



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



## **CERTIFICATION TEST**

- [CertificationTest.net](https://CertificationTest.net) - Cheap & Quality Resources With Best Support

Examine the description of the PROMOTIONS table:

Name	Null?	Type
PROMO_ID	NOT NULL	NUMBER (6)
PROMO_NAME	NOT NULL	VARCHAR2 (30)
PROMO_CATEGORY	NOT NULL	VARCHAR2 (30)
PROMO_COST	NOT NULL	NUMBER (10, 2)

You want to display the unique promotion costs in each promotion category.

Which two queries can be used? (Choose two.)

- A. SELECT DISTINCT promo\_category || ' has ' || promo\_cost AS COSTS FROM promotions ORDER BY 1;
- B. SELECT DISTINCT promo\_cost || ' in ' || DISTINCT promo\_category FROM promotions ORDER BY 1;
- C. SELECT DISTINCT promo\_category, promo\_cost FROM promotions ORDER BY 1;
- D. SELECT promo\_category DISTINCT promo\_cost, FROM promotions ORDER BY 2;
- E. SELECT promo\_cost, promo\_category FROM promotions ORDER BY 1;

**Suggested Answer: AC**

Community vote distribution

AC (100%)

 **kartajin** 2 months ago

**Selected Answer: AC**

A-C is correct

upvoted 1 times

 **pushpaksarode** 3 months ago

**Selected Answer: AC**

A and C are correct options

upvoted 1 times

 **dici79** 5 months, 3 weeks ago

**Selected Answer: AC**

B: Has two distincts in the statement. Not correct.

D: The position of the distinct is not correct


E: it is not returning unique values.

upvoted 2 times

 **5cff9a8** 8 months ago

A and C is the correct answer.

upvoted 1 times

 **Drexan** 8 months, 1 week ago

**Selected Answer: AC**

Correct answers

upvoted 1 times

 **fuck\_india** 1 year, 1 month ago

**Selected Answer: AC**

A and C are correct.

upvoted 1 times

 **alelejaja** 1 year, 1 month ago

**Selected Answer: AC**

Agree that right answers are A and C

upvoted 1 times

 **Stevenciu** 1 year, 3 months ago

Selected Answer: AC

Agree that right answers are A and C  
upvoted 1 times

🗨️ 👤 **Lee\_jong\_suk** 1 year, 6 months ago

Selected Answer: AC

Distinct is used to get distinct set of values for one or more columns mentioned in select statement  
upvoted 2 times

🗨️ 👤 **KizTine** 1 year, 9 months ago

Selected Answer: AC

The correct answers  
upvoted 1 times

🗨️ 👤 **tradersmj1** 2 years ago

Selected Answer: AC

AC are correct  
upvoted 1 times

🗨️ 👤 **Oracle\_373** 2 years, 1 month ago

-- Distinct keyword is used to select the unique values of the column/combination of columns mentioned after the DISTINCT clause.

Thus Option E is incorrect , as this does not use DISTINCT clause and might include duplicates.

-- Distinct should be used once in a SELECT statement and should be used IMMEDIATELY after the SELECT clause. NOT to be used before every column.

Thus Options D and B are incorrect, as the positioning of the DISTINCT clause is incorrect.

-- Option A and C are correct.

Option A: Even though there is a string ' has ' in everyrow returned , the combination of promo\_category || ' has ' || promo\_cost will still be distinct.

Option C: correct use case. ORDER BY clause does not affect the uniqueness in the returned result set.

upvoted 4 times

🗨️ 👤 **alic\_alex** 2 years, 3 months ago

Selected Answer: AC

AC are correct. DB have Distinct in a wrong position in the query and E returns all because of the missing DISTINCT  
upvoted 2 times

🗨️ 👤 **subtype11** 2 years, 5 months ago

Option A is not correct because the query uses the "||" operator to concatenate the promo\_category and promo\_cost columns with a string ' has ' in between. This creates a new column "COSTS" that combines the values from the two columns in a way that is not useful for the task of displaying unique promotion costs in each promotion category. The SELECT DISTINCT statement is used on this combined column, so it will return unique values for the concatenated column, but it does not address the requirement of showing unique promotion costs in each promotion category. Additionally, it uses an ORDER BY 1 which in this case doesn't make sense since the column COSTS doesn't exist.

upvoted 1 times

🗨️ 👤 **J4vi** 2 years, 2 months ago

I tried all the options.The ORDER BY still works in this concatenations and the table shows what its required

upvoted 2 times

🗨️ 👤 **elchiko** 2 years, 7 months ago

distinct must come soon after the select keyword and only one per select statement

upvoted 1 times

🗨️ 👤 **Rik92** 2 years, 8 months ago

AC are correct. DB have Distinct in a wrong position in the query and E returns all because of the missing DISTINCT  
upvoted 2 times

🗨️ 👤 **Artur1991** 2 years, 9 months ago

ac are correct answers

upvoted 2 times

Examine the description of the PRODUCTS table:

Name	Null?	Type
PRODUCT_ID	NOT NULL	NUMBER (2)
PRODUCT_NAME		VARCHAR2 (10)
UNIT_PRICE		NUMBER (3)
SURCHARGE		VARCHAR2 (2)
EXPIRY_DATE		DATE
DELIVERY_DATE		DATE

Which three queries use valid expressions? (Choose three.)

- A. SELECT product\_id, unit\_price, S "Discount", unit\_price + surcharge - discount FROM products;
- B. SELECT product\_id, (unit\_price \* 0.15 / (4.75 + 552.25)) FROM products;
- C. SELECT product\_id, (expiry\_date - delivery\_date) \* 2 FROM products;
- D. SELECT product\_id, unit\_price || 5 "Discount", unit\_price + surcharge - discount FROM products;
- E. SELECT product\_id, expiry\_date \* 2 FROM products;
- F. SELECT product\_id, unit\_price, unit\_price + surcharge FROM products;

**Suggested Answer: BCF**

Community vote distribution

BCF (100%)

**GunduzHuseynli** Highly Voted 2 years, 9 months ago

BCE should be correct answer cause number+varchar in this case impossible  
upvoted 7 times

**extopic01** 2 years, 9 months ago

Have u tried it yourself in livesql.oracle.com?  
Number + varchar is possible due to implicit casting, so BCF is correct.  
upvoted 6 times

**zzsqlcoder** 9 months, 1 week ago

D executes also. Implicit casting applies  
upvoted 1 times

**gioves28** 2 years, 9 months ago

E is wrong inconsistent datatypes: expected NUMBER got DATE  
upvoted 2 times

**dici79** Most Recent 5 months, 3 weeks ago

**Selected Answer: BCF**

A: Discount field doesn't exist in the table.

B: OK

C: OK

D: Discount field doesn't exist in the table.

E: Date \* Number is not a valid operation

F: Depending on the values, it may or may not be OK. if the entered values are numeric, this statement will work. Because the question requires 3 answers, I pass this option as correct.

upvoted 1 times

**Drexan** 8 months, 1 week ago

**Selected Answer: BCF**

F works (although is VARCHAR2) if its content is numeric,

upvoted 1 times

**alelejaja** 1 year, 1 month ago

**Selected Answer: BCF**

B and C clearly correct.

F is tricky. It can be correct only if the column Surcharge contains numeric characters defined as varchar2

upvoted 2 times

🗳️ 👤 **tom2992** 1 year, 2 months ago

**Selected Answer: BCF**

A and D wrong cause they have invalid identifier "Discount".

E wrong cause date can't multiple with number.

upvoted 1 times

🗳️ 👤 **Lee\_jong\_suk** 1 year, 6 months ago

BCF is correct answer

Please note:

date+date is not allowed

upvoted 2 times

🗳️ 👤 **TMerlo** 1 year, 9 months ago

**Selected Answer: BCF**

It's not a right decision plus a NUMBER with VARCHAR2 fields.

upvoted 1 times

🗳️ 👤 **jm9999** 1 year, 9 months ago

Seems like a sketchy question, have to assume surcharge contains only numerical characters in the string for F

upvoted 3 times

🗳️ 👤 **fantastic\_side\_eye** 2 years ago

BCF

The three queries that use valid expressions are:

B. SELECT product\_id, (unit\_price \* 0.15 / (4.75 + 552.25)) FROM products;

C. SELECT product\_id, (expiry\_date - delivery\_date) \* 2 FROM products;

F. SELECT product\_id, unit\_price, unit\_price + surcharge FROM products;

Explanation:

A. This query is invalid because the alias "S" is not defined anywhere in the query.

B. This query is valid. It calculates a value based on the unit price of each product using a mathematical expression.

C. This query is valid. It calculates the difference between the expiry date and delivery date for each product and multiplies it by 2.

D. This query is invalid because the alias "discount" is not defined anywhere in the query.

E. This query is invalid because you cannot multiply a date value by a number.

F. This query is valid. It selects the product ID, unit price, and the sum of the unit price and surcharge for each product.

upvoted 3 times

🗳️ 👤 **zouve** 2 years ago

select '2'+3 from dual;

it produces the SUM correctly so F is correct

upvoted 2 times

🗳️ 👤 **Nelita** 2 years, 1 month ago

B-C-F IS CORRECT IF the varchar-column has stored number.

upvoted 2 times

🗳️ 👤 **Indhumathishan** 2 years, 1 month ago

why is D incorrect?

upvoted 3 times

🗨️ 👤 **VoidWanderer** 10 months, 1 week ago

they are subtracting discount but there's no such column in table.

upvoted 1 times

🗨️ 👤 **holdfaststrongly** 2 years, 9 months ago

Is "-discount" in AD supposed to be an alias for the column?

upvoted 1 times

🗨️ 👤 **MZEnuovazelandia** 2 years, 9 months ago

IT is correct BCF because in the Datatype Varchar2 there may be a number. Oracle makes it an implicit cast and uses it as a number

upvoted 4 times

🗨️ 👤 **gabotka** 2 years, 9 months ago

you cant multiply date, BCF correct if varchar2 is number

upvoted 4 times

🗨️ 👤 **Efstathia** 2 years, 9 months ago

Just to say the same things in other words.... If surcharge column contains numeric values (even as varchar2 type) can be implicitly converted

upvoted 3 times

🗨️ 👤 **CyberP** 2 years, 1 month ago

If you cant multiply date, then why you chose c ? since c the date is also multiply ?

upvoted 1 times

🗨️ 👤 **MySekQL** 1 year, 9 months ago

I'm guessing because if you subtract two dates that are relatively close, you will get a number not a date. And that CAN be multiplied.

upvoted 1 times

🗨️ 👤 **Averagenoob** 2 years, 10 months ago

bce should be


upvoted 2 times

What is true about non-equijoin statement performance? (Choose two.)

- A. The BETWEEN condition always performs less well than using the >= and <= conditions.
- B. The BETWEEN condition always performs better than using the >= and <= conditions.
- C. The Oracle join syntax performs better than the SQL:1999 compliant ANSI join syntax.
- D. Table aliases can improve performance.
- E. The join syntax used makes no difference to performance.

**Suggested Answer:** DE

Community vote distribution



**oracledev** Highly Voted 2 years, 4 months ago

**Selected Answer:** CD

D. Table aliases can improve performance. Table aliases can improve performance by reducing the amount of parsing needed to execute the query.

C. The Oracle join syntax performs better than the SQL:1999 compliant ANSI join syntax in Oracle 18c. Oracle's join syntax can result in better performance in certain scenarios as it allows the optimizer to understand the join order and access paths more effectively.

Therefore, options C and D are correct. Options A, B, and E are incorrect.

upvoted 6 times

**RogerGadelha** Most Recent 1 month, 3 weeks ago

**Selected Answer:** E

A and B are false, no discussion.

C. Oracle's query optimizer internally rewrites both syntaxes into the same execution plan. The decision on how to execute the join—nested loop, hash join, merge join, etc.—is based on statistics, indexes, and optimizer settings, not on the join syntax itself.

D. Oracle's SQL parser and optimizer are designed to handle fully qualified table names just as efficiently as aliases. The use of aliases affects only the textual complexity of the query, not the underlying parsing or optimization time in any meaningful or measurable way.

upvoted 1 times

**dan\_bj** 2 months, 1 week ago

**Selected Answer:** DE

Oracle's official stance: Oracle documentation explicitly states that the SQL:1999-compliant join syntax does not offer any performance benefits over the Oracle-proprietary join syntax that existed in prior releases

upvoted 1 times

**highwater** 3 months ago

**Selected Answer:** DE

D, E with no doubt

upvoted 1 times

**Thameur01** 5 months ago

**Selected Answer:** DE

D. Table aliases can improve performance. Most Voted

E. The join syntax used makes no difference to performance.

upvoted 1 times

**Thameur01** 5 months ago

**Selected Answer:** DE

D. Table aliases can improve performance. Most Voted

E. The join syntax used makes no difference to performance.

upvoted 1 times

**dici79** 5 months, 3 weeks ago

**Selected Answer:** DE

A and B states ALWAYS - which is not always the case.  
C: The join syntax doesn't make much difference with performance.  
D: That's correct; it helps the optimizer  
E: this makes sense to me - The joint syntax doesn't make a difference.  
upvoted 1 times

🗨️ **Drexan** 8 months, 1 week ago

**Selected Answer: CD**

According to some papers and IA, these must be the better options...  
upvoted 1 times

🗨️ **hidro** 8 months, 2 weeks ago

options D and E are correct  
Options C are incorrect because the SQL:1999-compliant join syntax does not offer any performance benefits over the Oracle-proprietary join syntax that existed in the earlier releases.  
upvoted 1 times

🗨️ **VoidWanderer** 10 months, 3 weeks ago

i find D and E a little bit contradictory, do aliases come under join syntax?  
upvoted 1 times

🗨️ **alelejaja** 1 year, 1 month ago

**Selected Answer: DE**

D. Table aliases can improve performance.  
E. the difference related to performance that can be experimented between the two join syntax is minimum since Oracle optimizes internally both manners in a similar way  
upvoted 1 times

🗨️ **Misi\_Oracle** 1 year, 1 month ago

C and D is true.  
C is true because oracle Join does perform better.  
D is true. Table aliases improve the I/O. Refer:  
[https://docs.oracle.com/middleware/1221/bip/BIPDM/best\\_practices.htm](https://docs.oracle.com/middleware/1221/bip/BIPDM/best_practices.htm)  
E is false. It does affect performance. Refer:  
[https://docs.oracle.com/cd/F49540\\_01/DOC/server.815/a67781/c20c\\_joi.htm](https://docs.oracle.com/cd/F49540_01/DOC/server.815/a67781/c20c_joi.htm)  
upvoted 1 times

🗨️ **Lee\_jong\_suk** 1 year, 6 months ago

**Selected Answer: DE**

Table alias makes it easy for the parser to look up the columns in their respective tables.  
"Join" or "Where T1.id=T2.id" conditions have no difference in performance, though join is recommended to be used as the new syntax.  
upvoted 1 times

🗨️ **KizTine** 1 year, 9 months ago

**Selected Answer: CE**

I vote for answers C and E. To be honest, I don't know how good they are, so vague are their contents... I just know that answers A, B and C are wrong.  
So I vote by elimination.  
upvoted 1 times

🗨️ **captainkul** 1 year, 9 months ago

C and E correct.  
upvoted 1 times

🗨️ **dilshod** 1 year, 12 months ago

**Selected Answer: DE**

D E are correct  
upvoted 3 times

🗨️ **musafir** 2 years ago

**Selected Answer: DE**

C D E are all correct it seems  
<https://www.examttopics.com/discussions/oracle/view/10257-exam-1z0-071-topic-1-question-253-discussion/>  
upvoted 1 times



Which two are true? (Choose two.)

