rows=125 loops=1)
     -> Nested loop inner join (cost=241.12 rows=205) (actual time=0.141..7.787 rows=125 loops=1)
        -> Filter: (world.country.Continent = 'Asia') (cost=25.40 rows=34) (actual time=0.064..0.820 rows=51 loops=1)
           -> Table scan on country (cost=25.40 rows=239) (actual time=0.059..0.359 rows=239 loops=1)
              -> Filter: (world.city.Population > 1000000) (cost=4.53 rows=6) (actual time=0.030..0.131 rows=2 loops=51)
                 -> Index lookup on city using CountryCode (CountryCode=world.country.`Code`) (cost=4.53 rows=18) (actual
time=0.023..0.096 rows=35 loops=51)

1 row in set (0.0094 sec)
```

Which two statements are true? (Choose two.)

- A. The country table is accessed as the first table, and then joined to the city table.
- B. It takes more than 8 milliseconds to sort the rows.
- C. The optimizer estimates that 51 rows in the country table have Continent = 'Asia'.
- D. 35 rows from the city table are included in the result.
- E. The query returns exactly 125 rows.

Suggested Answer: CD

Community vote distribution

AE (100%)

🗳️ 👤 **agara** 5 days, 1 hour ago

My guess: Answer is B&E

A is wrong because CountryCode was 1st accessed.

C is wrong, 34 rows was estimated.

D is wrong because the 35 rows were filtered out later

upvoted 1 times

🗳️ 👤 **xjllll** 9 months ago

Selected Answer: AE

A not sure

upvoted 1 times

Examine this command, which executes successfully:

```
mysqlpump --user=root --password > full_backup.sql
```

Which two databases will be excluded from this dump? (Choose two.)

- A. information_schema
- B. world
- C. employee
- D. sys
- E. mysql

Suggested Answer: AD

Community vote distribution

AD (100%)

 **Zehir** 10 months, 3 weeks ago

Selected Answer: AD

Mysqump does not dump the performance_schema, ndbinfo, or sys schema by default.

Mysqump does not dump the INFORMATION_SCHEMA schema.

<https://dev.mysql.com/doc/refman/8.0/en/mysqlpump.html#mysqlpump-restrictions>

upvoted 1 times

 **marklv** 11 months ago

AD

Does not dump performance_schema, ndbinfo, or sys, schema_information

upvoted 1 times

Examine this MySQL client command to connect to a remote database: `mysql -h remote.example.org -u root -p --protocol=TCP --ssl-mode=`
Which two `--ssl-mode` values will ensure that an X.509-compliant certificate will be used to establish the SSL/TLS connection to MySQL?

- A. REQUIRED
- B. VERIFY_CA
- C. VERIFY_IDENTITY
- D. PREFERRED
- E. DISABLED

Suggested Answer: A

Community vote distribution

B (100%)

🗨️ 👤 **FelipeK** 1 year ago

Selected Answer: B

The two `--ssl-mode` values that will ensure that an X.509-compliant certificate will be used to establish the SSL/TLS connection to MySQL are `VERIFY_CA` and `VERIFY_IDENTITY`. So, the correct answers are B and C.

upvoted 1 times

🗨️ 👤 **xexeu** 1 year, 4 months ago

correct B and C

upvoted 2 times

🗨️ 👤 **devml** 1 year, 7 months ago

Selected Answer: B

The correct answer should be B and C . Please correct it. The question is asking for 2 ssl option but in answer its selected moreover the required ssl mode allow secure connection and fails if its not secure . But it need not be an x.509 based authentication.

<https://dev.mysql.com/doc/refman/5.7/en/using-encrypted-connections.html>

upvoted 2 times

You want to log only the changes made to the database objects and data on the MySQL system.
Which log will do this by default?

- A. general query log
- B. audit log
- C. slow query log
- D. binary log
- E. error log

Suggested Answer: D

Community vote distribution



  **devml** 1 year, 7 months ago

Selected Answer: D

Answer D is correct binary log. Becoz other logs are useful for finding performance bottlenecks and auditing suspicious activities.
upvoted 2 times

Which two statements are true about using backups of the binary log? (Choose two.)

- A. Multiple binary logs can be used to restore data.
- B. Multiple binary logs can be applied in parallel for faster data restoration.
- C. Binary logs are relatively small, and therefore, excellent for long-term storage and disaster recovery.
- D. Binary logs can always be used to unapply unwanted schema changes.
- E. They allow for point-in-time recovery of the data.

Suggested Answer: *DE*

Community vote distribution

AE (100%)

🗨️ 👤 **FelipeK** 1 year ago

Selected Answer: AE

Binary logs can be used to restore data to a specific point in time. This is known as point-in-time recovery.

