1Z0-117 Q&As

Oracle Database 11g Release 2: SQL Tuning Exam

Pass home 1Z0-117 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/1z0-117.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by home Official Exam Center

Instant Download After Purchase
100% Money Back Guarantee
365 Days Free Update
800,000+ Satisfied Customers

QUESTION 1

Examine the query:

```
SQL> SELECT /*+ RESULT_CACHE */ dept, AVG (sal)
FROM emp
GROUP BY deptno;
```

The RESULT_CACHE_MODE parameter is set to MANUAL for the database.

Which two statements are true about the usage of the result cache?

A. The SQL runtime environment checks whether the query result is cached in the result cache; if the result exists, the optimizer fetches the result from it.

B. The SQL runtime environment does check for the query result in the result cache because the RESULT_CACHE_MODE parameter is set to MANUAL.

C. The SQL runtime environment checks for the query result in the result cache only when the query is executed for the second time.

D. If the query result does not exist in the cache and the query is executed, the result is generated as output, and also sorted in the result cache.

Correct Answer: AD

Note:

*

result_cache_mode: the result cache can be enabled in three ways: via hint, alter session or alter system. Default is MANUAL which means that we need to explicitly request caching via the RESULT_CACHE hint;

*

As its name suggests, the query result cache is used to store the results of SQL queries for re- use in subsequent executions. By caching the results of queries, Oracle can avoid having to repeat the potentially time-consuming and intensive operations that generated the resultset in the first place (for example, sorting/ aggregation, physical I/O, joins etc). The cache results themselves are available across the instance (i.e. for use by sessions other than the one that first executed the query) and are maintained by Oracle in a dedicated area of memory. Unlike our homegrown solutions using associative arrays or global temporary tables, the query result cache is completely transparent to our applications. It is also maintained for consistency automatically, unlike our own caching programs.

*

RESULT_CACHE_MODE specifies when a ResultCache operator is spliced into a query\\'s execution plan.

Values:

/ MANUAL

The ResultCache operator is added only when the query is annotated (that is, hints).

/FORCE

The ResultCache operator is added to the root of all SELECT statements (provided that it is valid to do so).

For the FORCE setting, if the statement contains a NO_RESULT_CACHE hint, then the hint takes precedence over the parameter setting.

QUESTION 2

Your database supports a workload consisting of three categories of SQL statements:

Statements that should execute in less than one minute

Statement that may execute for up to 15 minutes

Statements that may be executed for more than 15 minutes

You set PARALLEL_DEGREE_POLICY to Auto.

You plan to prioritize queued statements by using the Database Resource manager.

Which two are true about parallelism prioritization by a consumer group?

A. PARALLEL_TARGET_PERCENTAGE is used to prioritize a consumer group\\'s use of the overall PARALLEL_SERVER_TARGET.

B. Queuing is done for a consumer group exceeding its percentage, even if the number of busy PX servers in the instance has not reached PARALLEL_SERVERS_TARGET.

C. PARALLEL_TARGET_PERCENTAGE us used to prioritize a consumer group\\'s use of the overall SPAN CLASS = `OracleCode\\'> PARALLEL_MAX_SERVERS.

D. Having separate queues for consumer groups requires the use of management attributes (MGMT_P1, MGMT_P2 etc. . .)

E. Separate queue timeout using PARALLEL_QUEUE_TIMEOUT require the use of management attributes (MGMT_P1, MGMT_P2 etc . . . groups)

Correct Answer: AD

A: Parallel Target Percentage

It is possible for a single consumer group to launch enough parallel statements to use all the available parallel servers. If this happens, when a high-priority parallel statement from a different consumer group is run, no parallel servers are available to allocate to this group. You can avoid such a scenario by limiting the number of parallel servers that can be used by a particular consumer group.

Use the PARALLEL_TARGET_PERCENTAGE directive attribute to specify the maximum percentage of the parallel server pool that a particular consumer group can use. The number of parallel servers used by a particular consumer group is counted as the sum of the parallel servers used by all sessions in that consumer group.