- A. ADD\_MONTHS adds a number of calendar months to a date.
- B. CEIL requires an argument which is a numeric data type.
- C. CEIL returns the largest integer less than or equal to a specified number.
- D. LAST\_DAY returns the date of the last day of the current month only.
- E. LAST\_DAY returns the date of the last day of the month for the date argument passed to the function.
- F. LAST\_DAY returns the date of the last day of the previous month only.


**Suggested Answer:** AC

Reference:

<https://docs.oracle.com/database/121/SQLRF/functions025.htm>

[https://docs.oracle.com/cd/B28359\\_01/server.111/b28286/functions004.htm](https://docs.oracle.com/cd/B28359_01/server.111/b28286/functions004.htm)

Community vote distribution



**Peter\_Chen** Highly Voted 2 years, 10 months ago

maby AE is Right...

C. CEIL returns the largest integer less than or equal to a specified number.

Can change : CEIL returns the minmum integer more than or equal to a specified number

upvoted 13 times

**Eren1998** Most Recent 4 months ago

Selected Answer: AE

B is wrong because [select ceil('12,4') from dual;] . '12.4' is char type, Oracle use implicit conversion to make the argument is valid if they can be.

upvoted 1 times

**Drexan** 8 months, 1 week ago

Selected Answer: AE

The only possible responses

upvoted 1 times

**fantastic\_side\_eye** 9 months ago

The two true statements are:

AE

Explanation:

B. CEIL is an Oracle SQL function that requires an argument which is a numeric data type. This statement is false, as CEIL can also take input arguments of type DATE or TIMESTAMP.

C. CEIL is an Oracle SQL function that returns the smallest integer greater than or equal to a specified number. This statement is false, as CEIL actually returns the largest integer less than or equal to a specified number.

D. LAST\_DAY is an Oracle SQL function that returns the date of the last day of the current month only. This statement is false, as LAST\_DAY can be used with any date argument to return the last day of that month.

E. This statement is true. LAST\_DAY is an Oracle SQL function that can be used with a date argument to return the date of the last day of the month for that date. For example, LAST\_DAY('2023-06-26') would return '2023-06-30'.

F. This statement is false. LAST\_DAY can be used with any date argument to return the last day of that month, not just the previous month.

upvoted 3 times

**highwater** 3 months ago

CEIL works only with numeric data types, it also works with implicit conversions  
for instance CEIL('123.45').

upvoted 1 times

🗳️ 👤 **KizTine** 1 year, 9 months ago

Are you sure that CEIL function can also take input arguments of type DATE or TIMESTAMP ?

I think it can take nonnumeric data type (that can be implicitly converted to a numeric data type) as an argument.

<https://www.w3resource.com/oracle/oracle-numeric-functions/oracle-ceil-function.php>

upvoted 2 times

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

AE is correct since ceil returns minimum integer greater than or equal to specified value.

upvoted 1 times

🗳️ 👤 **durvaljr20** 1 year ago

**Selected Answer: AE**

<https://docs.oracle.com/database/121/SQLRF/functions025.htm>

"CEIL returns the smallest integer that is greater than or equal to n. "

upvoted 1 times

🗳️ 👤 **sitttttt** 1 year ago

A and E is the correct answer. C is incorrect because ceil always return smallest integer greater than or equal to values.

upvoted 1 times

🗳️ 👤 **sitttttt** 1 year ago

B also so A,B and E

upvoted 1 times

🗳️ 👤 **ac77458** 1 year, 2 months ago

**Selected Answer: AB**

Based on Oracle's definitions, even though they do not make sense. Since when common sense is common practice? Anyhow for test purposes. According to that information I believe the correct answers are ABE. Although you can only choose 2.

Oracle's Definition of CEIL: CEIL returns the smallest integer that is greater than or equal to n. ( <https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/CEIL.html#GUID-6DCC9AFB-9B80-4C27-AF63-5AA3B1E43660> )

C is incorrect as it says: CEIL returns the largest integer less than or equal to a specific number. This is the definition of FLOOR ( <https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/FLOOR.html#GUID-67F61AC7-C097-4397-A122-213157BF584F> )

E: LAST\_DAY ( [https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/LAST\\_DAY.html#GUID-296C7C02-7FB9-4AAC-8927-6A79320CE0C6](https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/LAST_DAY.html#GUID-296C7C02-7FB9-4AAC-8927-6A79320CE0C6) )

upvoted 1 times

🗳️ 👤 **tom2992** 1 year, 2 months ago

No. B isn't correct cause CEIL can use with date and interval too.

upvoted 1 times

🗳️ 👤 **ac77458** 1 year, 2 months ago

B also could an answer as if you execute select ceil('abc') from dual; you get ORA-01722: invalid number.

Making this even more confusing.

upvoted 2 times

🗳️ 👤 **Lee\_jong\_suk** 1 year, 6 months ago

**Selected Answer: AE**

AE is the correct answer.

SELECT LAST\_DAY(TO\_DATE('09-SEP-2023')) FROM DUAL;

//op- 30-09-2023

SELECT CEIL('100'.2) FROM DUAL;

//op 101

select add\_months(TO\_DATE('25-DEC-2023'),3) from dual;

//op 25-03-2024

upvoted 1 times



🗳️ 👤 **Roneros** 1 year, 9 months ago

A and E are the correct answers.

A - ADD\_MONTHS returns the date date plus integer months. ([https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/ADD\\_MONTHS.html#GUID-B8C74443-DF32-4B7C-857F-28D557381543](https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/ADD_MONTHS.html#GUID-B8C74443-DF32-4B7C-857F-28D557381543))

E - LAST\_DAY returns the date of the last day of the month that contains date. The return type is always DATE, regardless of the datatype of date. ([https://docs.oracle.com/cd/B19306\\_01/server.102/b14200/functions072.htm](https://docs.oracle.com/cd/B19306_01/server.102/b14200/functions072.htm))

upvoted 1 times

  **Guru\_C** 1 year, 11 months ago

C. CEIL returns the largest integer less than or equal to a specified number.



How does Option:C is correct. CEIL always return greater than or equal to specified number.

upvoted 3 times

  **Fredderik91** 2 years, 1 month ago



AE is correct, C is false CEIL return SMALL NOT LARGEST <https://docs.oracle.com/en/database/oracle/oracle-database/21/sqlrf/CEIL.html#GUID-6DCC9AFB-9B80-4C27-AF63-5AA3B1E43660>

upvoted 1 times

  **jokri** 2 years, 1 month ago

por sure AE are correct. C cannot be correct by "CEIL returns the largest integer...". Correct answer should be "CEIL returns the smallest integer..."

upvoted 1 times

  **PinoZ** 2 years, 3 months ago

**Selected Answer: AE**

C is the definition of FLOOR()!

upvoted 3 times

  **Augustoaq** 2 years, 4 months ago

I think that these are correct: A, B, E.

C is wrong because CEIL returns the smallest integer that is greater than or equal to the specified number

upvoted 1 times

  **Orxan\_H** 2 years, 2 months ago

Why B incorrect:

This function takes as an argument any numeric data type or any nonnumeric data type that can be implicitly converted to a numeric data type.

The function returns the same data type as the numeric data type of the argument.

<https://docs.oracle.com/database/121/SQLRF/functions025.htm#SQLRF00614>

upvoted 4 times

  **jfc1** 2 years, 5 months ago

**Selected Answer: AE**

C is wrong because

CEIL returns the smallest integer greater than or equal to a specified number.

<https://docs.oracle.com/database/121/SQLRF/functions025.htm#SQLRF00614>

E is good

Because

select last\_day(('01-FEB-2023')) as "xxxgg" from dual;

Answer : 2/28/2023

Finally the answer is AE

upvoted 1 times

Which three statements are true about Oracle synonyms? (Choose three.)

- A. A synonym cannot be created for a PL/SQL package.
- B. A synonym can be available to all users.
- C. A SEQUENCE can have a synonym.
- D. Any user can drop a PUBLIC synonym.
- E. A synonym created by one user can refer to an object belonging to another user.

**Suggested Answer:** BCE

Community vote distribution

BCE (100%)

🗳️ 👤 **Drexan** 8 months, 1 week ago

**Selected Answer:** BCE

Almost all objects can have synonyms. And drop implies privileges  
upvoted 1 times

🗳️ 👤 **dilshod** 2 years ago

**Selected Answer:** BCE

bce, ecb  
upvoted 1 times

🗳️ 👤 **alic\_alex** 2 years, 3 months ago

**Selected Answer:** BCE

BCE is the correct answer!  
upvoted 1 times

🗳️ 👤 **thai\_27** 2 years, 9 months ago

BCE is the correct answer!  
upvoted 1 times

🗳️ 👤 **holdfaststrongly** 2 years, 9 months ago

D) Any user can drop a PUBLIC synonym. False. Must have system privilege to create or drop a PUBLIC SYNONYM.  
upvoted 2 times

🗳️ 👤 **Artur1991** 2 years, 9 months ago

bce are the correct answers  
upvoted 1 times

Which two are true? (Choose two.)

- A. CONCAT joins two character strings together.
- B. CONCAT joins two or more character strings together.
- C. FLOOR returns the largest positive integer less than or equal to a specified number.
- D. INSTR finds the offset within a character string, starting from position 0.
- E. INSTR finds the offset within a string of a single character only.
- F. FLOOR returns the largest integer less than or equal to a specified number.

**Suggested Answer:** AF

Community vote distribution

AF (100%)

 **highwater** 3 months ago

**Selected Answer:** AF

- A. CONCAT joins two character strings together
  - F. FLOOR returns the largest integer less than or equal to a specified number
- upvoted 2 times

 **LuisLikes** 4 months, 2 weeks ago

**Selected Answer:** BF

B is correct, CONCAT joins two or more arguments.

--Example query

```
select concat('this ', 'is ', 'a ', 'test') concat from dual;
```

Output: "this is a test"

F is also correct, since it accepts negative numbers as well.

--Example query

```
select floor(-5.2526) from dual;
```

Output: -6 (-6 is smaller than -5, it also caught me offguard the first time).

upvoted 1 times


 **highwater** 3 months ago

B is not true

if you want to use CONCAT() with many arguments then it looks like the example below

```
select concat('this', concat(' is', concat(' a', ' test')))) from dual;
```

upvoted 1 times

 **15eeda4** 5 months, 3 weeks ago

**Selected Answer:** BF

Starting from Oracle 23c, CONCAT can take multiple arguments:

<https://docs.oracle.com/en/database/oracle/oracle-database/21/sqlrf/CONCAT.html>

<https://docs.oracle.com/en/database/oracle/oracle-database/23/sqlrf/CONCAT.html>

upvoted 1 times

 **5cff9a8** 8 months ago

I thought B was correct, whereas it depends on the database. Oracle defers from Microsoft.

upvoted 1 times

 **Drexan** 8 months, 1 week ago

**Selected Answer:** AF

floor(12.3) returns 12 and ceil(12.3) returns 13

upvoted 1 times

🗨️ **hmatinnn** 1 year, 7 months ago

How F can be correct?! Okay, in oracle documents it is mentioned like : FLOOR returns the largest integer equal to or less than n.

But why it returns less than inputted value?? (select floor(15,7) from dual => output is 15)

That statement is proper for CEIL function for my point of view.

Can anyone help me to overcome this simply seem but complicated question?

upvoted 1 times

🗨️ **Lee\_jong\_suk** 1 year, 6 months ago

Floor takes the minimum integer value while CEIL takes the largest integer value.

```
SELECT FLOOR('100.2') FROM DUAL; //op: 100
```

```
SELECT CEIL('100.2242') FROM DUAL; //op: 101
```

upvoted 2 times

🗨️ **Guru\_C** 1 year, 11 months ago

How F is correct.

In case argument is passed in a negative. eg -1.4. The output is -2.

But -1 is greater than -2. And, its contrast with option F.