Multiple binary logs can be used to restore data, as they contain a record of all changes made to the database

upvoted 2 times

🗨️ 👤 **FelipeK** 7 months, 2 weeks ago

A- Multiple binary logs can be used to restore data. Binary logs record all changes to the database, and multiple logs can be used in sequence to restore the database to a particular point in time1.

E- They allow for point-in-time recovery of the data. Binary logs can be used to perform point-in-time recovery, which involves replaying changes from the logs to the database up to a specific point in time1

upvoted 1 times

Where is the default data directory located after installing MySQL using RPM on Oracle Linux 7?

- A. /usr/mysql
- B. /usr/bin
- C. /etc/my.cnf
- D. /var/lib/mysql
- E. /usr

Suggested Answer: B

Community vote distribution

D (100%)

🗨️ 👤 **FelipeK** 1 year ago

Selected Answer: D

After installing MySQL using RPM on Oracle Linux 7, the default data directory is located at /var/lib/mysql. This is where MySQL stores its database files and related data by default.

upvoted 1 times

🗨️ 👤 **Dhanushka** 1 year, 1 month ago

Selected Answer: D

/var/lib/mysql

upvoted 2 times

🗨️ 👤 **devml** 1 year, 7 months ago

Selected Answer: D

The correct answer would be D. The following are the respective directories

Files or Resources Location

Client programs and scripts /usr/bin

mysqld server /usr/sbin

Configuration file /etc/my.cnf

Data directory /var/lib/mysql

upvoted 2 times

A user wants to connect without entering his or her username and password on the Linux command prompt.
Which three locations can be used to store the user's mysql credentials to satisfy this requirement? (Choose three.)

- A. \$HOME/.my.cnf file
- B. \$MYSQL_HOME/my.cnf file
- C. DATADIR/mysql-auto.cnf file
- D. \$HOME/.mylogin.cnf file
- E. \$HOME/.mysql/auth/login file
- F. /etc/my.cnf file
- G. \$HOME/.mysqlrc file

Suggested Answer: BCF

Community vote distribution

ADF (100%)

🗨️ 👤 **FelipeK** 11 months, 2 weeks ago

Selected Answer: ADF

A. \$HOME/.my.cnf file

This is a common location for storing MySQL client configuration files, including user credentials.

D. \$HOME/.mylogin.cnf file

The mylogin.cnf file is used to securely store credentials for MySQL client programs. It is designed for this purpose.

F. /etc/my.cnf file

The /etc/my.cnf file can be used to store MySQL client configuration options, including credentials, but it's typically used for global MySQL configuration and not individual user credentials.

upvoted 1 times

🗨️ 👤 **marklv** 1 year ago

A is true. Not sure of B.

upvoted 1 times

t is a non-empty InnoDB table.

Examine these statements, which are executed in one session:

```
BEGIN;
```

```
SELECT * FROM t FOR UPDATE;
```

Which is true?

- A. If `OPTIMIZE TABLE`; is invoked, it will create a table lock on t and force a transaction rollback.
- B. If `OPTIMIZE LOCAL TABLE t`; is invoked from another session, it executes normally and returns the status.
- C. `mysqlcheck --analyze --all-databases` will execute normally on all tables and return a report.
- D. If `ANALYZE TABLE`; is invoked from the same session, it hangs until the transaction is committed or rolled back.

Suggested Answer: B

Community vote distribution

C (100%)

🗨️ 👤 **FelipeK** 7 months, 3 weeks ago

Selected Answer: C

The right answer is C

upvoted 1 times

🗨️ 👤 **xjllll** 9 months ago

Selected Answer: C

C is right

upvoted 2 times

🗨️ 👤 **marklv** 1 year ago

B is not right. D is. It locks until the first connection is done.

upvoted 2 times

🗨️ 👤 **marklv** 2 months ago

Sorry, if same session, D will not lock. D isn't right either

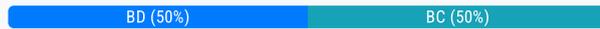
upvoted 1 times

Which two MySQL Server accounts are locked by default? (Choose two.)

- A. any user set as DEFINER for stored programs
- B. any internal system accounts
- C. any new ROLE accounts
- D. any user created without a password
- E. any user created with a username, but missing the host name

Suggested Answer: AB

Community vote distribution



🗨️ 👤 **marklv** 11 months ago

BC

role and internal accounts

upvoted 1 times

🗨️ 👤 **FelipeK** 11 months, 2 weeks ago

Selected Answer: BD

B. any internal system accounts

D. any user created without a password

upvoted 1 times

🗨️ 👤 **Dhanushka** 1 year, 1 month ago

Selected Answer: BC

Definer can be application user as well hence it will not be lock

upvoted 1 times

Examine this SQL statement:

```
UPDATE world.city
SET Population = Population * 1.1
WHERE CountryCode IN (SELECT Code FROM world.country
                      WHERE Continent = 'Asia')
```

Which set of privileges will allow Tom to execute this SQL statement?

