

IBM

Exam 000-614

DB2 10.1 Advanced DBA for Linux UNIX and Windows

Version: 6.0

[Total Questions: 108]

Topic 1, Volume A**Question No : 1 - (Topic 1)**

Which constraints are used to tell the DB2 Optimizer to consider business rules that are not enforced by the database manager?

- A. Check constraints
- B. Default constraints
- C. Informational constraints
- D. Referential integrity constraints

Answer: C

Question No : 2 - (Topic 1)

Given the following database design:

- A partitioned database was created that uses automatic storage.
- A table space named TS_DIM was created on database partition 0.
- A dimension table named DATE_DIM was created in table space TS_DIM.
- A table space named TS_REPL was created on database partitions 1 thru 4.
- A fact table named SALES was created in table space TS_REPL.

A Explain plan shows that a broadcast join is being used by a query that joins the DATE_DIM dimension table and the SALES fact table. Which action should be performed to improve query performance?

- A. Create a replicated summary table for the SALES table using the table space TS_REPL.
- B. Create a replicated summary table for the DATE_DIM table using the table space TS_REPL.
- C. Create an index on both the DATE_DIM and the SALES table for the column that is used in the collocated join.
- D. Create a replicated summary table for the SALES table using the table space TS_DIM and create an index on the DATE_DIM table.

Answer: B

Question No : 3 - (Topic 1)

Which procedure is used to Explain an SQL statement using information obtained from the package cache or the package cache event monitor?

- A. EXPLAIN_FROM_DATA
- B. EXPLAIN_FROM_ACTIVITY
- C. EXPLAIN_FROM_SECTION
- D. EXPLAIN_FROM_CATALOG

Answer: C

Question No : 4 - (Topic 1)

Which statement about the initiation of parallel I/O is correct?

- A. For list prefetching, each list of data pages is divided into smaller lists according to extent boundaries and assigned to different I/O servers.
- B. For read-ahead prefetching, each list of data pages is divided into smaller lists according to extent boundaries and assigned to different I/O servers.
- C. For sequential prefetching, parallel I/O is initiated when the prefetch size is a multiple of the extent size for a table space. Each prefetch request is divided into smaller requests and assigned to different I/O servers.
- D. For sequential prefetching, parallel I/O is initiated when the extent size is a multiple of the prefetch size for a table space. Each prefetch request is divided into smaller requests and assigned to different I/O servers.

Answer: C

Question No : 5 - (Topic 1)

The following message is frequently seen in the DB2 Diagnostics Log file (db2diag.log):
MESSAGE : ADM5502W The escalation of "28581" locks on table "TEST.TABLE1" to lock intent "X" was successful. Which two database configuration parameter settings should be checked? (Choose two.)

- A. LOCKLIST
- B. MAXAPPLS
- C. MAXLOCKS
- D. MAXAGENTS
- E. LOCKTIMEOUT

Answer: A,C

Question No : 6 - (Topic 1)

Consider the following statement:

```
SELECT lastname, job, workdept FROM employee e, department d
WHERE e.workdept = d.deptno
AND d.deptname = 'SOFTWARE SUPPORT'
ORDER BY lastname
```

Which indexes should be created for optimal query performance?

- A. CREATE INDEX emp_idx ON employee (workdept);
CREATE INDEX dept_idx ON department (deptno);
- B. CREATE INDEX emp_idx ON employee (workdept) CLUSTER;
CREATE INDEX dept_idx ON department (deptname, deptno);
- C. CREATE INDEX emp_idx ON employee (workdept, lastname, job);
CREATE INDEX dept_idx ON department (deptname, deptno);
- D. CREATE INDEX emp_idx ON employee (lastname, job, workdept) CLUSTER;
CREATE INDEX dept_idx ON department (deptname);

Answer: C

Question No : 7 - (Topic 1)

What will occur when a database administrator adds a new automatic storage path to an existing storage group?

- A. New table spaces will start consuming storage on the new path.
- B. Existing table spaces will not consume storage on the new path.
- C. The storage group will be placed in "Rebalance Pending" mode.
- D. DB2 will verify that space is available on the storage path specified.

Answer: B

Question No : 8 - (Topic 1)

Which monitoring tool can be used to obtain stripe set information for database managed storage table spaces?

- A. The MON_GET_CONTAINER table function.
- B. The SNAP_GET_TBSP_PART table function.
- C. The SNAPTbsp_RANGE administrative view.
- D. The TS_GET_CONTAINER administrative view.

Answer: A

Question No : 9 - (Topic 1)

Which statement about value compression is correct?