F. FLOOR returns the largest integer less than or equal to a specified number.

```
select floor(-1.4) from dual;
```

o/p

---2

upvoted 3 times

🗨️ **Proctored\_Expert** 2 years, 6 months ago

**Selected Answer: AF**

A. CONCAT joins two character strings together.

F. FLOOR returns the largest integer less than or equal to a specified number.

upvoted 2 times

🗨️ **MahdiHamdii** 2 years, 6 months ago

**Selected Answer: AF**

A-F Correct

upvoted 1 times

🗨️ **boddoju** 2 years, 7 months ago

Why option-b is not correct? can any one explain

upvoted 1 times

🗨️ **J4vi** 2 years, 4 months ago

Try to concat 3 different strings and see... it's possible in Microsoft SQL, not Oracle

upvoted 2 times

🗨️ **Rik92** 2 years, 1 month ago

CONCAT as a function in Oracle can only do two strings, if you want more, use ||

The function and operator are not the same thing!

upvoted 1 times

🗨️ **pawani** 2 years, 9 months ago

how is F correct?

upvoted 1 times

Examine these SQL statements which execute successfully:

```
CREATE TABLE emp
(emp_no NUMBER(2) CONSTRAINT emp_emp_no_pk PRIMARY KEY,
ename VARCHAR2(15),
salary NUMBER(8,2),
mgr_no NUMBER(2));
```

```
ALTER TABLE emp ADD CONSTRAINT emp_mgr_fk
FOREIGN KEY (mgr_no)
REFERENCES emp(emp_no)
ON DELETE SET NULL;
```

```
ALTER TABLE emp
DISABLE CONSTRAINT emp_emp_no_pk
CASCADE;
```

```
ALTER TABLE emp
ENABLE CONSTRAINT emp_emp_no_pk;
```

Which two statements are true after execution? (Choose two.)

- A. The primary key constraint will be enabled and IMMEDIATE.
- B. The foreign key constraint will be enabled and DEFERRED.
- C. The primary key constraint will be enabled and DEFERRED.
- D. The foreign key constraint will be disabled.
- E. The foreign key constraint will be enabled and IMMEDIATE.

**Suggested Answer: AD**

Community vote distribution

AD (100%)

826cb95 2 months, 3 weeks ago

**Selected Answer: AE**

Its A & E not D.

why? read below

D. The foreign key constraint will be disabled.

This is incorrect because the question states that the SQL statements execute successfully, so the foreign key constraint would typically be enabled, not disabled.

upvoted 1 times

5cff9a8 8 months ago

Never knew what DEFERRED was. One keeps, learning every day.

upvoted 2 times

Drexan 8 months, 1 week ago

**Selected Answer: AD**

DEFERRED: allows to make insertions and updates without immediate validation, only when you commit

upvoted 1 times

Lee\_jong\_suk 1 year, 6 months ago

**Selected Answer: AD**

AD is the correct answer

upvoted 2 times

Andrew91 1 year, 11 months ago

What does it mean by IMMEDIATE tho?

upvoted 1 times

🗨️ 👤 **Lee\_jong\_suk** 1 year, 6 months ago

IMMEDIATE means the check will be done as soon as you make change in your session.

But DEFERRED means the check will be done only when you do commit.

upvoted 3 times

🗨️ 👤 **twing** 2 years, 3 months ago

is AD

PK enable again

upvoted 1 times

🗨️ 👤 **Chandra1104** 2 years, 3 months ago

AD

Cascade constraint will disable to primay key as well as foreign key.

upvoted 2 times

🗨️ 👤 **Rik92** 2 years, 8 months ago

AD. Due to the PK being disabled, the FK is disabled. PK afterwards is enabled again

upvoted 3 times



Examine this SQL statement:

```
UPDATE orders o
  SET customer_name =
    (SELECT cust_last_name
     FROM customers
     WHERE customer_id = o.customer_id);
```

Which two are true? (Choose two.)




- A. All existing rows in the ORDERS table are updated.
- B. The subquery is executed before the UPDATE statement is executed.
- C. The subquery is not a correlated subquery.
- D. The subquery is executed for every updated row in the ORDERS table.
- E. The UPDATE statement executes successfully even if the subquery selects multiple rows.

**Suggested Answer: AB**

Community vote distribution

AD (75%)

AE (25%)



  **meddheddy**  2 years, 9 months ago

A and D tested on Oracle 12c1  
upvoted 14 times

  **highwater**  3 months ago



**Selected Answer: AD**

Correct answers are A,D  
upvoted 1 times

  **Thameur01** 5 months, 1 week ago

**Selected Answer: AD**

Correct Answers:  
A. All existing rows in the ORDERS table are updated.  
D. The subquery is executed for every updated row in the ORDERS table.  
upvoted 2 times

  **Drexan** 8 months, 1 week ago

**Selected Answer: AD**

No WHEN in update: all rows will be updated. Correlated subquery, will execute for every row in the update  
upvoted 1 times

  **Ramileu** 12 months ago

SHOULD BE b AND d  
upvoted 1 times

  **Fredderik91** 1 year, 1 month ago

**Selected Answer: AD**

B (false), depending EXPLAIN PLAN. no always is executed first.  
C (false), ".o" is correlative  
E (false) is possible multiple rows, an error more than row.  
A and D is correct  
upvoted 2 times

  **hmatinnn** 1 year, 5 months ago

**Selected Answer: AD**

SHOULD BE a AND d  
upvoted 1 times

🗨️ 👤 **Lee\_jong\_suk** 1 year, 6 months ago

AD option is right  
upvoted 1 times

🗨️ 👤 **jm9999** 1 year, 9 months ago

AD; cannot be B, correlated subqueries are not executed before the parent (outer) query. Cannot be E, cannot set customer\_name to multiple values for each row  
upvoted 1 times

🗨️ 👤 **HassanMkhlalati** 1 year, 9 months ago

**Selected Answer: AD**

A is correct  
D :It will retrieve the matching value for each row from the table orders with the table customers  
upvoted 1 times

🗨️ 👤 **Fredderik91** 2 years, 1 month ago

BD,  
A is incorrect only matching rows are updated, not all  
C is invalid, statement correct structure  
E is incorrect, in case multiple rows, error multiple rows in subquery  
upvoted 4 times

🗨️ 👤 **HassanMkhlalati** 1 year, 9 months ago

Incorrect  
upvoted 1 times

🗨️ 👤 **GEN0123** 2 years ago

definitely agree. Oracle will only update the rows that are matched  
upvoted 1 times

🗨️ 👤 **TofOra** 1 year, 4 months ago

No, it will update unmatched rows with NULL  
upvoted 2 times

🗨️ 👤 **hmatinnn** 1 year, 6 months ago

Actually i though in that way. But most likely customer table is PK. Orders table is FK. So it means the data that in the orders table also have in the customers table. And all rows will be updated.  
upvoted 1 times

🗨️ 👤 **NB196** 2 years, 1 month ago

Can someone explain why A is correct?  
upvoted 1 times

🗨️ 👤 **raduM** 2 years ago

because it updates every row. If it does not have a match it will put the value null  
upvoted 4 times

🗨️ 👤 **hmatinnn** 1 year, 6 months ago

most likely customer table is PK. Orders table is FK. So it means the data that in the orders table also have in the customers table. And all rows will be updated.  
upvoted 1 times

🗨️ 👤 **CyberP** 2 years, 1 month ago

I think B is correct, subquery will execute for each row its right but also update will update for each raw after fulfilling the condition of subquery :)  
Then how update will update each raw ? update will check its resources "subquery" is condition met ? yes ? then update after execute subquery and fully checking it.  
upvoted 1 times

🗨️ 👤 **TofOra** 1 year, 4 months ago

The STATMENT update will execute first, then the action of updating will execute after the subquery. explain plan for  
update ordersQ8 o  
set customer\_name =  
(select cust\_last\_name  
from customersQ8  
where customer\_id = o.customer\_id);

```
SELECT plan_table_output  
FROM TABLE(DBMS_XPLAN.DISPLAY('plan_table',null,'basic +note'));
```

upvoted 2 times

🗲️ 👤 **Rik92** 2 years, 1 month ago

Selected Answer: AD

AD are correct

upvoted 2 times

🗲️ 👤 **kuff** 2 years, 1 month ago

Selected Answer: AD

AD is the answer

upvoted 1 times

🗲️ 👤 **kaplanek** 2 years, 2 months ago

Selected Answer: AD

A nd D are correct

upvoted 2 times

🗲️ 👤 **alic\_alex** 2 years, 3 months ago

Selected Answer: AD

E - wrong because will cause an error "ORA-01427: Subquery returns more than one row"

upvoted 3 times

Which two statements are true about TRUNCATE and DELETE? (Choose two.)

- A. DELETE can use a WHERE clause to determine which row(s) should be removed.
- B. TRUNCATE can use a WHERE clause to determine which row(s) should be removed.
- C. TRUNCATE leaves any indexes on the table in an UNUSABLE state.
- D. The result of a TRUNCATE can be undone by issuing a ROLLBACK.
- E. The result of a DELETE can be undone by issuing a ROLLBACK.

**Suggested Answer: AE**

Community vote distribution

AE (100%)

  **devml**  2 years, 5 months ago


- A. DELETE can use a WHERE clause to determine which row(s) should be removed. ---> Is correct
  - B. TRUNCATE can use a WHERE clause to determine which row(s) should be removed. ---> Its wrong bcoz truncate will delete every rows in a table except the table definitions.
  - C. TRUNCATE leaves any indexes on the table in an UNUSABLE state. ----> its wrong truncate will remove everything except table definition or skeleton of the table
  - D. The result of a TRUNCATE can be undone by issuing a ROLLBACK.-----> Its wrong bcoz TRUNCATE is a ddl statement whereas it can be done for DELETE since its dml statement
  - E. The result of a DELETE can be undone by issuing a ROLLBACK. ----> Its correct bcoz DELETE is DML statement and each delete of row are recorded in transaction log or the redo logs.
- upvoted 5 times

  **pushpaksarode**  3 months ago

**Selected Answer: AE**

We can use WHERE clause in DELETE, We can ROLLBACK DELETE operation.


upvoted 1 times

  **Drexan** 8 months, 1 week ago

**Selected Answer: AE**

Truncate is forever...

upvoted 1 times

  **alelejaja** 1 year, 1 month ago

**Selected Answer: AE**

A, E correct answers

upvoted 1 times

  **Chandra1104** 2 years, 3 months ago

AE  
is the correct ans.  
With truncate we can not user where clause (B)  
with truncate even if we rollback data will not be rollback. (D)

upvoted 2 times

  **heoj10272** 2 years, 8 months ago

AE correct answers

upvoted 2 times

  **Artur1991** 2 years, 9 months ago

AE are the correct answers

upvoted 2 times

The STORES table has a column START\_DATE of data type DATE, containing the date the row was inserted.  
 You only want to display details of rows where START\_DATE is within the last 25 months.  
 Which WHERE clause can be used?

- A. WHERE TO\_NUMBER(start\_date - SYSDATE) <= 25
- B. WHERE MONTHS\_BETWEEN(start\_date, SYSDATE) <= 25
- C. WHERE MONTHS\_BETWEEN(SYSDATE, start\_date) <= 25
- D. WHERE ADD\_MONTHS(start\_date, 25) <= SYSDATE

**Suggested Answer: C**

Community vote distribution

C (100%)

🗳️ 👤 **Drexan** 8 months, 1 week ago

**Selected Answer: C**

months\_between.. The interval is important, and if I add 25 months to a date 3 years ago, it will not be in the desired interval... disregard my previous comment  
 upvoted 1 times

🗳️ 👤 **Drexan** 8 months, 1 week ago

**Selected Answer: C**

ADD\_Months  
 upvoted 1 times

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

c is the answer using ADD\_MONTHS(date, n)  
 upvoted 1 times

🗳️ 👤 **guretto** 1 year, 8 months ago

START\_DATE stores "the date the row was inserted". Therefore cannot be > SYSDATE (that is "now") and C is the right one  
 upvoted 1 times

🗳️ 👤 **PuErh** 2 years, 3 months ago

I can't get why the B is incorrect.  
 A negative value is still less than 25.  
 I think I have to check it in DB.  
 upvoted 1 times

🗳️ 👤 **TofOra** 1 year, 4 months ago

"A negative value is still less than 25." Yes, but -30 (meaning 30 months ago) would be less than 25.  
 upvoted 1 times

🗳️ 👤 **Rik92** 2 years, 1 month ago

The statement in B would indeed result in -25 and thus be correct. But C is the 'cleaner' answer here and thus the one you should pick.  
 upvoted 1 times

🗳️ 👤 **Rik92** 2 years ago

The question has been edited and this is no longer an option!  
 upvoted 1 times

🗳️ 👤 **dexdinh91** 2 years, 5 months ago

**Selected Answer: C**

C is correct  
 upvoted 2 times

🗳️ 👤 **LAT12022** 2 years, 8 months ago

D is false, only put an example between '01/01/200' and '01/02/2002' 25 months but isn't in the last 25 months  
 upvoted 1 times

🗨️ 👤 **Rik92** 2 years, 8 months ago

C is true. MONTHS\_BETWEEN(date 1, date 2) Oracle gives a positive value when date 1 > date 2 and a negative when date 1 < date 2

[https://www.techonthenet.com/oracle/functions/months\\_between.php](https://www.techonthenet.com/oracle/functions/months_between.php)

upvoted 2 times

🗨️ 👤 **meddheddy** 2 years, 9 months ago

why D is not true ?

upvoted 1 times

🗨️ 👤 **kakamusili** 2 years, 8 months ago

It would work as well,

upvoted 1 times

🗨️ 👤 **Peter\_Chen** 2 years, 8 months ago

Because.....less than 25 months,ex:~26.27.28.29.30 month?

upvoted 2 times

🗨️ 👤 **Antreu** 2 years, 8 months ago

if start\_date+25 is less than sysdate, then start\_date is not in the last 25 months

upvoted 1 times

🗨️ 👤 **zzsqlcoder** 9 months, 2 weeks ago

Yes because if we try a date that is 30 months before, if we add 25 months it is still less than SYSDATE so the condition will result true. That's

why C is the only alternative

upvoted 1 times

Which three are true about scalar subquery expressions? (Choose three.)


- A. They can be nested.
- B. They cannot be used in the VALUES clause of an INSERT statement.
- C. A scalar subquery expression that returns zero rows evaluates to zero.
- D. They can be used as default values for columns in a CREATE TABLE statement.
- E. A scalar subquery expression that returns zero rows evaluates to NULL.
- F. They cannot be used in GROUP BY clauses.