- A. GRANT ALL PRIVILEGES ON 'world'.city TO 'tom'@'%';
GRANT SELECT ('code') ON 'world'.country TO 'tom'@'%';
- B. GRANT UPDATE ON 'world'.* TO 'tom'@'%';
GRANT ALL PRIVILEGES ON 'world'.country TO 'tom'@'%';
- C. GRANT UPDATE ON 'world'.city TO 'tom'@'%';
GRANT SELECT ON 'world'.* TO 'tom'@'%';
- D. GRANT UPDATE ON 'world'.city TO 'tom'@'%';
GRANT SELECT ON 'world'.country TO 'tom'@'%';

Suggested Answer: D

Community vote distribution

C (100%)

 **LrnsTgh** 8 months, 1 week ago

Selected Answer: C

agree with C. I took the least privileges only.

C -> Query OK, 1766 rows affected (0.0211 sec)

Rows matched: 1766 Changed: 1766 Warnings: 0

D -> ERROR: 1143: SELECT command denied to user 't***h'@'localhost' for column 'CountryCode' in table 'city'
upvoted 1 times

 **marklv** 9 months, 1 week ago

Sorry for changing my answers.

C. read and write to city, read from country, and two columns from country need to be read. Only C works.

A won't work because you need to read from two columns in country, not just one.

B won't work because you need to read from city.

D won't work because you need to read from city as well.

upvoted 1 times

 **marklv** 11 months ago

city need to read and write, it needs to read CountryCode. Thus A, all privs includes read.

BCD and all work if you just need to update city and read from country.

upvoted 1 times

 **FelipeK** 11 months, 2 weeks ago

Selected Answer: C

The right answer is C

upvoted 2 times

 **marklv** 1 year ago

Not sure why B and C won't work. Especially C, since it includes D

upvoted 1 times

Examine these commands and output:

```
mysql> SHOW FULL PROCESSLIST;
+-----+-----+-----+-----+-----+
| Id | User      | ... | State                               | Info                               |
+-----+-----+-----+-----+-----+
| 6  | event_scheduler | ... | Waiting on empty queue             | NULL                               |
| 20 | root      | ... |                                     | NULL                               |
| 21 | root      | ... |                                     | NULL                               |
| 22 | root      | ... | Waiting for table metadata lock     | optimize table test.demo_test     |
| 24 | root      | ... | Waiting for table metadata lock     | select * from test.demo_test      |
| 25 | root      | ... | starting                            | SHOW FULL PROCESSLIST             |
+-----+-----+-----+-----+-----+

mysql> SELECT object_type, object_schema, object_name, lock_type, lock_status, owner_thread_id, owner_event_id
-> FROM performance_schema.metadata_locks WHERE object_schema != 'performance_schema';
+-----+-----+-----+-----+-----+-----+-----+
| OBJECT_TYPE | OBJECT_SCHEMA | OBJECT_NAME | LOCK_TYPE           | LOCK_STATUS | OWNER_THREAD_ID | OWNER_EVENT_ID |
+-----+-----+-----+-----+-----+-----+-----+
| TABLE      | test          | demo_test   | SHARED_READ        | GRANTED     | 60              | 7              |
| TABLE      | test          | demo_test   | SHARED_WRITE       | GRANTED     | 60              | 9              |
| SCHEMA      | test          | NULL        | INTENTION_EXCLUSIVE | GRANTED     | 62              | 6              |
| TABLE      | test          | demo_test   | SHARED_NO_READ_WRITE | PENDING     | 62              | 6              |
+-----+-----+-----+-----+-----+-----+-----+

mysql> SELECT thread_id, processlist_id, processlist_user, parent_thread_id
-> FROM performance_schema.threads WHERE processlist_user='root';
+-----+-----+-----+-----+
| THREAD_ID | PROCESSLIST_ID | PROCESSLIST_USER | PARENT_THREAD_ID |
+-----+-----+-----+-----+
| 60        | 20              | root              | NULL              |
| 61        | 21              | root              | NULL              |
| 62        | 22              | root              | 1                 |
| 64        | 24              | root              | 1                 |
| 65        | 25              | root              | NULL              |
+-----+-----+-----+-----+
```

Which connection ID is holding the metadata lock?