Incorrect:

Not B: PARALLEL_SERVERS_TARGET specifies the number of parallel server processes allowed to run parallel statements before statement queuing will be used. When the parameter PARALLEL_DEGREE_POLICY is set to AUTO,

Oracle will queue SQL statements that require parallel execution, if the necessary parallel server processes are not available. Statement queuing will begin once the number of parallel server processes active on the system is equal to or greater than PARALLEL_SERVER_TARGET. Not C: Would be true if we replace PARALLEL_MAX_SERVERS with PARALLEL_SERVER_TARGET. Not E: The PARALLEL_QUEUE_TIMEOUT directive attribute enables you to specify the maximum time, in seconds, that a parallel statement can wait in the parallel statement queue before it is timed out. The PARALLEL_QUEUE_TIMEOUT attribute can be set for each consumer group. This attribute is applicable even if you do not specify other management attributes (mgmt_p1, mgmt_p2, and so on) in your resource plan. Note:

*

PARALLEL_DEGREE_POLICY AUTO Enables automatic degree of parallelism, statement queuing, and in-memory parallel execution.

*

The PARALLEL_TARGET_PERCENTAGE attribute enables you to specify when parallel statements from a consumer group can be queued. Oracle Database maintains a separate parallel statement queue for each consumer group.

*

PARALLEL_SERVERS_TARGET specifies the number of parallel server processes allowed to run parallel statements before statement queuing will be used. When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle will queue SQL statements that require parallel execution, if the necessary parallel server processes are not available. Statement queuing will begin once the number of parallel server processes active on the system is equal to or greater than PARALLEL_SERVER_TARGET.

By default, PARALLEL_SERVER_TARGET is set lower than the maximum number of parallel server processes allowed on the system (PARALLEL_MAX_SERVERS) to ensure each parallel statement will get all of the parallel server resources required and to prevent overloading the system with parallel server processes.

Note that all serial (non-parallel) statements will execute immediately even if statement queuing has been activated.

* Oracle Database Resource Manager (the Resource Manager) enables you to optimize resource allocation among the many concurrent database sessions.

The elements of the Resource Manager are:

/ Resource consumer group A group of sessions that are grouped together based on resource requirements. The Resource Manager allocates resources to resource consumer groups, not to individual sessions.

/ Resource plan A container for directives that specify how resources are allocated to resource consumer groups. You specify how the database allocates resources by activating a specific resource plan.

/ Resource plan directive Associates a resource consumer group with a particular plan and specifies how resources are to be allocated to that resource consumer group.

QUESTION 3

Which three are benefits of In-Memory Parallel Execution?

A. Reduction in the duplication of block images across multiple buffer caches

B. Reduction in CPU utilization

C. Reduction in the number of blocks accessed

D. Reduction in physical I/O for parallel queries

E. Ability to exploit parallel execution servers on remote instance

Correct Answer: ACD

Note: In-Memory Parallel Execution

When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle Database decides if an object that is accessed using parallel execution would benefit from being cached in the SGA (also called the buffer cache). The decision to cache an object is based on a well-defined set of heuristics including the size of the object and frequency on which it is accessed. In an Oracle RAC environment, Oracle Database maps pieces of the object into each of the buffer caches on the active instances. By creating this mapping, Oracle Database automatically knows which buffer cache to access to find different parts or pieces of the object. Using this information, Oracle Database prevents multiple instances from reading the same information from disk over and over again, thus maximizing the amount of memory that can cache objects. If the size of the object is larger than the size of the buffer cache (single instance) or the size of the buffer cache multiplied by the number of active instances in an Oracle RAC cluster, then the object is read using direct-path reads.

Reference: Oracle Database VLDB and Partitioning Guide 11g, How Parallel Execution Works

QUESTION 4

A database supports three applications: CRM, ERP, and ACC. These applications connect to the database by using three different services: CRM_SRV for the

CRM application, ERP_SRV for the ERP application, and ACC_SRV for the ACC application.