**Suggested Answer:** AEF

Reference:

[https://docs.oracle.com/cd/B12037\\_01/server.101/b10759/expressions010.htm](https://docs.oracle.com/cd/B12037_01/server.101/b10759/expressions010.htm)

Community vote distribution



🗳️ **highwater** 3 months ago

**Selected Answer:** AEF

- A. They can be nested.
  - E. A scalar subquery expression that returns zero rows evaluates to NULL
  - F. They cannot be used in GROUP BY clauses
- upvoted 1 times

🗳️ **alelejaja** 7 months, 1 week ago

**Selected Answer:** AEF

I also think AEF are the correct answers

upvoted 1 times

🗳️ **Ashkush** 8 months, 3 weeks ago

Why is D incorrect?

upvoted 1 times

🗳️ **NSilva** 8 months, 2 weeks ago

D. They can be used as default values for columns in a CREATE TABLE statement.

>> True, but there might be one of two issues.

- 1) The website has marked it incorrectly as "False"
- 2) It is an SQL exam, where wording matters. "You can use it as a default value (IN columns) / (for ROWS) during CREATE TABLE statement"

upvoted 1 times

🗳️ **alelejaja** 7 months, 1 week ago

Sorry to say I don't share your point here. D is falso since, as official doc from Oracle highlights, scalar subqueries are not valid expressions as default values for columns.

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/Scalar-Subquery-Expressions.html#GUID-475D80C3-C873-4475-AB1A-8837C5CF8CE4>

upvoted 4 times

🗳️ **jm9999** 1 year, 3 months ago

**Selected Answer:** BEF

I don't think they can be nested. Tried:

```
select employee_id,
(select department_name, (select city from locations l where l.location_id=d.location_id)
from departments d where d.department_id=e.department_id)
from employees e;
```

But got 'too many values error'.

upvoted 1 times

🗳️ **jm9999** 1 year, 2 months ago

Changing my answer to AEF. I found B to be incorrect. Still not sure how to nest them (A).

upvoted 2 times

  **Tof0ra** 11 months ago



```
select order_id, order_total_value, (select avg(order_total_value from orders where order_date > (select sysdate - 30 from dual)) as
average_30_days, order_total_value - (select avg(order_total_value from orders where order_date > (select sysdate - 30 from dual)) as
above_average_30_days from orders where order_total_value > (select avg(order_total_value) from orders where order_date > (select sysdate -
30 from dual))
```

upvoted 1 times

  **Tof0ra** 11 months ago

This query is returning more than 1 row: select city from locations l where l.location\_id=d.location\_id so 'too many values error', try select  
DISTINCT city from locations l where l.location\_id=d.location\_id

upvoted 1 times

  **zouve** 1 year, 5 months ago

**Selected Answer: AEF**

A. They can be nested. Scalar subqueries can be nested within other scalar subqueries or within other SQL expressions.

E. A scalar subquery expression that returns zero rows evaluates to NULL. If a scalar subquery returns no rows, it evaluates to NULL.

F. They cannot be used in GROUP BY clauses. Scalar subqueries cannot be used in GROUP BY clauses because they do not return a single value that  
can be used to group rows.

upvoted 3 times

  **Chandra1104** 1 year, 9 months ago

Its ABE.

A. They can be nested. --> True

B. They cannot be used in the VALUES clause of an INSERT statement. --> True

C. A scalar subquery expression that returns zero rows evaluates to zero. --> False

D. They can be used as default values for columns in a CREATE TABLE statement.--> False

E. A scalar subquery expression that returns zero rows evaluates to NULL. --> True

F. They cannot be used in GROUP BY clauses.--> False

<https://www.oratable.com/scalar-subquery/>

upvoted 1 times

  **J4vi** 1 year, 7 months ago

AEF

There is no way you could use it in GROUP BY clauses, but it says nothing about VALUES clause.

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/Scalar-Subquery-Expressions.html#GUID-475D80C3-C873-4475-AB1A-8837C5CF8CE4>

upvoted 2 times

  **holdfaststrongly** 2 years, 2 months ago

According to both of your links, wouldn't B also correct because INSERT is part of a DML statement? B says "CANNOT" be used in the value of an  
INSERT statement. Double negative makes the statement positive. Just saying...

upvoted 1 times

  **Artur1991** 2 years, 3 months ago

AEF are the correct answers

[https://docs.oracle.com/cd/B19306\\_01/server.102/b14200/expressions010.htm](https://docs.oracle.com/cd/B19306_01/server.102/b14200/expressions010.htm)

upvoted 2 times

  **heoj10272** 2 years, 2 months ago

I agree with you, but this would be better docs.

<https://docs.oracle.com/database/121/SQLRF/expressions014.htm#SQLRF52093>

this is 12c docs.

upvoted 2 times



Examine this query:


```
SELECT employee_id, first_name, salary
FROM employees
WHERE hire_date > '&1';
```

Which two methods should you use to prevent prompting for a hire date value when this query is executed? (Choose two.)

- A. Use the DEFINE command before executing the query.
- B. Replace '&1' with '&&1' in the query.
- C. Use the UNDEFINE command before executing the query.
- D. Execute the SET VERIFY OFF command before executing the query.
- E. Execute the SET VERIFY ON command before executing the query.
- F. Store the query in a script and pass the substitution value to the script when executing it.

**Suggested Answer: AF**

Community vote distribution



**ArslanAltaf** 8 months, 2 weeks ago

B cannot be an option because the question says "Prevent prompting of hire date value". && will still prompt it once.  
upvoted 3 times

**yanoolthecool** 1 year ago

**Selected Answer: AB**

Still cant get why F is the answer, it literally says 'pass a value', doesnt that translate to prompting the user? as for B, && does an implicit DEFINE on first value entered via the prompt, but if it was already defined then it will just use the value that was already defined, I say AB.  
upvoted 1 times

**yanoolthecool** 1 year ago

looking this up was a bit difficult, but GPT says as I say.  
upvoted 1 times

**winston** 1 year, 2 months ago

Why is B not correct?..and someone explain how o store a query in a script, i cant get the concept..  
upvoted 1 times

**Lee\_jong\_suk** 1 year ago

vi sql\_script.sql  
select \* from employees where employee\_id = '&1';  
to execute the file -> @sql\_script.sql  
then it will prompt for value -> Enter value for 1:101  
here you can enter whatever id you want.  
These are used to generate reports for debugging.  
upvoted 2 times

**Rik92** 1 year, 7 months ago

AF are correct.  
B will prompt once, but question asks for no prompting at all  
C. UNDEFINE is not a thing  
D/E Verify has nothing to do with prompting  
upvoted 2 times

**J4vi** 1 year, 7 months ago

**Selected Answer: AF**

It's confusing, if you replace &1 with &&1 it would still be asking for the value every time you run the query, it doesn't make sense to me... but if you previously used DEFINE to define that variable, it wouldn't. That's why I think AF is more accurate  
upvoted 2 times

🗨️ 👤 **tangerine141** 1 year, 8 months ago

B. By replacing '&1' with '&&1' in the query, it creates a double-ampersand substitution variable which will remember the value of the variable throughout the session and will not prompt for a value every time it is used.

F. By storing the query in a script and passing the substitution value to the script when executing it, you can provide a value for the variable before the query is executed. This prevents the query from prompting for a value when it is executed.

upvoted 3 times

🗨️ 👤 **Orxan\_H** 1 year, 8 months ago

Why not B?

upvoted 1 times

🗨️ 👤 **ninjax\_m** 1 year, 10 months ago

[https://www.oreilly.com/library/view/oracle-sqlplus-](https://www.oreilly.com/library/view/oracle-sqlplus-the/1565925785/ch04s06.html#:~:text=The%20UNDEFINE%20command%20deletes%20a,it%20is%20no%20longer%20needed)

[the/1565925785/ch04s06.html#:~:text=The%20UNDEFINE%20command%20deletes%20a,it%20is%20no%20longer%20needed](https://www.oreilly.com/library/view/oracle-sqlplus-the/1565925785/ch04s06.html#:~:text=The%20UNDEFINE%20command%20deletes%20a,it%20is%20no%20longer%20needed).

The UNDEFINE command deletes a variable definition. If you have created a variable containing sensitive information, such as a password,...

upvoted 1 times

🗨️ 👤 **dexdinh91** 1 year, 11 months ago

**Selected Answer: AF**

There is no UNDEFINE keyword

upvoted 1 times

You need to allow user ANDREW to:

1. Modify the TITLE and ADDRESS columns of your CUSTOMERS table.
2. GRANT that permission to other users.

Which statement will do this?

- A. GRANT UPDATE ON customers.title, customers.address TO andrew;
- B. GRANT UPDATE (title, address) ON customers TO andrew;
- C. GRANT UPDATE (title, address) ON customers TO andrew WITH GRANT OPTION;
- D. GRANT UPDATE ON customers.title, customers.address TO andrew WITH ADMIN OPTION;
- E. GRANT UPDATE ON customers.title, customers.address TO andrew WITH GRANT OPTION;
- F. GRANT UPDATE (title, address) ON customers TO andrew WITH ADMIN OPTION;

**Suggested Answer: C**

Community vote distribution

E (100%)



  **nejak**  2 years, 8 months ago

C is correct: WITH GRANT OPTION is necessary in this case  
upvoted 7 times

  **highwater**  3 months ago

**Selected Answer: C**



C. GRANT UPDATE (title, address) ON customers TO andrew WITH GRANT OPTION;  
upvoted 1 times

  **15eeda4** 5 months, 3 weeks ago

**Selected Answer: C**

C is correct (GRANT option is for object privileges):  
GRANT UPDATE (title, address) ON customers TO andrew WITH GRANT OPTION

E is not correct (ADMIN option is for system privileges):  
GRANT UPDATE (title, address) ON customers TO andrew WITH ADMIN OPTION  
gives the error ORA-00993: missing GRANT keyword.  
upvoted 1 times

  **braintop** 7 months, 1 week ago

**Selected Answer: C**


C is the correct answer.

[https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/GRANT.html#GUID-20B4E2C0-A7F8-4BC8-A5E8-BE61BDC41AC3\\_\\_I2126048](https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/GRANT.html#GUID-20B4E2C0-A7F8-4BC8-A5E8-BE61BDC41AC3__I2126048)  
upvoted 1 times

  **kay000001** 7 months, 3 weeks ago

**Selected Answer: C**

The answer is C.  
C is for Cat.  
upvoted 1 times

  **Drexan** 8 months, 1 week ago

**Selected Answer: C**

Yes, this is the correct syntax according to some sites  
upvoted 1 times

  **a947739** 11 months, 2 weeks ago

C is correct  
upvoted 1 times

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

Why is E not correct?

upvoted 1 times

🗨️ 👤 **MIGHTYSTUNNERS** 1 year, 2 months ago

**Selected Answer: E**

E is also correct

upvoted 2 times

🗨️ 👤 **NSilva** 1 year, 2 months ago

Whats the difference between C & E ?

C. GRANT UPDATE (title, address) ON customers TO andrew WITH GRANT OPTION;

E. GRANT UPDATE ON customers.title, customers.address TO andrew WITH GRANT OPTION;

upvoted 1 times

🗨️ 👤 **Lee\_jong\_suk** 1 year, 6 months ago

C is correct!

WITH GRANT OPTION

now the user ANDREW can give grants to other users.

GRANT UPDATE(TITLE) ON CUSTOMERS TO HR;

If REVOKE UPDATE(TITLE,ADDRESS) ON CUSTOMERS FROM ANDREW

is executed then eventually HR will also lose the grants.

upvoted 1 times

🗨️ 👤 **Chandra1104** 2 years, 3 months ago

C is correct.

The owner of an object can grant it to another user by specifying the WITH GRANT OPTION clause in the GRANT statement. In this case, the new grantee can then grant the same level of access to other users or roles. Here are three points to keep in mind about the WITH GRANT OPTION clause:

[https://www.oreilly.com/library/view/oracle-database-](https://www.oreilly.com/library/view/oracle-database-administration/1565925165/ch06s01s05s01.html#:~:text=The%20owner%20of%20an%20object,to%20other%20users%20or%20roles.)

[administration/1565925165/ch06s01s05s01.html#:~:text=The%20owner%20of%20an%20object,to%20other%20users%20or%20roles.](https://www.oreilly.com/library/view/oracle-database-administration/1565925165/ch06s01s05s01.html#:~:text=The%20owner%20of%20an%20object,to%20other%20users%20or%20roles.)

upvoted 1 times

🗨️ 👤 **tawa\_z58** 2 years, 5 months ago

C is correct

upvoted 2 times

You own table DEPARTMENTS, referenced by views, indexes, and synonyms.

Examine this command which executes successfully:

DROP TABLE departments PURGE;

Which three statements are true? (Choose three.)


- A. It will remove the DEPARTMENTS table from the database.
- B. It will drop all indexes on the DEPARTMENTS table.
- C. It will remove all views that are based on the DEPARTMENTS table.
- D. It will remove all synonyms for the DEPARTMENTS table.
- E. Neither can it be rolled back nor can the DEPARTMENTS table be recovered.
- F. It will delete all rows from the DEPARTMENTS table, but retain the empty table.

**Suggested Answer:** ABE

Reference:

[https://docs.oracle.com/cd/B28359\\_01/server.111/b28310/tables010.htm#ADMIN01505](https://docs.oracle.com/cd/B28359_01/server.111/b28310/tables010.htm#ADMIN01505)

Community vote distribution



Lee\_jong\_suk **Highly Voted** 1 year ago

ABE.