- A. 20
- B. 24
- C. 21
- D. 25
- E. 22
- F. 6

Suggested Answer: D

 **marklv** 9 months, 1 week ago

A. 20 is holding read and write locks which blocks the other commands.
upvoted 2 times

 **marklv** 1 year ago

It is not D, it wouldn't block access to the table. Its either id 20 or 21
upvoted 1 times

Which two are true about binary logs used in asynchronous replication? (Choose two.)

- A. They are pushed from the master to the slave.
- B. They contain events that describe database changes on the master.
- C. They contain events that describe all queries run on the master.
- D. They contain events that describe only administrative commands run on the master.
- E. They are pulled from the master to the slave.

Suggested Answer: AB

Community vote distribution

BE (100%)

 **roobenj** 7 months, 1 week ago

Selected Answer: BE

A is wrong. Replica pulls from master.

upvoted 1 times

 **FelipeK** 11 months, 2 weeks ago

Selected Answer: BE

B. They contain events that describe database changes on the master.

Binary logs in MySQL contain events that describe changes to the data in the master database. These changes can include INSERTs, UPDATEs, DELETEs, and other data modifications.

E. They are pulled from the master to the slave.

In asynchronous replication, the slave server typically pulls the binary logs from the master server. The slave connects to the master and requests the binary log events it needs to replicate changes from the master.

Option A is not entirely accurate because binary logs are not "pushed" from the master to the slave; instead, they are made available on the master for the slave to pull.

Option C and D are incorrect because binary logs do not necessarily contain all queries or administrative commands run on the master; they specifically contain events related to database changes, not all SQL queries or administrative commands.

upvoted 1 times

 **Dhanushka** 1 year, 1 month ago

Selected Answer: BE

Each replica that connects to the source requests a copy of the binary log. That is, it pulls the data from the source, rather than the source pushing the data to the replica. The replica also executes the events from the binary log that it receives.

upvoted 2 times

Which two queries are examples of successful SQL injection attacks? (Choose two.)

- A. `SELECT user,passwd FROM members - WHERE user = '?';INSERT INTO members('user','passwd') VALUES ('bob@example.com','secret');--;`
- B. `SELECT id, name FROM user WHERE user.id=(SELECT members.id FROM members);`
- C. `SELECT id, name FROM user WHERE id=23 OR id=32 OR 1=1;`
- D. `SELECT id, name FROM user WHERE id=23 OR id=32 AND 1=1;`
- E. `SELECT email,passwd FROM members WHERE email = 'INSERT INTO members('email','passwd') VALUES ('bob@example.com', 'secret');--';`
- F. `SELECT user, phone FROM customers WHERE name = '\; DROP TABLE users; --';`

Suggested Answer: DE

Community vote distribution

AE (50%)

AC (50%)

🗨️ 👤 **LrmsTgh** 8 months, 1 week ago

Selected Answer: AE

F only possible if session has all privileges;
upvoted 1 times

🗨️ 👤 **xjlllll** 9 months ago

Selected Answer: AC

A and C
upvoted 1 times

🗨️ 👤 **marklv** 1 year ago

E and F, the other two don't write to the database. A might be valid. Injection is a query inside another query.
upvoted 1 times

You execute this command:

```
shell> mysqlpump --exclude-databases=% --users
```

Which statement is true?

- A. It creates a logical backup of all MySQL user accounts.
- B. It creates a logical backup of all metadata, but contains no table data.
- C. It returns an error because the mysqldump command should have been used.
- D. It creates a logical backup of only the users database.

Suggested Answer: B

Community vote distribution

A (100%)

🗨️ **FelipeK** 11 months, 2 weeks ago

Selected Answer: A

The correct answer should be A

upvoted 1 times

🗨️ **devml** 1 year, 7 months ago

Selected Answer: A

The correct answer should be A. It only take the dump of user accounts.

upvoted 2 times

Which two are contained in the InnoDB system tablespace (ibdata1) by default? (Chose two.)

- A. table data
- B. primary indexes
- C. user privileges
- D. InnoDB Data Dictionary
- E. change buffer
- F. doublewrite buffer

Suggested Answer: DF

Community vote distribution

AE (75%)

EF (25%)

 **roobenj** 7 months, 1 week ago

Selected Answer: AE

A and E. See documentation: <https://dev.mysql.com/doc/refman/8.0/en/innodb-system-tablespace.html>
upvoted 2 times

 **xjllll** 9 months ago

Selected Answer: EF

InnoDB system tablespace contains the doublewrite buffer and the change buffer.
• Each database (including the mysql system database) has a single directory under the data directory that contains storage-engine specific data files for the database.
– For example: tablename.ibd for InnoDB, which contains the table data and metadata
upvoted 1 times

 **FelipeK** 11 months, 2 weeks ago

Selected Answer: AE

In previous MySQL versions, the system tablespace contained the InnoDB data dictionary, but in MySQL 8.0, InnoDB stores metadata in the MySQL data dictionary.
