



Free Questions for 1Z0-888 by [braindumpscollection](#)

Shared by [Rosales](#) on 15-04-2024

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

Question 1

Question Type: MultipleChoice

You are using the Performance Schema to investigate replication on a slave which has a single master. The option `slave-parallel-type` is set to `DATABASE`.

```
mysql> SELECT THREAD_ID, threads.NAME, SUM(COUNT_STAR) AS TotalCount, SUM
(SUM_TIMER_WAIT) AS TotalTime
  → FROM
    performance_schema.events_waits_summary_by_thread_by_event_name
  → INNER JOIN performace_schema.threads USING (THREAD_ID)
  → WHERE threads.NAME LIKE 'thread/sql/slave\_%'
  → GROUP BY THREAD_ID, threads.NAME;
```

THREAD_ID	NAME	TotalCount	TotalTime
20	thread/sql/slave_io	5785	654785731198
21	thread/sql/slave_sql	3875	96931638913
22	thread/sql/slave_worker	0	0
23	thread/sql/slave_worker	0	0
24	thread/sql/slave_worker	346730	7262131209667
25	thread/sql/slave_worker	597127	15498842906584

Assume that all instruments and consumers are enabled and all threads are instrumented.

Which two facts can be concluded from the given output? (Choose two.)

Options:

- A- The slave has two intermediate relay slaves connected to it.
- B- The slave is configured with `slave_parallel_workers = 4`
- C- At most, two schemas are being updates concurrently.
- D- `THREAD_ID 21` has stopped running.
- E- The slave cannot process the relay log fast enough to use all threads.
- F- The server needs more cores to use all slave threads.

Answer:

B, E

Question 2

Question Type: MultipleChoice

Which two statements are true regarding the creating of new MySQL physical and logical backups? (Choose two.)

Options:

- A- Physical backups can be used to recover from data corruption.
- B- Logical backups are human-readable whereas physical backups are not.
- C- Logical backups are always larger than physical backups.
- D- Physical backups are usually slower than text backups.
- E- Physical backups are usually faster than text backups.

Answer:

A, E

Question 3

Question Type: MultipleChoice

The MySQL error log shows:

InnoDB: Warning: a long semaphore wait:

The relevant parts of the InnoDB monitor output shows:

```
--Thread 140259946129152 has waited at btr0sea.cc line 658 for
241.00 seconds the semaphore:

X-lock (wait_ex) on RW-latch at 0x2a5581378 created in file
btr0sea.cc line 173 a writer (thread id 140259946129152) has
reserved it in mode wait exclusive number of readers 1, waiters
flag 1, lock_word: ffffffff

Last time read locked in file btr0sea.cc line 907
Last time write locked in file /pb2/build/sb_0-10188268-
1378799520.26/rpm/BUILD/mysqlcom-pro-5.7.14/mysqlcom-pro-
5.7.14/storage/innobase/btr/btr0sea.cc line 658

...
---TRANSACTION 1935115BA, ACTIVE 942 sec, process no 20643, OS
thread id 140223541274368

mysql tables in use 3, locked 0
, holds adaptive hash latch

MySQL thread id 3631102, query id 141949524 localhost 127.0.0.1
world Waiting for query cache lock

...
```

Which two options would help avoid the long wait in the future?

Options:

- A- Increase the value of the `innodb_lock_wait_timeout` option.
- B- Increase the value of the `innodb_read_io_threads` option.

- C-** Change the table to use HASH indexes instead of BTREE indexes.
- D-** Set the value of innodb_adaptive_hash_index to zero.
- E-** Deactivate the query cache.
- F-** Increase the size of the InnoDB buffer pool.

Answer:

B, F

Question 4

Question Type: MultipleChoice

A MySQL server was initialized with separate UNDO tablespaces. Users complain that when they roll back large transactions, the time to process the request takes too long. The DBA would like to move the MySQL InnoDB UNDO tablespace to a solid-state drive (SSD) for better performance. Is this possible and how?