Note: Synonyms and views will get invalid, Indexes will be dropped along with the table and data will be purged and free space will be regained, dropped table can not be recovered even if you use flashback table command

upvoted 5 times

pushpaksarode **Most Recent** 3 months ago

**Selected Answer:** ABE

It will drop DEPARTMENTS table along with all indexes as we have used PURGE in DROP statement so table can't be rollback or recovered.

upvoted 1 times

alelejaja 7 months, 1 week ago

**Selected Answer:** ABE

ABE correct answers

upvoted 1 times

Rik92 1 year, 6 months ago

**Selected Answer:** ABE

ABE are correct. Views and synonyms are not dropped with the table

upvoted 2 times

Nelita 1 year, 7 months ago

**Selected Answer:** ABE

ABE IS CORRECT.

upvoted 1 times

kuff 1 year, 7 months ago

**Selected Answer:** ABE

ABE is the answer

upvoted 1 times

Orxan\_H 1 year, 8 months ago

ABE is correct answers

upvoted 1 times

Ac37 1 year, 9 months ago

ABE is correct

upvoted 1 times

🗨️ 👤 **jfc1** 1 year, 9 months ago

**Selected Answer: ABE**

[https://docs.oracle.com/database/121/SQLRF/statements\\_9003.htm#SQLRF01806](https://docs.oracle.com/database/121/SQLRF/statements_9003.htm#SQLRF01806)

upvoted 1 times

🗨️ 👤 **alic\_alex** 1 year, 10 months ago

**Selected Answer: ADE**

A,D,E - correct answer

upvoted 1 times

🗨️ 👤 **subtype11** 1 year, 11 months ago

[https://docs.oracle.com/cd/E18283\\_01/server.112/e17120/tables010.htm](https://docs.oracle.com/cd/E18283_01/server.112/e17120/tables010.htm)

upvoted 1 times

🗨️ 👤 **subtype11** 1 year, 11 months ago

In ORACLE SQL, you own table DEPARTMENTS, referenced by views, indexes, and synonyms.

Examine this command which executes successfully:

DROP TABLE departments PURGE;

Which three statements are true? (Choose three.)

A. It will remove the DEPARTMENTS table from the database. (True)

B. It will drop all indexes on the DEPARTMENTS table. (True)

C. It will remove all views that are based on the DEPARTMENTS table. (False, it will not remove the views that are based on the DEPARTMENTS table)

D. It will remove all synonyms for the DEPARTMENTS table. (False, it will not remove the synonyms for the DEPARTMENTS table)

E. Neither can it be rolled back nor can the DEPARTMENTS table be recovered. (True)

F. It will delete all rows from the DEPARTMENTS table, but retain the empty table. (False, it will delete all rows and the table itself)

upvoted 2 times

🗨️ 👤 **heoj10272** 2 years, 2 months ago

A, B, E

[https://docs.oracle.com/database/121/SQLRF/statements\\_9003.htm#SQLRF01806](https://docs.oracle.com/database/121/SQLRF/statements_9003.htm#SQLRF01806)

upvoted 1 times

Which three statements are true about Structured Query Language (SQL)? (Choose three.)

- A. It requires that data be contained in hierarchical data storage.
- B. It best supports relational databases.
- C. It provides independence for logical data structures being manipulated from the underlying physical data storage.
- D. It is the only language that can be used for both relational and object-oriented databases.
- E. It guarantees atomicity, consistency, isolation, and durability (ACID) features.
- F. It is used to define encapsulation and polymorphism for a relational table.

**Suggested Answer:** BCE

🗨️ **5cff9a8** 8 months ago

A needs to be heard as well.

upvoted 1 times

🗨️ **speedy\_vin** 8 months ago

I think the answer should be BCD as SQL is used for both RDBMS and Object Oriented DB. Regarding atomicity and consistency, I think that is facilitated by the RDBMS engine.

upvoted 1 times

🗨️ **kay000001** 8 months, 1 week ago

**Selected Answer: BCE**

Answers are B, C, E.

upvoted 1 times

🗨️ **ShahedOdeh** 9 months, 3 weeks ago

A is also true my reference is <https://docs.oracle.com/en/database/oracle/oracle-database/23/sqlrf/Hierarchical-Queries.html#GUID-0118DF1D-B9A9-41EB-8556-C6E7D6A5A84E>

upvoted 1 times

🗨️ **ShahedOdeh** 9 months, 3 weeks ago

so i think the true answer is ACE

upvoted 1 times

🗨️ **ulkaaa** 11 months ago

BCF CORRECT

upvoted 1 times

🗨️ **Lee\_jong\_suk** 1 year, 6 months ago

BCE is the correct answer

upvoted 4 times

🗨️ **steven2827902** 1 year, 9 months ago

answer is BCE

upvoted 4 times

🗨️ **anushamathew12071997** 1 year, 11 months ago

Answer is ACE


upvoted 1 times

Which two statements are true about Oracle synonyms? (Choose two.)

- A. Any user can create a PUBLIC synonym.
- B. A synonym has an object number.
- C. All private synonym names must be unique in the database.
- D. A synonym can be created on an object in a package.
- E. A synonym can have a synonym.

**Suggested Answer:** BE

Community vote distribution



**holdfaststrongly** Highly Voted 2 years, 9 months ago

- A. Any user can create a PUBLIC synonym. (x) Must have Create Public Syn Priv.
- B. A synonym has an object number.
- C. All private synonym names must be unique in the database. (X) Unique in the schema.
- D. A synonym can be created on an object in a package. (X) A schema object can't be in a package
- E. A synonym can have a synonym.

upvoted 16 times

**kenan\_yu** 2 years, 9 months ago

D answer is incorrect for another reason, you actually can create synonym on an object in a PL/SQL package but in won't be active(checked):

```
CREATE OR REPLACE PACKAGE long_package_name AS
FUNCTION give_me_zero RETURN NUMBER;
END;
/
```

```
CREATE OR REPLACE PACKAGE BODY long_package_name AS
FUNCTION give_me_zero RETURN NUMBER IS BEGIN RETURN 0; END;
END;
/
```

```
CREATE OR REPLACE SYNONYM pkg_func FOR long_package_name.give_me_zero;
```

```
SELECT pkg_func.give_me_zero FROM dual;
```

upvoted 2 times

**speedy\_vin** 8 months ago

basically oracle allows any name after for. doesn't mean that it is valid unless it works.

upvoted 1 times

**Lee\_jong\_suk** 1 year, 6 months ago

Right!

upvoted 1 times

**Deni61** Most Recent 2 weeks, 1 day ago

**Selected Answer:** BE

A synonym cannot be created on an object inside a package like procedure or function, but synonym can be created on whole package.

upvoted 1 times

**highwater** 3 months ago

**Selected Answer:** BE

- B. A synonym has an object number
- E. A synonym can have a synonym

upvoted 1 times



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

B,D,E are correct.

B and E are obvious.

Why D is correct : Following workd in oracle.livesql.com

```
CREATE PACKAGE employee_pkg1 AS
```

```
PROCEDURE add_employee(p_first_name VARCHAR2, p_last_name VARCHAR2);
```

```
END employee_pkg1;
```

```
CREATE SYNONYM emp_pkg_syn FOR employee_pkg1;
```

```
--Works
```

```
CREATE SYNONYM add_emp_addemp_syn FOR employee_pkg1.add_employee;
```

```
-- works
```

upvoted 1 times

🗨️ 👤 **JanThanh** 9 months, 2 weeks ago

**Selected Answer: BE**

B & E is correct

upvoted 2 times

🗨️ 👤 **ShahedOdeh** 9 months, 3 weeks ago

C is incorrect because private synonyms are need to be private in the schama not in the database

so

BDE

upvoted 1 times

🗨️ 👤 **ulkaaa** 11 months ago

BE IS CORRECT

upvoted 1 times

🗨️ 👤 **ArslanAltaf** 1 year ago

BE is correct. syn has object. here is to find object ID. Modify it or remove the owner

```
SELECT
```

```
s.synonym_name,
```

```
s.owner,
```

```
o.object_id,
```

```
o.object_name,
```

```
o.object_type
```

```
FROM
```

```
dba_synonyms s
```

```
JOIN
```

```
dba_objects o
```

```
ON
```

```
s.synonym_name = o.object_name
```

```
AND
```

```
s.owner = o.owner
```

```
and S.owner = 'HR'
```

```
;
```

upvoted 1 times

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

**Selected Answer: DE**

With no so many arguments, I would answer DE in an hypothetical exam since documentation says that a synonym can be created on a Stored procedure, function, or package. Besides, I don't find accurate the way B is written since the creation of a new synonym doesn't involve a new object id, but the synonim actually "has" the object id of the object it references...

upvoted 1 times

🗨️ 👤 **Fredderik91** 1 year, 1 month ago

D and E

Specify the object for which the synonym is created. The schema object for which you are creating the synonym can be of the following types:

Table or object table  
View or object view  
Sequence  
Stored procedure, function, or package  
Materialized view  
Java class schema object  
User-defined object type  
Synonym

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/CREATE-SYNONYM.html>

upvoted 1 times

🗲️ 👤 **hmatinnn** 1 year, 5 months ago

**Selected Answer: DE**

de should be

upvoted 1 times

🗲️ 👤 **ismoil** 1 year, 5 months ago

b is wrong

1000%

upvoted 1 times

🗲️ 👤 **id10111110** 1 year, 3 months ago

Assuming user has the "CREATE SYNONYM" grant, then:

select \* from newsyn2 -- produces error

create synonym newsyn2 for dual;

select \* from newsyn2; --returns X

select object\_name, object\_type, object\_id  
from user\_objects;

This seems to suggest that Oracle synonyms have object ids.

upvoted 1 times

🗲️ 👤 **yaya32** 1 year, 5 months ago

**Selected Answer: DE**

DE for me

upvoted 1 times

🗲️ 👤 **mavista** 1 year, 8 months ago

**Selected Answer: BE**

A - user need CREATE PUBLIC SYNONYM : false

C - must be unique in schema : false

D - The schema object cannot be contained in a package for synonyms

upvoted 1 times

🗲️ 👤 **jm9999** 1 year, 9 months ago

**Selected Answer: BE**

Synonym can be created for the whole package but not for components of the package.

upvoted 3 times

🗲️ 👤 **alelejaja** 1 year, 1 month ago

But for a procedure in a package? a synonym can be created in that case

upvoted 1 times

🗲️ 👤 **zouve** 1 year, 12 months ago

**Selected Answer: DE**

D. A synonym can be created on an object in a package. A synonym can be created for an object in a package, allowing users to reference the object using the synonym name instead of the package and object name 1.

E. A synonym can have a synonym. A synonym can be created for another synonym, allowing users to reference the underlying object using either synonym name

upvoted 1 times

  **dilshod** 2 years ago

**Selected Answer: DE**

D and E is true the rest is false

upvoted 1 times

Which is true about the ROUND, TRUNC and MOD functions?

- A.  $\text{TRUNC}(\text{MOD}(25,3),-1)$  is invalid.
- B.  $\text{ROUND}(\text{MOD}(25,3),-1)$  is invalid.
- C.  $\text{ROUND}(\text{MOD}(25,3),-1)$  and  $\text{TRUNC}(\text{MOD}(25,3),-1)$  are both valid and give the same result.
- D.  $\text{ROUND}(\text{MOD}(25,3),-1)$  and  $\text{TRUNC}(\text{MOD}(25,3),-1)$  are both valid but give different results.

**Suggested Answer:** C

Community vote distribution

C (100%)

  **dexdinh91** 11 months, 3 weeks ago

**Selected Answer: C**

Both return 0

upvoted 2 times

  **Artur1991** 1 year, 3 months ago

C is correct

Both return 0

upvoted 2 times

Which two are true about transactions in the Oracle Database? (Choose two.)

- A. DML statements always start new transactions.
- B. DDL statements automatically commit only data dictionary updates caused by executing the DDL.
- C. A session can see uncommitted updates made by the same user in a different session.
- D. A DDL statement issued by a session with an uncommitted transaction automatically commits that transaction.
- E. An uncommitted transaction is automatically committed when the user exits SQL\*Plus.

**Suggested Answer:** DE

Community vote distribution



**holdfaststrongly** Highly Voted 2 years, 9 months ago

- A. DML statements always start new transactions. (X). Not always
- B. DDL statements automatically commit only data dictionary updates caused by executing the DDL. (X) Not "only" data dictionary
- C. A session can see uncommitted updates made by the same user in a different session. (X) in the same session
- D. A DDL statement issued by a session with an uncommitted transaction automatically commits that transaction.
- E. An uncommitted transaction is automatically committed when the user exits SQL\*Plus.

upvoted 9 times

**highwater** Most Recent 3 months ago

**Selected Answer: DE**

- D. A DDL statement issued by a session with an uncommitted transaction automatically commits that transaction
- E. An uncommitted transaction is automatically committed when the user exits SQL\*Plus

upvoted 1 times

**JanThanh** 9 months, 2 weeks ago

**Selected Answer: AD**

hix chatGPT said that

upvoted 2 times

**acd3ea6** 9 months, 2 weeks ago

**Selected Answer: AD**

A : transaction contains of one or more dml statements or DDL statement

D : because DDL make commit before and after DDL statement

[https://docs.oracle.com/cd/E11882\\_01/server.112/e40540/transact.htm#CNCPT016](https://docs.oracle.com/cd/E11882_01/server.112/e40540/transact.htm#CNCPT016)

upvoted 1 times

**archit4321** 1 year, 1 month ago

**Selected Answer: DE**

The answer is DE

upvoted 1 times

**MIGHTYSTUNNERS** 1 year, 2 months ago

**Selected Answer: BD**

bd is correct

upvoted 2 times

**presido11111** 1 year, 4 months ago

The answer is BE

upvoted 1 times

**hmatinnn** 1 year, 5 months ago

**Selected Answer: DE**

should be DE

SET AUTOCOMMIT does not alter the commit behavior when SQL\*Plus exits. Any uncommitted data is committed by default.

upvoted 1 times

🗨️ 👤 **dilshod** 2 years ago

[https://docs.oracle.com/cd/E11882\\_01/server.112/e16604/ch\\_twelve040.htm#BABEGEGC](https://docs.oracle.com/cd/E11882_01/server.112/e16604/ch_twelve040.htm#BABEGEGC)

upvoted 3 times

🗨️ 👤 **DiplomiraniZnalac** 2 years, 9 months ago

True hehe

upvoted 1 times

Examine the description of the MEMBERS table:

Name	Null?	Type
MEMBER_ID	NOT NULL	VARCRAR2 (6)
FIRST_NAME		VARCHAR2 (50)
LAST_NAME	NOT NULL	VARCHAR2 (50)
ADDRESS		VARCHAR2 (50)
CITY		VARCHAR2 (25)

Examine the partial query:

```
SELECT city, last_name AS lname FROM members ...;
```

You want to display all cities that contain the string AN. The cities must be returned in ascending order, with the last names further sorted in descending order.