The doublewrite buffer storage area also resided in the system tablespace in previous MySQL releases, but this storage area resides in separate doublewrite files as of MySQL 8.0.201.
upvoted 1 times

 **FelipeK** 7 months, 3 weeks ago

<https://dev.mysql.com/doc/refman/8.0/en/innodb-system-tablespace.html>
upvoted 1 times

 **marklv** 1 year ago

A and B. data directory contains other things. double write now has its own files.
upvoted 1 times

 **marklv** 9 months, 1 week ago

E is stored in tablespace. Not sure why B isn't correct.
upvoted 1 times

Which two MySQL Shell commands are excluded from the InnoDB Cluster creation procedure? (Choose two.)

- A. `dba.configureInstance()`
- B. `cluster.setPrimaryInstance()`
- C. `dba.configureLocalInstance()`
- D. `cluster.forceQuorumUsingPartitionOf()`
- E. `cluster.addInstance()`
- F. `dba.createCluster()`
- G. `dba.checkInstanceConfiguration()`

Suggested Answer: *BD*

Community vote distribution

AB (100%)

 **dmees** 6 months, 3 weeks ago

Selected Answer: AB

- A. `dba.configureInstance()`
- B. `cluster.setPrimaryInstance()`

These commands are not used during the creation of an InnoDB Cluster.

The `dba.configureInstance()` command is used to configure an instance for InnoDB Cluster after it has been added to the cluster.

The `cluster.setPrimaryInstance()` command is used to manually set a specific instance as the primary instance in the cluster.

upvoted 1 times

 **marklv** 9 months, 1 week ago

- D. not used in creation. Used if there is an issue later.
 - G. not used because if there are no errors, you can check it after its been created.
- upvoted 1 times

Examine this statement:

```
mysql> DROP ROLE r_role1, r_role2;
```

Which two are true? (Choose two.)

- A. It fails if any of the roles is specified in the mandatory_roles variable.
- B. You must revoke r_role1 and r_role2 from all users and other roles before dropping the roles.
- C. Existing connections can continue to use the roles' privileges until they reconnect.
- D. You must revoke all privileges from r_role1 and r_role2 before dropping the roles.
- E. It fails if you do not have the ADMIN OPTION of the roles r_role1 and r_role2.
- F. It fails if at least one of the roles does not exist.

Suggested Answer: CE

Community vote distribution

AF (50%)

BC (50%)

🗳️ 👤 **roobenj** 7 months, 1 week ago

Selected Answer: AF

Tested it. A and F are correct options.

upvoted 1 times

🗳️ 👤 **FelipeK** 7 months, 2 weeks ago

Selected Answer: AF

A. It fails if any of the roles is specified in the mandatory_roles variable.

Roles named in the mandatory_roles system variable value cannot be dropped.

F. It fails if at least one of the roles does not exist.

By default, an error occurs if you try to drop a role that does not exist. However, if the IF EXISTS clause is given, the statement produces a warning for each named role that does not exist, rather than an error.

upvoted 1 times

🗳️ 👤 **marklv** 11 months ago

correction, A and F. You don't need to reload connection if role changes.

upvoted 1 times

🗳️ 👤 **marklv** 1 year ago

I tested it, you can drop a role if users have that role.

You need grant and revoke privs to manipulate roles.

I also tested deleting a role that didn't exist, it errored.

C and F.

upvoted 1 times

🗳️ 👤 **dieymir** 1 year, 2 months ago

A and F

upvoted 1 times

🗳️ 👤 **FelipeK** 1 year, 5 months ago

Selected Answer: BC

B. This is true. Before dropping a role, you need to make sure that the role is not assigned to any user or any other role. Otherwise, the DROP ROLE statement will fail with an error message.

C. This is also true. Existing connections that are currently using the privileges of the roles being dropped can continue to do so until they disconnect and reconnect.

upvoted 2 times

🗳️ 👤 **FelipeK** 7 months, 2 weeks ago

correction A and F

upvoted 1 times

Examine these statements and output:

```
mysql> GRANT PROXY ON accounting@localhost TO '@'%;
```

```
mysql> SELECT USER(), CURRENT_USER(), @@proxy_user;
```

USER()	CURRENT_USER()	@@proxy_user
rsmith@localhost	accounting@localhost	'@'%'

Which statement is true?

- A. The user is logged in with --user=accounting as an option.
- B. The user is authenticated as the anonymous proxy user '@%'.
- C. The user is authorized as the accounting@localhost user.
- D. The user is authorized as the rsmith@localhost user.