You enable tracing for the ACC_SRV service by issuing the following command:

SQL> EXECUTE DBMS for the ACC_SRV service by issuing the following command:

SQL> EXECUITIVE DBMS_MONITOR. SERV_MOD_ACT_TRACE_ENABLE

(service_name => `ACC_SRV\\', waits => TRUE, binds = > FALSE, instance_name = > `inst1\\');

Which statement is true?

A. All trace information for the service connection to inst1 will be stored in a single trace file.

B. A trace file is not created because the module name is not specified.

C. A single trace file is created for each session that uses the ACC_SRV service.

D. Only those SQL statements that are identified with the ACC_SRV service executed on the inst1 instance are recorded in trace files.

E. All trace information for the ACC_SRV service connected to inst1 is stored in multiple trace files, which can be consolidated by using the tkprof utility.

Correct Answer: C

SERV_MOD_ACT_TRACE_ENABLE

serv_mod_act_trace_enable and serv_mod_act_trace_disable, which enables and disables trace for given

service_name, module and action.

For example for a given service name you can trace all session started from SQL*Plus. Module and action in your own created application can be set using

dbms_application_info set_module and set_action procedures.

serv_mod_act_trace_enable fills sys table wri$_tracing_enabled and view dba_enabled_traces on top of this table as follows:

SQL> exec dbms_monitor.serv_mod_act_trace_enable(service_name=>\\'orcl\\', module_name=>\\'SQL*Plus\\')

PL/SQL procedure successfully completed.

SQL> select * from sys.wri$_tracing_enabled;

TRACE_TYPE PRIMARY_ID QUALIFIER_ID1 QUALIFIER_ID2 INSTANCE_NAME

FLAGS

---------- ---------- ------------- ------------- ------------- ----- 4 orcl SQL*Plus 8

SQL> select * from dba_enabled_traces;

TRACE_TYPE PRIMARY_ID QUALIFIER_ID1 QUALIFIER_ID2 WAITS BINDS

INSTANCE_NAME

-------------- ---------- ------------- ------------- ----- ----- ------------- SERVICE_MODULE orcl SQL*Plus TRUE FALSE

QUESTION 5

Examine the following command:

```
SQL > CREATE VIEW all_employees AS
        (SELECT employee_id, last_name, job_id, commission_pct, department_id FROM employees)
UNION
        (SELECT employee_id, last_name, job_id, commission_pct, department_id FROM contract_workers);
```

You query the view as follows:

```
SELECT last_name
FROM all_employee
WHERE department_id = 50;
```

Which query transformation technique is used by the optimizer in this case?

A. View merging

B. Filter push-down

C. Predicate pushing

D. Predicate move-around

Correct Answer: C

In predicate pushing, the optimizer "pushes" the relevant predicates from the containing query block into the view query block. For views that are not merged, this technique improves the subplan of the unmerged view because the database can use the pushed-in predicates to access indexes or to use as filters.

For example, suppose you create a view that references two employee tables. The view is defined with a compound query that uses the UNION set operator, as follows:

CREATE VIEW all_employees_vw AS

(SELECT employee_id, last_name, job_id, commission_pct, department_id FROM employees)

UNION

(SELECT employee_id, last_name, job_id, commission_pct, department_id FROM contract_workers);

You then query the view as follows:

SELECT last_name

FROM all_employees_vw

WHERE department_id = 50;

Because the view is a compound query, the optimizer cannot merge the view\\'s query into the accessing query block. Instead, the optimizer can transform the

accessing statement by pushing its predicate, the WHERE clause condition department_id=50, into the view\\'s compound query.

The equivalent transformed query is as follows:

SELECT last_name

FROM (SELECT employee_id, last_name, job_id, commission_pct, department_id FROM employees

WHERE department_id=50

UNION

SELECT employee_id, last_name, job_id, commission_pct, department_id FROM contract_workers

WHERE department_id=50);

Reference: Oracle Database Performance Tuning Guide, Predicate Pushing

Latest 1Z0-117 Dumps 1Z0-117 Practice Test 1Z0-117 Exam Questions