Which two clauses must you add to the query? (Choose two.)

- A. ORDER BY 1, 2
- B. ORDER BY 1, lname DESC
- C. WHERE city IN ('%AN%')
- D. WHERE city = '%AN%'
- E. WHERE city LIKE '%AN%'
- F. ORDER BY last\_name DESC, city ASC

**Suggested Answer:** BE

Community vote distribution

BE (82%)

Other

 **pushpaksarode** 3 months ago

**Selected Answer: BE**

For '%' or '\_' we have to use LIKE, not "="  
upvoted 1 times


 **Lee\_jong\_suk** 6 months, 3 weeks ago

Option BE is right!  
upvoted 2 times

 **randomArrow** 7 months ago


**Selected Answer: BE**

EF cannot be right. As the question asks us to first sort by cities in the ascending order. Only then further the sorting requires for lastname is descending order. F does vice versa  
upvoted 1 times

 **Nelita** 1 year, 1 month ago

**Selected Answer: BE**

B-E IS CORRECT.  
upvoted 2 times

 **alic\_alex** 1 year, 3 months ago

**Selected Answer: BE**

BE is correct  
upvoted 1 times

 **aghakhan** 1 year, 5 months ago

**Selected Answer: BE**

BE is ok , but i think EF is also right  
upvoted 1 times

 **dexdinh91** 1 year, 5 months ago

Selected Answer: BE

BE is correct


upvoted 1 times

  **trex\_fcs** 1 year, 6 months ago

Selected Answer: BE

Iname is alias and can be used on order by

upvoted 1 times

  **leongkaon** 1 year, 6 months ago

Selected Answer: BE

Iname is in the question

SELECT city, last\_name AS Iname FROM members

upvoted 2 times

  **leongkaon** 1 year, 6 months ago

Selected Answer: BF

Iname is in the question

SELECT city, last\_name AS Iname FROM members



upvoted 1 times

  **apolodor1** 1 year, 7 months ago

Selected Answer: EF

LNAME doesn't exist

upvoted 1 times

  **ninjax\_m** 1 year, 4 months ago

LNAME is an alias for last\_name and order by accepts aliases for sorting.

upvoted 3 times



Examine this partial command:

```
CREATE TABLE cust (
  cust_id NUMBER(2),
  credit_limit NUMBER(10)
)
ORGANIZATION EXTERNAL
```

Which two clauses are required for this command to execute successfully? (Choose two.)

- A. the access driver TYPE clause
- B. the DEFAULT DIRECTORY clause
- C. the REJECT LIMIT clause
- D. the LOCATION clause
- E. the ACCESS PARAMETERS clause

**Suggested Answer:** BD

Community vote distribution

BD (67%)

AD (33%)

🗳️ 👤 **826cb95** 2 months, 2 weeks ago

**Selected Answer:** AB

A. the access driver TYPE clause

✓ Required

This tells Oracle which access driver to use (typically ORACLE\_LOADER or ORACLE\_DATAPUMP).

Example:

sql

Copy

Edit

TYPE ORACLE\_LOADER

B. the DEFAULT DIRECTORY clause

✓ Required

This defines the Oracle directory object where the external file is located.

Example:

sql

Copy

Edit

DEFAULT DIRECTORY ext\_dir

upvoted 1 times

🗳️ 👤 **braintop** 7 months, 1 week ago

LOCATION is the only MUST clause. You can specify the ENTIRE physical path in LOCATION without having to have a DEFAULT\_DIRECTORY clause

upvoted 1 times

🗳️ 👤 **ShahedOdeh** 9 months, 3 weeks ago

The Right answers are :

B and D

as i tried the following create table statement but an error occurred to specify the default directory . The query is :

create table CUST(

cust\_id number (2),  
Credit\_limit Number (10))  
Organization External  
Type Oracle\_Loader  
Location ('c:\test.txt')) ;  
upvoted 1 times

🗨️ 👤 **yolanda\_suny** 10 months, 2 weeks ago

**Selected Answer: AD**

Option A is required because Oracle needs to know how to handle the external data file (driver type).  
Option D is required because Oracle needs to know the location of the external data file to access the data.  
upvoted 1 times

🗨️ 👤 **ulkaaa** 11 months ago

A D TRUE BECAUSE IF WE USE LOCATION THEN WE CAN USE DEFAULT DIRECTORY  
upvoted 2 times

🗨️ 👤 **TofOra** 1 year, 4 months ago

D and E.  
B is WRONG, there is no need to define the DEFAULT DIRECTORY, you can specify each of the files with a DIRECTORY, like:

create directory testmdir as 'c:\tests';

LOGFILE testmdir:'logfile.log'  
BADFILE testmdir:'badfile.log'

LOCATION (testmdir:'invoice\_data.txt')

Tried to upload the full example but the site does not accept it, gives a Cloudflare error.  
upvoted 1 times

🗨️ 👤 **yanoolthecool** 1 year ago

nah, in:  
<https://docs.oracle.com/en/database/oracle/oracle-database/23/sutil/oracle-external-tables-concepts.html#GUID-ACF1D3AA-1D61-4682-AEC5-42C944756E12>  
it says:  
Access parameters are optional. See Access Parameters.  
upvoted 1 times

🗨️ 👤 **J4vi** 2 years, 1 month ago

**Selected Answer: BD**

TYPE: "Oracle Database provides two access drivers: ORACLE\_LOADER and ORACLE\_DATAPUMP. If you omit the TYPE specification, ORACLE\_LOADER is the default access driver."  
[https://docs.oracle.com/cd/B28359\\_01/server.111/b28310/tables013.htm#ADMIN11705](https://docs.oracle.com/cd/B28359_01/server.111/b28310/tables013.htm#ADMIN11705)

ACCESS PARAMETERS: "To modify the default behavior of the access driver for external tables, specify access parameters."  
<https://docs.oracle.com/en/database/oracle/oracle-database/23/sutil/oracle-external-tables-concepts.html#GUID-ACF1D3AA-1D61-4682-AEC5-42C944756E12>

REJECT LIMIT : there are several examples that don't use this clause

LOCATION: "For ORACLE\_HIVE, the LOCATION clause is not used."  
<https://docs.oracle.com/en/database/oracle/oracle-database/23/sutil/oracle-external-tables-concepts.html#GUID-ACF1D3AA-1D61-4682-AEC5-42C944756E12>

That about LOCATION confuse me, but I tend to believe it's a very particular case, so the answer should be (BD) DEFAULT DIRECTORY and LOCATION  
upvoted 4 times

🗨️ 👤 **Sbablo** 2 years, 2 months ago

E is also correct since it species the data source of the external tables.  
upvoted 2 times

  **Orxan\_H** 2 years, 2 months ago

why not E?

upvoted 1 times

Which two are true about unused columns? (Choose two.)

- A. A query can return data from unused columns, but no DML is possible on those columns.
- B. Unused columns retain their data until they are dropped.
- C. Once a column has been set to unused, a new column with the same name can be added to the table.
- D. The DESCRIBE command displays unused columns.
- E. A primary key column cannot be set to unused.
- F. A foreign key column cannot be set to unused.

**Suggested Answer:** BC

Reference:

<https://decipherinfosys.wordpress.com/2007/11/15/back-to-the-basics-dropping-unused-columns-in-oracle/>

*Community vote distribution*

BC (100%)

🗳️ 👤 **yolanda\_suny** 10 months, 2 weeks ago

**Selected Answer: BC**

Option B: "Unused columns retain their data until they are dropped."

This is correct. Unused columns keep their data until they are explicitly removed from the table. Setting a column to unused does not immediately delete its data; it only marks the column as unavailable.

Option C: "Once a column has been set to unused, a new column with the same name can be added to the table."

This is correct. Once a column is marked as unused, you can add a new column with the same name to the table, because the unused column is effectively 'removed' from the table's structure in a logical sense.

upvoted 1 times

🗳️ 👤 **CMjer** 2 years, 1 month ago

**Selected Answer: BC**

Note on C: you can add the same name and data type as unused column even if it's not dropped. Tested and it works.

upvoted 4 times

🗳️ 👤 **Makub** 2 years, 4 months ago

I tested, Unused can be added on primary key.

Desc does not display unused column. not does select query.

upvoted 1 times

🗳️ 👤 **Rik92** 2 years, 5 months ago

B C.

Note on C. You can only add a new column with the same name as long as it's a different data type until the UNUSED is dropped.

upvoted 2 times

Which two are true about the precedence of operators and conditions? (Choose two.)

- A. || has a higher order of precedence than + (addition).
- B. + (addition) has a higher order of precedence than \* (multiplication).
- C. NOT has a higher order of precedence than AND and OR in a condition.
- D. AND and OR have the same order of precedence in a condition.
- E. Operators are evaluated before conditions.

**Suggested Answer:** CE

Community vote distribution

CE (100%)

 **CyberP** Highly Voted 1 year, 1 month ago

The rules are listed in descending order:

Parentheses (as well as brackets in lookup expressions and IN expressions). Note that you can freely add parentheses any time you want to impose an alternative precedence or to make precedence clearer.

\* /

+ -

= <> < > <= >=

IS (IS NULL, IS NOT NULL, IS EMPTY, IS NOT EMPTY)

BETWEEN

NOT

AND

OR

upvoted 11 times

 **826cb95** Most Recent 2 months, 2 weeks ago

**Selected Answer: AC**

✓ A. || (concatenation) vs + (addition)

In Oracle SQL, the concatenation operator || has higher precedence than the addition operator +.

✓ C. NOT vs AND/OR

Correct. In logical expressions, NOT is evaluated before AND, and AND is evaluated before OR.

Order: NOT > AND > OR

upvoted 1 times

 **LuisLikes** 4 months, 1 week ago

**Selected Answer: CE**

|| (OR) doesn't have a higher order of precedence than + (addition). Here I'll show you why:

SELECT 1+1||'1.9' AS addition\_operator FROM DUAL; --output: 21.9

SELECT '1.9'||1+1 AS or\_condition FROM DUAL; --output: 2.91

In the first SELECT statement is first adding up 1+1 (2) and then concatenating with '1.9', making the output '21.9'.

Lastly, the second SELECT statement concatenates '1.9' with 1 ('1.91') and then adding up 1, making the output '2.91'.

This proves that, in this case, whatever comes first in the statement will be calculated.

If there was an order of precedence, the outputs would've been:

OR first: '12.9' (1+(1||'1.9')) and '2.91' (('1.91'||1)+1).

+ first: '21.9' ((1+1)||'1.9') and '1.92' ('1.9'||(1+1)).

So the correct answers are C and E.

upvoted 1 times

🗳️ 👤 **bignosepig** 5 months, 3 weeks ago

**Selected Answer: AC**

|| has a higher order of precedence than + (addition).

upvoted 1 times

🗳️ 👤 **21b6ee4** 6 months, 3 weeks ago

**Selected Answer: AC**

select '9.9'||1+2 from dual

Result 11.91

Contamination is done first

upvoted 1 times

🗳️ 👤 **Szczurk3y** 3 months, 3 weeks ago

select '9.9'+1||2 from dual results in 10,92. Meaning both have the same order precedence.

upvoted 1 times

🗳️ 👤 **alic\_alex** 1 year, 3 months ago

**Selected Answer: CE**

[https://docs.oracle.com/cd/E49933\\_01/server.770/es\\_eql/src/ceql\\_expr\\_precedence\\_rules.html#:~:text=The%20rules%20are,the%20JOIN%20operators](https://docs.oracle.com/cd/E49933_01/server.770/es_eql/src/ceql_expr_precedence_rules.html#:~:text=The%20rules%20are,the%20JOIN%20operators).

upvoted 1 times

🗳️ 👤 **dexdinh91** 1 year, 5 months ago

**Selected Answer: CE**

CE is correct

upvoted 2 times

🗳️ 👤 **holdfaststrongly** 1 year, 9 months ago

This Oracle doc explains between OR, AND, and NOT. <https://docs.oracle.com/en/database/other-databases/nosql-database/22.2/sqlreference/nosql/operator-precedence.html>

upvoted 2 times

🗳️ 👤 **AGoodStuff** 1 year, 9 months ago

[https://www.tutorialspoint.com/plsql/plsql\\_operators\\_precedence.htm](https://www.tutorialspoint.com/plsql/plsql_operators_precedence.htm)

upvoted 1 times

In your session, the NLS\_DATE\_FORMAT is DD-MM-YYYY.

There are 86400 seconds in a day.

Examine this result:

DATE -

02-JAN-2020

Which statement returns this?

- A. SELECT TO\_CHAR(TO\_DATE('29-10-2019') + INTERVAL '2' MONTH + INTERVAL '4' DAY - INTERVAL '120' SECOND, 'DD-MON-YYYY') AS "date" FROM DUAL;
- B. SELECT TO\_CHAR(TO\_DATE('29-10-2019') + INTERVAL '3' MONTH + INTERVAL '7' DAY - INTERVAL '360' SECOND, 'DD-MON-YYYY') AS "date" FROM DUAL;
- C. SELECT TO\_CHAR(TO\_DATE('29-10-2019') + INTERVAL '2' MONTH + INTERVAL '5' DAY - INTERVAL '120' SECOND, 'DD-MON-YYYY') AS "date" FROM DUAL;
- D. SELECT TO\_CHAR(TO\_DATE('29-10-2019') + INTERVAL '2' MONTH + INTERVAL '5' DAY - INTERVAL '86410' SECOND, 'DD-MON-YYYY') AS "date" FROM DUAL;
- E. SELECT TO\_CHAR(TO\_DATE('29-10-2019') + INTERVAL '2' MONTH + INTERVAL '6' DAY - INTERVAL '120' SECOND, 'DD-MON-YYYY') AS "date" FROM DUAL;

**Suggested Answer: C**

Community vote distribution

C (100%)

**zouve** Highly Voted 1 year, 12 months ago

**Selected Answer: C**

Option C is the statement that returns the result 02-JAN-2020. The statement converts the string '29-10-2019' to a date using the TO\_DATE function, then adds an interval of 2 months and an interval of 5 days, and subtracts an interval of 120 seconds. The resulting date is then converted to a string using the TO\_CHAR function with the format model 'DD-MON-YYYY', which produces the result 02-JAN-2020.