- E. The user failed to define a username and the connecting username defaulted to '@%'.

Suggested Answer: D

Community vote distribution

B (67%)

C (33%)

🗨️ **xjlll** 9 months ago

Selected Answer: C

C is right

upvoted 1 times

🗨️ **marklv** 1 year ago

incorrect. You authenticate as your own account but are authorized as another. There B is not true, C is.

upvoted 2 times

🗨️ **FelipeK** 1 year, 5 months ago

Selected Answer: B

Explanation:

The first statement grant proxy on accounting@localhost to '@%' grants the permission to the anonymous proxy user '@%' to use the accounting@localhost user as a proxy.

The second statement Select user(),current_user(), @@proxy_user returns the output with three columns:

USER() returns the current user as rsmith@localhost.

CURRENT_USER() returns the authenticated user as accounting@localhost, which is the user that was granted the proxy permission in the first statement.

@@proxy_user returns the current anonymous proxy user as '@%'.

Therefore, statement B is true, which states that the user is authenticated as the anonymous proxy user '@%'.

upvoted 2 times

Which two statements are true about InnoDB data-at-rest encryption? (Choose two.)

- A. It supports only non-blob datatypes.
- B. It does not support the transportable tablespaces feature.
- C. It supports all indexes transparently.
- D. It decrypts data for use in memory.
- E. It enforces encryption from disk to memory and over network transmission.

Suggested Answer: CD

Community vote distribution



🗨️ **JESUSBB** 8 months, 2 weeks ago

InnoDB data-at-rest encryption supports all indexes transparently and With InnoDB data-at-rest encryption, in-memory data is decrypted, which provides complete transparency.

ANS: C and D

<https://dev.mysql.com/doc/refman/8.0/en/faqs-tablespace-encryption.html>

upvoted 1 times

🗨️ **FelipeK** 10 months, 3 weeks ago

Selected Answer: CD

C. It supports all indexes transparently.

D. It decrypts data for use in memory.

upvoted 1 times

🗨️ **marklv** 11 months ago

can't be encrypted over network, where is the documentation for that? I guess C and D. D because if decrypts as needed "for use".

upvoted 1 times

🗨️ **FelipeK** 1 year, 5 months ago

Selected Answer: CE

Explanation:

InnoDB data-at-rest encryption provides built-in encryption for data at the tablespace level, which is enforced from disk to memory and over network transmission, ensuring that the data is encrypted at all times.

It supports all indexes transparently, including primary, secondary, unique, full-text, and spatial indexes, which are all encrypted along with the table data.

InnoDB data-at-rest encryption supports all datatypes, including blob datatypes.

InnoDB data-at-rest encryption does not affect the transportable tablespaces feature, which can be used to move InnoDB tables between servers, as long as the tablespace encryption key is available on the destination server.

InnoDB data-at-rest encryption does not decrypt data for use in memory; instead, it keeps the data encrypted in memory until it is required by the user, at which point it is decrypted temporarily for processing.

upvoted 1 times

You plan to install MySQL Server by using the RPM download.

Which two statements are true? (Choose two.)

- A. You can provide the root password interactively.
- B. You must manually initialize the data directory.
- C. The MySQL RPM package installation supports deploying multiple MySQL versions on the same host.
- D. MySQL uses the RPM relocatable installation target feature.
- E. The functionality is split among several RPM package files.
- F. You can find the root password in the error log after the first start.

Suggested Answer: CF

Community vote distribution

EF (50%)

BE (50%)

🗨️ 👤 **LmsTgh** 8 months, 1 week ago

Selected Answer: EF

agree E & F

upvoted 1 times

🗨️ 👤 **FelipeK** 1 year, 5 months ago

Selected Answer: BE

Explanation:

A. You cannot provide the root password interactively during the installation of the MySQL RPM package. Instead, you can set the root password after the installation is complete by using the `mysql_secure_installation` command.

B. After you install the MySQL RPM package, you must manually initialize the data directory by running the `mysqld --initialize` command.

C. The MySQL RPM package installation does not support deploying multiple MySQL versions on the same host.

D. MySQL does not use the RPM relocatable installation target feature.

E. The functionality of MySQL is split among several RPM package files, including the `mysql-server`, `mysql-lib`, and `mysql-client` packages.

F. The root password is not written to the error log after the first start. Instead, you can set the root password after the installation is complete by using the `mysql_secure_installation` command.

upvoted 1 times

🗨️ 👤 **FelipeK** 10 months, 3 weeks ago

E. The functionality is split among several RPM package files. The MySQL RPM packages include `mysql-community-client`, `mysql-community-server`, `mysql-community-lib`, and others¹.

