

Oracle 1Z0-918

MySQL Database Administration Professional

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/1z0-918>

Latest Version: 6.0

Question: 1

On a SaaS platform's server, an analyst opened a REPEATABLE READ transaction, ran a SELECT, and left the session idle for several hours. Heavy updates continued on the same tables throughout. The DBA sees the InnoDB history list length climbing steadily and the undo tablespaces growing.

Which two statements are true?

(Choose two.)

- A. Rerunning the SELECT in that session returns rows as of its first read
- B. The idle transaction holds shared locks on the rows it read, which is what slows the updates
- C. Purge cannot discard undo records that the idle transaction's read view may still need, so history keeps accumulating until it ends
- D. Old row versions live in the redo log, which grows
- E. Adding purge threads lets purge remove records that the idle read view still needs

Answer: A,C

Question: 2

A university's student records server stores its binary logs on the same volume as its data. The DBA must estimate how much disk the volume will need in twelve months, covering both table growth and binary log usage.

Which approach gives a sound estimate?

- A. Track DATA_LENGTH plus INDEX_LENGTH in information_schema.TABLES over several weeks, and add daily binary log volume times the retention period
- B. Multiply the current COUNT(*) of each table by the expected row growth, and leave out binary logs because the server purges them after every full backup
- C. Project growth from Innodb_buffer_pool_pages_data, and size the binary logs from max_binlog_size
- D. Use today's datadir size as the baseline, and add today's binary log files as a fixed allowance

Answer: A

Question: 3

Players of a game studio's online title intermittently fail to log in, and the application logs show Too many connections. During an incident, SHOW GLOBAL STATUS reports Threads_connected equal to the configured max_connections value of 800 and Threads_running of 6. Most rows in the process list show the command Sleep.

Which statement is true?

- A. Threads_running counts the sessions queued for a free connection slot, so raising thread_cache_size will clear the errors
- B. Idle sessions, not active queries, are using up the connection limit
- C. The server is CPU-bound, because every connected thread is executing a statement and new logins cannot be scheduled
- D. Threads_connected counts the connections refused since startup, so the value will fall only after a server restart

Answer: B

Question: 4

An online retailer has approved a larger InnoDB buffer pool for a production server as part of a scheduled capacity change. The application must stay connected, so during business hours the DBA runs SET GLOBAL innodb_buffer_pool_size with the new value.

Which statement is true about how this change is carried out?

- A. It takes effect only at the next restart, because innodb_buffer_pool_size is read-only
- B. InnoDB evicts the cached pages, then allocates a larger pool
- C. Only sessions that connect after the change use the larger pool, while sessions that already exist keep using the old one
- D. InnoDB grows the pool online in background chunk-sized steps, and Innodb_buffer_pool_resize_status reports its progress

Answer: D

Question: 5

A payments company runs gtid_mode=ON and enforce_gtid_consistency=ON on source pay-src1. A new replica has been restored from a backup that set gtid_purged correctly. The team wants replication to start without anyone looking up a binary log file or position.

What should you do on the replica?

- A. Run CHANGE REPLICATION SOURCE TO with SOURCE_HOST, the credentials, SOURCE_LOG_FILE='binlog.000001' and SOURCE_LOG_POS=4, then START REPLICA
- B. Run CHANGE REPLICATION SOURCE TO with only SOURCE_HOST and the credentials, as GTID mode positions the replica automatically
- C. Run CHANGE REPLICATION SOURCE TO with SOURCE_HOST, the credentials and SOURCE_AUTO_POSITION=1, then START REPLICA
- D. Set gtid_mode to OFF_PERMISSIVE on the replica, then START REPLICA

Answer: C

Question: 6

Six months ago a help-desk vendor ran ANALYZE TABLE tickets UPDATE HISTOGRAM ON priority, with no other histogram options, on this non-indexed and heavily skewed column, and row estimates for filters on it became accurate. Since then a new triage process has changed the distribution of priority values substantially, and the estimates are now far from the actual row counts. innodb_stats_auto_recalc is enabled.

Which statement is true?

- A. InnoDB refreshes the histogram automatically when innodb_stats_auto_recalc detects enough changed rows, so the estimates will correct themselves
- B. The histogram changes only when it is rebuilt with ANALYZE TABLE ... UPDATE HISTOGRAM
- C. Running ANALYZE TABLE tickets with no histogram clause rebuilds the histogram together with the index statistics
- D. The optimizer uses histograms only for indexed columns, so an index on priority must be added before any estimate can improve

Answer: B

Question: 7

At 14:07 a developer at an online retailer ran DROP TABLE orders by mistake. The last full backup ran at 02:00 and recorded binary log file binlog.000412, position 1540. Binary logging is enabled and every log since the backup is available.

What should you do to recover orders with the least data loss?

- A. Restore the 02:00 backup, then use mysqlbinlog to replay from binlog.000412 position 1540 to the end of the latest log, so that no transaction committed after the backup is lost
- B. Replay binlog.000412 from its first event onto the running server with mysqlbinlog
- C. Restore the 02:00 backup, replay with mysqlbinlog from binlog.000412 position 1540 to just before the DROP TABLE event, then resume replay after that event
- D. Restore the 02:00 backup and start the server, as the redo log in the backup covers later changes

Answer: C

Question: 8

An analytics team adds a second mysqld instance to a Linux host that already runs one, using a separate systemd instance unit and its own option group. The DBA has already given the new instance its own datadir and pid-file.

Which two other settings must also differ from the first instance?
(Choose two.)

- A. The innodb_buffer_pool_size value configured for this new instance's workload
- B. The TCP port the new instance listens on
- C. The basedir holding the MySQL server binaries and support files
- D. The character_set_server setting
- E. The Unix socket file path it creates

Answer: B,E

Question: 9

A SaaS provider's compliance audit requires that every client connection over TCP/IP to its MySQL server be encrypted. The server has more than 300 accounts, new accounts are created weekly by several teams, and the DBA team wants a control that does not depend on each team remembering an extra clause.

The server already has a valid certificate and key configured.
What should you do?

- A. Set require_secure_transport to ON with SET PERSIST, so that the server rejects unencrypted TCP/IP connections for every account, including accounts created later
- B. Set ssl_ca, ssl_cert and ssl_key in the [mysqld] section, then restart
- C. Run ALTER USER ... REQUIRE SSL for each existing account, so that those accounts are refused unless they connect with TLS
- D. Add ssl-mode=REQUIRED to the [client] section of the server's option file

Answer: A

Question: 10

At a telecom operator, the DBA suspects that an index merge chosen for one reporting SELECT on net.call_records (alias c) is slower than using a single index. The production application sends many other statements over the same pooled connections, and the DBA must not change the optimizer's behaviour for any of them.

What should you do to test the alternative plan?

- A. Add index_merge=off to optimizer_switch in the option file and restart the server
- B. Run SET GLOBAL optimizer_switch='index_merge=off' and compare the plan of that reporting SELECT
- C. Run that SELECT with EXPLAIN ANALYZE, which executes both the index merge and the single-index plan and reports the timing of each
- D. Add the optimizer hint /*+ NO_INDEX_MERGE(c) */ to that one SELECT and compare its EXPLAIN output and timing

Answer: D

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>