Here is the calculation: TO\_DATE('29-10-2019') = October 29, 2019 + INTERVAL '2' MONTH = December 29, 2019 + INTERVAL '5' DAY = January 3, 2020 - INTERVAL '120' SECOND = January 2, 2020

upvoted 5 times

**yanoolthecool** Most Recent 1 year ago

note to self: study the calendar to know when months end, because if Dec ends on 30 instead of 31, the answer would've been A or D

upvoted 1 times

**mrkr01** 1 year, 6 months ago

C because: when we add INTERVAL '2' MONTH + INTERVAL '5' DAY to 29-10-2019, we get 03-01-2020, as of midnight, and then we subtract - INTERVAL '120' SECOND i.e. 2 minutes, hence we get answer as '02-JAN-2020'

upvoted 3 times

**dexdinh91** 2 years, 5 months ago

Correct answer should be

SELECT TO\_CHAR(TO\_DATE('29-OCT-2019') + INTERVAL '2' MONTH + INTERVAL '5' DAY - INTERVAL '120' SECOND, 'DD-MON-YYYY') AS "date" FROM DUAL;

upvoted 4 times

**dexdinh91** 2 years, 5 months ago

NLS\_DATE\_FORMAT is DD-MM-YYYY

upvoted 1 times

**holdfaststrongly** 2 years, 9 months ago

C is correct after testing, although I don't know why it is INTERVAL '5' instead of '4'.

upvoted 2 times

🗨️ 👤 **kakamusili** 2 years, 8 months ago

its 5 cause we are substracting the seconds from the 5th day.

upvoted 2 times

🗨️ 👤 **Rik92** 2 years, 5 months ago

THIS. It's easy to read over the minus operator.

Also the way this is written gives ORA-01843 error.

Correct syntax would be:

```
SELECT to_char(to_date ('29-10-2019', 'DD-MM-YYYY') + INTERVAL '2' MONTH + INTERVAL '5' DAY - INTERVAL '120' SECOND, 'DD-MM-YYYY')  
AS "date" FROM dual;
```

upvoted 1 times

🗨️ 👤 **jr\_C** 2 years, 4 months ago

No error just do ALTER SESSION SET NLS\_DATE\_FORMAT = 'DD-MM-YYYY';

upvoted 2 times



Examine the data in the INVOICES table:

INVOICE_ID	CURRENCY_CODE	RAISED_DATE
1	EUR	01-JAN-2019
2	USD	01-FEB-2019
3	JPY	01-MAR-2019

Examine the data in the CURRENCIES table:

CURRENCY_CODE
JPY
GPB
CAD
EUR
USD

Which query returns the currencies in CURRENCIES that are not present in INVOICES?

A.

```
SELECT * FROM currencies
WHERE NOT EXISTS (
  SELECT NULL FROM invoices WHERE currency_code = currency_code
);
```

B.

```
SELECT * FROM currencies
MINUS
SELECT * FROM invoices;
```

C.

```
SELECT currency_code FROM currencies
MINUS
SELECT currency_code FROM invoices;
```

D.

```
SELECT currency_code FROM currencies
INTERSECT
SELECT currency_code FROM invoices;
```

**Suggested Answer:** C

 **TheOracleWasTaken**  11 months ago

For anyone confused A isn't gonna work because you didn't use aliases so its not comparing the right things.

Here is the correct query

```
create table invoices(
invoice_id number,
currency_code varchar2(10),
raised_date date
);
insert into invoices values(1, 'EUR', to_date('01-jan-2019'));
insert into invoices values(2, 'USD', to_date('01-feb-2019'));
insert into invoices values(3, 'JPY', to_date('01-mar-2019'));
```

```
create table currencies(
currency_code varchar2(10)
);
```

```
insert into currencies values('JPY');
insert into currencies values('GPD');
insert into currencies values('CAD');
```

```
insert into currencies values('EUR');
insert into currencies values('USD');
```

```
select *
from invoices;
```

```
select *
from currencies;
```

```
select *
from currencies c
where not exists (
select null from invoices i where i.currency_code = c.currency_code
);
upvoted 6 times
```

  **babyjaan** Most Recent 1 year, 1 month ago

D IS CORRECT OPTION...AS THE NUMBER OF COLUMN ARE NOT SAME IN BOTH TABLES..PRACTICALLY IMPLEMENTED..CORRECT OPTION D  
upvoted 1 times



  **Rik92** 1 year ago

D uses INTERSECT.

The Oracle INTERSECT operator compares the result of two queries and returns the distinct rows that are output by BOTH queries.

The question was to find the currencies only in ONE of the two tables. So C (MINUS)

upvoted 2 times

  **Rik92** 1 year, 5 months ago

for minus operator one needs to specify the columns.

Intersect gives you the matching data.

So C

upvoted 3 times

  **MahdiHamdii** 1 year, 6 months ago

C is correct.

upvoted 2 times

The SALES table has columns PROD\_ID and QUANTITY\_SOLD of data type NUMBER.

Which two queries execute successfully? (Choose two.)

- A. SELECT prod\_id FROM sales WHERE quantity\_sold > 55000 AND COUNT(\*) > 10 GROUP BY COUNT(\*) > 10;
- B. SELECT prod\_id FROM sales WHERE quantity\_sold > 55000 GROUP BY prod\_id HAVING COUNT(\*) > 10;
- C. SELECT COUNT(prod\_id) FROM sales GROUP BY prod\_id WHERE quantity\_sold > 55000;
- D. SELECT prod\_id FROM sales WHERE quantity\_sold > 55000 AND COUNT(\*) > 10 GROUP BY prod\_id HAVING COUNT(\*) > 10;
- E. SELECT COUNT(prod\_id) FROM sales WHERE quantity\_sold > 55000 GROUP BY prod\_id;

**Suggested Answer:** BE

Community vote distribution

BE (100%)

🗲️ 👤 **jfc1** 11 months ago

**Selected Answer: BE**

BE is the correct one

upvoted 2 times

🗲️ 👤 **Rik92** 11 months, 1 week ago

You cannot have a COUNT() in a WHERE clause. That is what the HAVING clause is for

upvoted 1 times

Which three statements are true about single-row functions? (Choose three.)

- A. They return a single result row per table.
- B. They can be nested to any level.
- C. They can accept only one argument.
- D. The argument can be a column name, variable, literal or an expression.
- E. They can be used only in the WHERE clause of a SELECT statement.
- F. The data type returned can be different from the data type of the argument.

**Suggested Answer:** BDF

Community vote distribution

BDF (100%)

826cb95 2 months, 2 weeks ago

Selected Answer: BDF

BDF is correct. Can be nested to any level.  
upvoted 1 times

dexdinh91 11 months, 3 weeks ago

Selected Answer: BDF

BDF is correct  
upvoted 4 times

holdfaststrongly 1 year, 2 months ago

BDF is correct. Can be nested to any level.  
upvoted 1 times

heoj10272 1 year, 2 months ago

Answer is BDF.

check the discussion of the link.

<https://www.examttopics.com/discussions/oracle/view/20182-exam-1z0-071-topic-2-question-46-discussion/>

upvoted 2 times

holdfaststrongly 1 year, 3 months ago

E is not right either because of the word "only".  
upvoted 2 times

holdfaststrongly 1 year, 3 months ago

DEF.

B is wrong. Nested only to 255 levels.

upvoted 1 times

Orxan\_H 8 months, 1 week ago

this applies to subquery (255 levels). Functions unlimited

upvoted 2 times

Which two statements are true about \*\_TABLES views? (Choose two.)

- A. USER\_TABLES displays all tables owned by the current user.
- B. You must have ANY TABLE system privileges, or be granted object privileges on the table, to view a table in USER\_TABLES.
- C. All users can query DBA\_TABLES successfully.
- D. You must have ANY TABLE system privileges, or be granted object privileges on the table, to view a table in DBA\_TABLES.
- E. ALL\_TABLES displays all tables owned by the current user.
- F. You must have ANY TABLE system privileges, or be granted object privileges on the table, to view a table in ALL\_TABLES.

**Suggested Answer:** AD

Community vote distribution

AF (100%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer:** AD

dba\_tables - requires privileges for querying, all\_tables - does not require privileges for querying.  
upvoted 1 times

🗳️ 👤 **826cb95** 2 months, 2 weeks ago

**Selected Answer:** AD

Not F because ALL\_TABLES shows only those tables the user already has access to, so no special privileges are needed beyond access. Also, the question says 'You must have.'

✓ D. You must have ANY TABLE system privileges, or be granted object privileges on the table, to view a table in DBA\_TABLES.  
True.

To query DBA\_TABLES, you need appropriate system privileges, like SELECT\_CATALOG\_ROLE.  
upvoted 1 times

🗳️ 👤 **Lee\_jong\_suk** 1 year, 6 months ago

AF is the right answer anyday.

USER\_TABLE - dictionary view which has list of tables owned by the user

ALL\_TABLE - dictionary view which has list of tables owned by the user + tables which user has grants for.

DBA\_TABLE - All tables in database (but you should have dba privilege to see them)

upvoted 1 times

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

F is incorrect because if i have only my own tables and no access to others. i don't need ANY TABLE privileges to see my own table in all tables  
upvoted 1 times

🗳️ 👤 **Hrishi41** 1 year, 9 months ago

AD, the right answers.

The data dictionary views that begin with DBA\_ are restricted. These views can be accessed only by users with the SELECT\_ANY\_TABLE privilege. This privilege is assigned to the DBA role when the system is initially installed.

[https://docs.oracle.com/cd/A57673\\_01/DOC/server/doc/SRF73/ch2a.htm](https://docs.oracle.com/cd/A57673_01/DOC/server/doc/SRF73/ch2a.htm)  
upvoted 3 times

🗳️ 👤 **jm9999** 1 year, 8 months ago

granting select any table to my user did not allow him to query dba\_tables  
upvoted 4 times

🗳️ 👤 **HassanMkhlalati** 1 year, 9 months ago

**Selected Answer:** AF

if you don't have dba role then you will get message that DBA\_TABLES doesn't exist.

DBA\_TABLES is a data dictionary view

upvoted 2 times

🗨️ 👤 **TheOracleWasTaken** 1 year, 11 months ago

I tried D and it doesn't work for me!!!!

But its probably A. idk about the 2nd option

upvoted 1 times

🗨️ 👤 **Shrimathi** 1 year, 11 months ago

AD

to view dba\_table should have privilege.

upvoted 2 times

🗨️ 👤 **zouve** 1 year, 12 months ago

**Selected Answer: AF**

i am for AF

upvoted 3 times

🗨️ 👤 **Rik92** 2 years, 5 months ago

AD.

Tried on ORA 19.

GRANT SELECT ANY\_TABLE to gives access to DBA\_TABLES

F is incorrect because querying ALL\_TABLES you can always do but only see those tables you have the rights to view, which would be a subset of DBA\_TABLES

upvoted 4 times

🗨️ 👤 **dexdinh91** 2 years, 5 months ago

**Selected Answer: AF**

AF

cannot grant dba\_tables privilege

upvoted 3 times

🗨️ 👤 **thejivsniper** 2 years, 5 months ago

Explain all answers

upvoted 1 times

🗨️ 👤 **Viviana3184** 2 years, 8 months ago

A, F are correct

upvoted 1 times

🗨️ 👤 **holdfaststrongly** 2 years, 9 months ago

Don't get USER\_TABLE and ALL\_TABLE views confused.

A) USER\_TABLES displays all tables owned by the current user. (True)

E) ALL\_TABLES displays all tables owned by the current user. (False)

The ALL\_TABLES view contains description of tables accessible to the user.

upvoted 3 times

Which two statements are true about conditional INSERT ALL? (Choose two.)

- A. Each row returned by the subquery can be inserted into only a single target table.
- B. A single WHEN condition can be used for multiple INTO clauses.
- C. Each WHEN condition is tested for each row returned by the subquery.
- D. It cannot have an ELSE clause.
- E. The total number of rows inserted is always equal to the number of rows returned by the subquery.

**Suggested Answer:** BC

🗉  **Artur1991** Highly Voted 2 years, 9 months ago

BC are correct

<https://oracle-base.com/articles/9i/multitable-inserts>

upvoted 7 times

🗉  **826cb95** Most Recent 2 months, 2 weeks ago

**Selected Answer: BC**

Not A Because:

✗ A. Each row returned by the subquery can be inserted into only a single target table.

False.

A row can be inserted into multiple tables if multiple WHEN clauses are true, or if a WHEN clause targets multiple INTO tables

upvoted 1 times

🗉  **kay000001** 7 months, 3 weeks ago

**Selected Answer: BC**

B and C are correct.

upvoted 1 times

🗉  **usarovsherzod68** 1 year, 4 months ago

A,c are correct

upvoted 2 times

Which two statements are true about the COUNT function? (Choose two.)

- A. COUNT(\*) returns the number of rows in a table including duplicate rows and rows containing NULLs in any column.
- B. It can only be used for NUMBER data types.
- C. COUNT(DISTINCT inv\_amt) returns the number of rows excluding rows containing duplicates and NULLs in the INV\_AMT column.
- D. COUNT(inv\_amt) returns the number of rows in a table including rows with NULL in the INV\_AMT column
- E. A SELECT statement using the COUNT function with a DISTINCT keyword cannot have a WHERE clause.

**Suggested Answer:** AC

🗒️ 👤 **kakamusili** Highly Voted 👍 2 years, 8 months ago

Group functions do not consider NULL values, except the COUNT(\*)  
upvoted 9 times

🗒️ 👤 **826cb95** Most Recent ⌚ 2 months, 2 weeks ago

Selected Answer: AC  
tested  
upvoted 1 times

🗒️ 👤 **kay000001** 7 months, 3 weeks ago

Selected Answer: AC  
A and C are correct.  
upvoted 1 times

🗒️ 👤 **TheOracleWasTaken** 1 year, 11 months ago

A and C seem correct!  
upvoted 4 times



The EMPLOYEES table contains columns EMP\_ID of data type NUMBER and HIRE\_DATE of data type DATE. You want to display the date of the first Monday after the completion of six months since hiring. The NLS\_TERRITORY parameter is set to AMERICA in the session and, therefore, Sunday is the first day of the week. Which query can be used?