F. You can find the root password in the error log after the first start. After the first start, the root password is generated and stored in the error log.

upvoted 1 times

You plan to take daily full backups, which include the ndbinfo and sys (internal) databases. Which command will back up the databases in parallel?

- A. `mysqldump --single-transaction > full-backup-$(date +%Y%m%d).sql`
- B. `mysqlpump --include-databases=% > full-backup-$(date +%Y%m%d).sql`
- C. `mysqlpump --all-databases > full-backup-$(date +%Y%m%d).sql`
- D. `mysqldump --all-databases > full_backup-$(date +%Y%m%d).sql`

Suggested Answer: D

Community vote distribution

B (100%)

🗨️ **marklv** 1 year ago

<https://dev.mysql.com/doc/refman/8.0/en/mysqldump.html>

mysqldump does not dump ndbinfo, so D is incorrect

upvoted 1 times

🗨️ **marklv** 1 year ago

B might be correct because the wildcard might be explicitly listing it. Otherwise mysqlpump won't back it up.

upvoted 1 times

🗨️ **marklv** 1 year ago

I tested, it does.

upvoted 2 times

🗨️ **HSong** 1 year, 1 month ago

Should the answer be B?

mysqlpump does not dump the performance_schema, ndbinfo, or sys schema by default. To dump any of these, name them explicitly on the command line. You can also name them with the --databases or --include-databases option.

<https://dev.mysql.com/doc/refman/8.0/en/mysqlpump.html#mysqlpump-restrictions>

upvoted 1 times

🗨️ **devml** 1 year, 7 months ago

Selected Answer: B

The correct answer is B. mysqlpump will allow you to backup the databases parallelly .

<https://dev.mysql.com/doc/refman/8.0/en/mysqlpump.html>

By default --all-databases won't backup ndbinfo schema this can be done via --include-database=%

<https://dev.mysql.com/doc/refman/8.0/en/mysqlpump.html#mysqlpump-restrictions>

upvoted 4 times

What does the binlog dump thread do?

- A. It monitors and schedules the rotation/deletion of the binary logs.
- B. It reads the relay log and executes the events contained in them.
- C. It acquires a lock on the binary log for reading each event to be sent to the slave.
- D. It connects to the master and asks it to send updates recorded in its binary logs.

Suggested Answer: *D*

Community vote distribution

C (100%)

🗨️ 👤 **marklv** 1 year ago

correct, C

upvoted 1 times

🗨️ 👤 **HSong** 1 year, 1 month ago

Selected Answer: C

Answer is C

Binary log dump thread. The source creates a thread to send the binary log contents to a replica when the replica connects. This thread can be identified in the output of SHOW PROCESSLIST on the source as the Binlog Dump thread.

The binary log dump thread acquires a lock on the source's binary log for reading each event that is to be sent to the replica. As soon as the event has been read, the lock is released, even before the event is sent to the replica.

<https://dev.mysql.com/doc/refman/8.0/en/replication-threads.html>

upvoted 2 times

Which two commands will display indexes on the parts table in the manufacturing schema? (Choose two.)

- A. DESCRIBE manufacturing.parts;
- B. SELECT * FROM information_schema.statistics WHERE table_schema='manufacturing' AND TABLE_NAME='parts';
- C. SHOW INDEXES FROM manufacturing.parts;
- D. SELECT * FROM information_schema.COLUMN_STATISTICS;
- E. EXPLAIN SELECT INDEXES FROM manufacturing.parts;

Suggested Answer: *BD*

Community vote distribution

BC (100%)

🗨️ 👤 **FelipeK** 10 months, 3 weeks ago

Selected Answer: BC

These commands will provide information about the indexes on the specified table in the given schema.

The SHOW INDEXES command is a simple way to view the index information of a table, while the SELECT statement from the information_schema.statistics table allows for more complex queries and can be used to get index information across multiple tables or schemas. Please note that you need to have the appropriate privileges to run these commands.

upvoted 1 times

🗨️ 👤 **HSong** 1 year, 1 month ago

Selected Answer: BC

Answer is B,C

upvoted 2 times

Your MySQL server is running on the Microsoft Windows platform.
Which three local connection protocols are available to you? (Choose three.)

- A. named pipes
- B. shared memory
- C. SOCKET
- D. X Protocol
- E. UDP
- F. TCP/IP

Suggested Answer: ABF

Community vote distribution

ABF (100%)

  **devml** 1 year, 7 months ago

Selected Answer: ABF

The correct answer is A,B,F.

<https://dev.mysql.com/doc/refman/8.0/en/transport-protocols.html>

upvoted 3 times

Which two statements are true about using MySQL Enterprise Monitor Query Analyzer? (Choose two.)