- A. To use value compression, you must have a license for the DB2 Storage Optimization Feature.
- B. Value compression uses a page-level dictionary-based compression algorithm to compress data based on repetition of values found in a single column.
- C. Value compression involves removing duplicate entries for a value, only storing one copy, and keeping track of the location of any references to the stored copy.
- D. When value compression is used, the byte count of compressed NULLs, zeros, and blanks might be larger than that of the uncompressed version of the same values.

Answer: C

Question No : 10 - (Topic 1)

Which statement about compression dictionaries is FALSE?

- A. Page-level dictionaries are stored in hidden rows within the table.
- B. Table-level dictionaries are stored in hidden rows within the table.
- C. Page-level dictionaries are automatically created or recreated by the database manager as pages are filled.
- D. Table-level dictionaries are automatically created or recreated as by the database manager as pages are filled.

Answer: D

Question No : 11 - (Topic 1)

What advantage do block-based buffer pools offer compared to regular buffer pools?

- A.** They improve the performance of sequential prefetching by reading contiguous pages from disk and writing them into contiguous pages in a buffer pool.
- B.** They improve the performance of random I/O by reading non-contiguous pages from disk and writing them into near-contiguous free space in a buffer pool.
- C.** They improve the performance of vectored reads by retrieving non-contiguous pages from disk, combining them into extent-sized blocks, and writing the blocks into contiguous pages in a buffer pool.
- D.** They improve the performance of scattered writes by retrieving non-contiguous pages from a buffer pool, combining them into extent-sized blocks, and writing the blocks into contiguous pages on disk.

Answer: A

Question No : 12 - (Topic 1)

The following query is executed frequently against the SALES table: `SELECT sales_person, region FROM sales WHERE region = 'Quebec'` Which statement will create an index for optimal query performance?

- A.** `CREATE INDEX sales_idx ON sales (region)`
- B.** `CREATE INDEX sales_idx ON sales (sales_person)`
- C.** `CREATE INDEX sales_idx ON sales (region, sales_person)`
- D.** `CREATE INDEX sales_idx ON sales (region) INCLUDE (sales_person)`

Answer: C

Question No : 13 - (Topic 1)

Which monitoring tool can be used to find out how many data pages have been read from the table space containers for regular, large, and temporary table spaces?

- A.** The `MON_DB_SUMMARY` administrative view.
- B.** The `MON_BP_UTILIZATION` administrative view.
- C.** The `MON_GET_ACTIVITY_DETAILS` table function.
- D.** The `MON_GET_EXTENT_MOVEMENT_STATUS` table function.

Answer: B

Question No : 14 - (Topic 1)

What is the FAILARCHPATH database configuration parameter used for?

- A. To specify a path that DB2 will use to look for archive log files that are needed for a rollforward recovery operation.
- B. To specify a path to which DB2 will attempt to store a copy of archive log files to protect against accidental deletion.
- C. To specify a path to which DB2 will attempt to store a copy of both active and archive log files to protect against disk failure and power outages.
- D. To specify a path to which DB2 will attempt to store archive log files if they cannot be written to the primary or secondary archive destination because of a media problem.

Answer: D

Question No : 15 - (Topic 1)

An auxiliary HADR standby database must use which HADR sync mode?

- A. SYNC
- B. ASYNC
- C. NEARSYNC
- D. SUPERASYNC

Answer: D

Question No : 16 - (Topic 1)

What is the recommended method for removing old backup images and log files for a DB2 10.1 database?

- A. Manually delete old backups and logs.
- B. Utilize the PRUNE LOGFILE PRIOR TO command to remove old backups and logs.
- C. Utilize the PRUNE HISTORY AND DELETE command to remove old backups and logs.
- D. Set the DB2PRUNE_ACTION_LOG configuration parameter to a desired date before

generating a new backup.

Answer: C

Question No : 17 - (Topic 1)

Which diagnostic tool can be used to collect metrics on individual SQL statement executions?

- A. db2dart
- B. A snapshot monitor
- C. A statement event monitor
- D. The MON_GET_UNIT_OF_WORK_DETAILS table function

Answer: C

Question No : 18 - (Topic 1)

Which statement about DB2 Merge Backup is FALSE?

- A. It can be used to merge offline backups to build full offline table space backups.
- B. It can be used to merge SNAPSHOT and TSM tape backup images to create a full database backup image.
- C. It can be used merge multiple offline database backups to build a cumulative single offline database backup.
- D. It can be used to merge backups on a machine that is different from the machine where the backup has been created in stand-alone mode.

Answer: B

Question No : 19 - (Topic 1)

What two types of information can be obtained with the MON_GET_TABLESPACE monitor table function? (Choose two.)

- A. Table space available space
- B. Table space read and write activity