Options:

- A-** Yes. Shut down the mysql process, enable the transportable_tablespace option, and move the UNDO directory to the SSD.

- B-** Yes. Shut down, copy the UNDO tablespaces to the new location, and change the innodb_undo_directory value in your my.cnf.
- C-** No. The UNDO tablespaces must remain on the same file system as the system tablespaces.
- D-** No. The sequential write pattern of the UNDO tablespaces is not supported on modern SSD block devices.

Answer:

C

Question 5

Question Type: MultipleChoice

The following grants were executed:

```
GRANT CREATE ROUTING ON sales.* TO 'webadmin'@'%';
```

```
GRANT ALTER ON PROCEDURE sales.myproc TO 'webadmin'@'%';
```

A user successfully connects to the database as webadmin and created a stored procedure named get_reports. The next day, the user logs in again as webadmin and wants to delete the stored procedure named get_reports, and therefore, issues the following statement:

```
USE sales;
```

```
DROP PROCEDURE IF EXISTS get_reports;
```


What is the result of executing the statement?

Options:

- A- The user will get an error because he or she did not use the ALTER statement to drop the stored procedure.
- B- The user will get an error because he or she did not put the database name in front of the stored procedure name.
- C- The stored procedure named get_reports will be dropped.
- D- The user will get an error because he or she does not have the permission to drop stored procedures.

Answer:

C

Question 6

Question Type: MultipleChoice

A single InnoDB table has been dropped by accident. You are unable to use an additional intermediate MySQL instance to restore the table. Which two backup methods can be used to restore the single table without stopping the MySQL instance? (Choose two.)

Options:

- A- a backup created with mysqldump --all-databases
- B- a backup created using FLUSH TABLES ... FOR EXPORT
- C- an up-to-date replication slave
- D- a file system-level snapshot
- E- a file system copy created while MySQL was shut down.

Answer:

A, D

Question 7

Question Type: MultipleChoice

Consider the join_buffer_size parameter in MySQL Server.

Which two statements are true about the join buffer? (Choose two.)

Options:

- A- The value should be increased if the client performs several SELECT operations.
- B- The join buffer is set per connection.
- C- The join buffer is used to process sorts when complex joins are being performed.
- D- The value should be increased from the default if the query joins large rows without using an index.
- E- The join buffer is global and can be changed only by restarting the server.

Answer:

C, D

Question 8

Question Type: MultipleChoice

When you examine a new MySQL installation with default configuration, you find a file called ibdata1 in the database directory. Which two statements are true about this file? (Choose two.)

Options:

- A- it contains the binary log.

- B-** it contains a general tablespace.
- C-** it is the default location for all new tables that you create.
- D-** it contains the system tablespace.
- E-** it contains the redo log.
- F-** it contains the undo log.

References:

Answer:

C, D

Question 9

Question Type: MultipleChoice

old_alter_table is disabled as shown.

```
mysql> SELECT @@old_alter_table;
```

```
+-----+
| @@old_alter_table |
+-----+
| 0 |
+-----+
1 row in set (0.00 sec)
```

Consider this statement on a RANGE-partitioned table:

```
mysql> ALTER TABLE orders DROP PARTITION p1, p3;
```

What is the outcome of executing this statement?

Options:

- A-** All data in p1 and p3 partitions is removed and the table definition is changed.
- B-** All data in p1 and p3 partitions is removed, but the table definition remains unchanged.
- C-** Only the first partition (p1) will be dropped because only one partition can be dropped at any time.
- D-** It results in a syntax error because you cannot specify more than one partition in the same statement.

Answer:

B

To Get Premium Files for 1Z0-888 Visit

<https://www.p2pexams.com/products/1z0-888>

For More Free Questions Visit

<https://www.p2pexams.com/oracle/pdf/1z0-888>