- A. The single query QRTi pie chart in the Query Analyzer view is based on the average execution of all statements.
- B. It is possible to retrieve a normalized statement, but never the exact statement that was executed.
- C. It is possible to configure the Query Analysis built-in advisor to get notified about slow query execution.
- D. It is possible to list and analyze statements in an arbitrary graph range selection from timeseries graphs.
- E. It is possible to import data into the Query Analyzer from heterogeneous sources, such as CSV.

Suggested Answer: *CD*

  **vinhv** 10 months, 4 weeks ago

A maybe one of the answers:

<https://dev.mysql.com/doc/mysql-monitor/8.0/en/mem-features-qrti.html>

upvoted 1 times

Which two are use cases of MySQL asynchronous replication? (Choose two.)

- A. You can scale writes by creating a replicated mesh.
- B. It guarantees near real-time replication between a master and a slave.
- C. You can scale reads by adding multiple slaves.
- D. MySQL Enterprise Backup will automatically back up from an available slave.
- E. It allows backup to be done on the slave without impacting the master.

Suggested Answer: CE

Community vote distribution

CE (100%)

🗨️ 👤 **FelipeK** 7 months, 1 week ago

Selected Answer: CE

Option C: Asynchronous replication allows you to add multiple slaves to distribute read queries, thus scaling reads horizontally.

Option E: With asynchronous replication, you can perform backups on the slave without impacting the master, ensuring that backup processes do not affect the performance of the primary database operations.

upvoted 1 times

🗨️ 👤 **marklv** 11 months ago

C and E

upvoted 1 times

Examine this list of MySQL data directory binary logs:

```
binlog.000001
binlog.000002
.....
binlog.000289
binlog.000300
binlog.000301
binlog.index
```

Now examine this command, which executes successfully:

```
mysqldump --delete-master-logs --all-databases > /backup/db_backup.sql
```

Which two are true? (Choose two.)

- A. All databases are backed up to the output file.
- B. All non-active binary logs are removed from the master.
- C. All binary logs are deleted from the master.
- D. All binary logs are backed up and then deleted.
- E. All databases, excluding master metadata, are backed up to the output file.
- F. All details regarding deleted logs and master metadata are captured in the output file.

Suggested Answer: CE

Community vote distribution

AC (67%)

AD (33%)

🗨️ **dmees** 6 months, 3 weeks ago

AC.

D. All binary logs are backed up and then deleted (FALSE): The mysqldump command doesn't have the capability to back up binary logs. It only deals with backing up database schemas and data.

upvoted 1 times

🗨️ **LrnsTgh** 8 months, 1 week ago

Selected Answer: AC

AC.

>8.0.26 use --delete-source-logs

upvoted 1 times

🗨️ **FelipeK** 10 months, 3 weeks ago

Selected Answer: AC

A-All databases are backed up to the output file. This is true because the --all-databases option tells mysqldump to dump all databases.

C-All binary logs are deleted from the master. This is true because the --delete-master-logs option tells mysqldump to delete all binary logs on the master server2. However, it's important to note that deleting the MySQL binary logs with mysqldump --delete-master-logs can be dangerous if your server is a replication source server, because replicas might not yet fully have processed the contents of the binary log

upvoted 1 times

🗨️ **marklv** 1 year ago

incorrect not D, it doesn't backup them up

upvoted 1 times

🗨️ **marklv** 11 months ago

B, non active logs i think, there is always one active log, after mysqldump it rotates the log and will purge the old ones

upvoted 1 times

🗨️ **marklv** 2 months ago

sorry me again,

A. All databases are backed up to the output file.

No, sys, information_schema, performance_schema are not backed up.

But I think they mean yes to this.

B. All non-active binary logs are removed from the master.

Yes, the last binlog is active and does not get deleted.

C. All binary logs are deleted from the master.

No, last one is not. Tested this.

D. All binary logs are backed up and then deleted.

No, they are not backed up.

E. All databases, excluding master metadata, are backed up to the output file.

Maybe, if sys, performance_schema, information_schema are all meta, then yes, otherwise no.

My guess is no. Performance_schema is not meta.

F. All details regarding deleted logs and master metadata are captured in the output file.

No.

upvoted 1 times

  **FelipeK** 1 year, 5 months ago

Selected Answer: AD

https://dev.mysql.com/doc/refman/8.0/en/mysqldump.html#option_mysqldump_delete-master-logs

upvoted 1 times

  **marklv** 9 months, 1 week ago

A. yes, all databases are backedup.

Not D, it doesn't backup binary logs.

B or C, probably B, because there might logs made after the dump is started. So you are not deleting all logs.

upvoted 1 times