- A. SELECT emp\_id, ADD\_MONTHS(hire\_date, 6), NEXT\_DAY('MONDAY') FROM employees;
- B. SELECT emp\_id, NEXT\_DAY(ADD\_MONTHS(hire\_date, 6), 1) FROM employees;
- C. SELECT emp\_id, NEXT\_DAY(MONTHS\_BETWEEN(hire\_date, SYSDATE), 6) FROM employees;
- D. SELECT emp\_id, NEXT\_DAY(ADD\_MONTHS(hire\_date, 6), 'MONDAY') FROM employees;

**Suggested Answer: D**

Community vote distribution

D (100%)

🗳️ 👤 **826cb95** 2 months, 2 weeks ago

**Selected Answer: D**

Oracle's NEXT\_DAY(date, 'DAY') function returns the next specified weekday after the given date. Since the NLS\_TERRITORY is set to AMERICA, Oracle understands 'MONDAY' correctly.  
upvoted 1 times

🗳️ 👤 **yolanda\_suny** 10 months, 2 weeks ago

**Selected Answer: D**

This query first uses ADD\_MONTHS(hire\_date, 6) to calculate the date six months after the hire date, and then uses NEXT\_DAY to find the first Monday after this date. This is the correct approach and meets the requirements of the problem.  
upvoted 1 times

🗳️ 👤 **CMjer** 2 years, 1 month ago

**Selected Answer: D**

D is correct.  
[https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/NEXT\\_DAY.html#GUID-01B2CC7A-1A64-4A74-918E-26158C9096F6](https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/NEXT_DAY.html#GUID-01B2CC7A-1A64-4A74-918E-26158C9096F6)  
upvoted 3 times

Which three statements are true about GLOBAL TEMPORARY TABLES? (Choose three.)

- A. GLOBAL TEMPORARY TABLE space allocation occurs at session start.
- B. GLOBAL TEMPORARY TABLE rows inserted by a session are available to any other session whose user has been granted select on the table.
- C. A TRUNCATE command issued in a session causes all rows in a GLOBAL TEMPORARY TABLE for the issuing session to be deleted.
- D. Any GLOBAL TEMPORARY TABLE rows existing at session termination will be deleted.
- E. A DELETE command on a GLOBAL TEMPORARY TABLE cannot be rolled back.
- F. A GLOBAL TEMPORARY TABLE'S definition is available to multiple sessions.

**Suggested Answer:** CDF

Community vote distribution

CDF (100%)

 **zouve**  12 months ago

**Selected Answer:** CDF

A global temporary table is a type of table in Oracle that holds data only for the duration of a session or transaction. It is a persistent database object, visible to all sessions, until the user drops it explicitly using the DROP TABLE statement. Data is only visible to the session which inserts it. The data stored in the global temporary table is private to the session1.

In contrast, a regular table is a permanent database object that stores data on disk and is visible to all sessions. The data stored in a regular table is not private to any session and can be accessed by any session with the appropriate privileges.

In summary, the main differences between a global temporary table and a regular table are:

A global temporary table holds data only for the duration of a session or transaction, while a regular table holds data permanently. Data stored in a global temporary table is private to the session, while data stored in a regular table is not private to any session.

upvoted 5 times

 **826cb95**  2 months, 2 weeks ago

**Selected Answer:** CDF

✓ F. A GLOBAL TEMPORARY TABLE'S definition is available to multiple sessions.

True. The table structure is shared among all users, but data is private to each session

upvoted 1 times

 **jm9999** 9 months, 2 weeks ago

When I delete from GTT and then Rollback, response is "Rollback complete" but the rollback does not occur. The rows are still gone from the table

upvoted 3 times

 **TheOracleWasTaken** 11 months ago

**Selected Answer:** CDF

Delete operation can totally be rolled back!

upvoted 1 times

 **holdfaststrongly** 1 year, 9 months ago

<https://www.oracletutorial.com/oracle-basics/oracle-global-temporary-table/> will explain the differences between session or transaction specific GTT.

upvoted 3 times

 **karols32** 1 year, 9 months ago

- A. Can someone explain me when a space allocation occurs?
- B. It's not true because GTT are available for everyone even without permissions like SELECT.
- C. Changing in private session causes only for private data in GTT.
- D. Any rows existing at GTT retain after close a session but are not available.
- E. It can be rolled back, GTT works like a normal table.

F. GTT for everyone LOCAL TEMPORARY TABLE only for a current session.

If what I wrote is not true, please correct me.

upvoted 1 times

  **holdfaststrongly** 1 year, 9 months ago

A) The space is allocated at the time of the first DML operation.


upvoted 5 times

Which two statements are true about the SET VERIFY ON command? (Choose two.)

- A. It displays values for variables used only in the WHERE clause of a query.
- B. It displays values for variables created by the DEFINE command.
- C. It can be used only in SQL\*Plus.
- D. It displays values for variables prefixed with &&.
- E. It can be used in SQL Developer and SQL\*Plus.

**Suggested Answer:** DE

Community vote distribution



826cb95 2 months, 2 weeks ago

**Selected Answer:** CD

✓ C. It can be used only in SQL\*Plus.

True. SET VERIFY ON is a command specific to SQL\*Plus (and tools that emulate it).

It is not supported in GUI tools like SQL Developer (though similar behavior can sometimes be simulated).

✓ D. It displays values for variables prefixed with &&.

True. When SET VERIFY ON is enabled, SQL\*Plus displays the old and new values of substitution variables (like &&var) before executing the SQL.

✗ Incorrect options:

A. It displays values for variables used only in the WHERE clause of a query.

➤ False. It shows substituted values regardless of where they are used, not just in the WHERE clause.

B. It displays values for variables created by the DEFINE command.

➤ False. SET VERIFY ON shows values of substitution variables, not necessarily those created by DEFINE.

E. It can be used in SQL Developer and SQLPlus.\*

➤ False. SET VERIFY ON works only in SQL\*Plus. SQL Developer ignores it.

upvoted 2 times

highwater 3 months ago

**Selected Answer:** BE

BE works well

upvoted 1 times

maksud020 7 months ago

**Selected Answer:** CD

ANSWER CD

upvoted 1 times

acd3ea6 9 months, 2 weeks ago

**Selected Answer:** BE

after testing it does not work with &&

upvoted 1 times

highwater 3 months ago

it does

upvoted 1 times

ShahedOdeh 9 months, 2 weeks ago

The Correct Answer Is : D & E

upvoted 1 times

ArslanAltaf 1 year, 2 months ago

BDE ---All true  
upvoted 1 times

🗳️ 👤 **hmatinnn** 1 year, 5 months ago

**Selected Answer: DE**

Should be DE  
upvoted 2 times

🗳️ 👤 **yaya32** 1 year, 5 months ago

**Selected Answer: BE**

BE for me  
upvoted 1 times

🗳️ 👤 **yanoolthecool** 1 year, 6 months ago

**Selected Answer: DE**

id say DE, B is kinda right?  
upvoted 1 times

🗳️ 👤 **jm9999** 1 year, 9 months ago

BDE all worked for me.  
upvoted 2 times

🗳️ 👤 **jm9999** 1 year, 8 months ago

I think it is just BE. If I make sure the variable is undefined before I use && then I don't get the new and old references in script output.  
upvoted 3 times

🗳️ 👤 **KizTine** 1 year, 8 months ago

I did the test and it also worked for me for answers B, D and E.  
upvoted 1 times

🗳️ 👤 **zouve** 2 years ago

**Selected Answer: BE**

correct is BE  
upvoted 2 times

🗳️ 👤 **zouve** 1 year, 12 months ago

The SET VERIFY command controls whether or not SQLPlus displays before and after images of each line that contains a substitution variable<sup>1</sup>. The correct answers to your question are B and E. The SET VERIFY ON command will enable the display of any substitution variable used in our last executed SQL statement. When you use the DEFINE clause to create a variable, and the variable is not used in the next SQL statement then that variable will not be "shown" by the SET VERIFY command<sup>2</sup>. It can be used in both SQL Developer and SQLPlus.  
upvoted 1 times

🗳️ 👤 **zouve** 2 years ago

correct is BE  
upvoted 2 times

🗳️ 👤 **ESZ** 2 years ago

Is C also correct?  
upvoted 1 times

🗳️ 👤 **SantiBZ\_07032022\_1744** 2 years, 5 months ago

B no right: VERIFY ON shows all substituted variables although you dont use DEFINE (sqlplus/developer ask for them).  
upvoted 2 times

🗳️ 👤 **holdfaststrongly** 2 years, 9 months ago

BDE are true but if I have to narrow it down to 2 answers then DE.  
upvoted 2 times

🗳️ 👤 **karols32** 2 years, 9 months ago

BDE are correct, when SET VERIFY ON it shows us variables defined by a DEFINE and a variables in &&.  
upvoted 4 times

🗳️ 👤 **KizTine** 1 year, 8 months ago

I did the test and it also worked for me for answers B, D and E.  
upvoted 1 times

🗳️ 👤 **Averagenoob** 2 years, 9 months ago

correct is BE  
upvoted 3 times

Examine this list of requirements for a sequence:

1. Name: EMP\_SEQ
2. First value returned: 1
3. Duplicates are never permitted.
4. Provide values to be inserted into the EMPLOYEES.EMPLOYEE\_ID column.
5. Reduce the chances of gaps in the values.

Which two statements will satisfy these requirements? (Choose two.)

- A. CREATE SEQUENCE emp\_seq START WITH 1 INCREMENT BY 1 CYCLE;
- B. CREATE SEQUENCE emp\_seq START WITH 1 INCREMENT BY 1 CACHE;
- C. CREATE SEQUENCE emp\_seq;
- D. CREATE SEQUENCE emp\_seq START WITH 1 INCREMENT BY 1 NOCACHE;
- E. CREATE SEQUENCE emp\_seq NOCACHE;
- F. CREATE SEQUENCE emp\_seq START WITH 1 CACHE;

**Suggested Answer:** DE

Community vote distribution

DE (100%)

🗳️ 👤 **826cb95** 2 months, 2 weeks ago

**Selected Answer:** DE

Reduce chances of gaps: so NOCACHE is preferred, since cached sequences can create gaps if the database shuts down before cache is used.  
upvoted 1 times

🗳️ 👤 **dexdinh91** 11 months, 3 weeks ago

**Selected Answer:** DE

NO CACHE  
upvoted 2 times

🗳️ 👤 **mikom80** 1 year, 3 months ago

if there is cache option sequence get some numbers (for example 10). if oracle used some of this (less than 10) and restart instance occurs, then will be gap, because oracle will cache next 10 numbers.  
upvoted 2 times

🗳️ 👤 **karols32** 1 year, 3 months ago

How reduce the chances of gaps in the values by syntax?  
upvoted 1 times

🗳️ 👤 **Viviana3184** 1 year ago

using NOCACHE  
upvoted 1 times



Which three queries execute successfully? (Choose three.)

- A. SELECT 1 - SYSDATE - DATE '2019-01-01' FROM DUAL;
- B. SELECT SYSDATE - DATE '2019-01-01' - 1 FROM DUAL;
- C. SELECT SYSDATE / DATE '2019-01-01' - 1 FROM DUAL;
- D. SELECT SYSDATE - 1 - DATE '2019-01-01' FROM DUAL;
- E. SELECT (SYSDATE - DATE '2019-01-01') / 1 FROM DUAL;
- F. SELECT 1 / SYSDATE - DATE '2019-01-01' FROM DUAL;

**Suggested Answer:** BDE

Community vote distribution

BDE (100%)

  **826cb95** 2 months, 2 weeks ago

**Selected Answer:** ABD

A is also right why?...

✓ A. SELECT 1 - SYSDATE - DATE '2019-01-01' FROM DUAL;

Valid: Oracle allows arithmetic with dates and numbers.

1 - SYSDATE gives a negative number (interval), and subtracting another DATE results in a valid numeric result.

✓ Executes successfully



upvoted 1 times

  **TheOracleWasTaken** 1 year, 4 months ago

**Selected Answer:** BDE

I am assuming it is supposed to be to\_date

upvoted 2 times

  **ArslanAltaf** 8 months, 2 weeks ago

it worked in 19c exactly like below syntax

SELECT SYSDATE - DATE '2019-01-01' - 1 FROM DUAL

upvoted 1 times

  **J4vi** 1 year, 7 months ago

**Selected Answer:** BDE

Tried all the alternatives

upvoted 4 times




Which two are true about granting object privileges on tables, views, and sequences? (Choose two.)

- A. INSERT can be granted only on tables and sequences.
- B. DELETE can be granted on tables, views, and sequences.
- C. SELECT can be granted on tables, views, and sequences.
- D. ALTER can be granted only on tables and sequences.
- E. REFERENCES can be granted only on tables.

**Suggested Answer:** CD

Community vote distribution



**holdfaststrongly** Highly Voted 2 years, 2 months ago

**Selected Answer:** CD

The following OBJECT PRIV'S can be granted on tables, views and sequences---

UPDATE: tables, views.

REFERENCES: tables, views

ALTER: tables, sequences.

SELECT: tables, views, sequences

upvoted 10 times

**826cb95** Most Recent 2 months, 2 weeks ago

**Selected Answer:** CE

D incorrect for Oracle - only for IBM hierarchial dbms

Per Oracle doc

REFERENCES

Table or materialized view

Enables a user to create a foreign key dependency on a table or materialized view.

The REFERENCES privilege on a parent table implicitly grants SELECT privilege on the parent table.

upvoted 1 times

**ragdoll\_raincoat** 5 months, 2 weeks ago

**Selected Answer:** CE

Object Privilege

Table SELECT, INSERT, UPDATE, DELETE, REFERENCES, ALTER, DROP, INDEX

View SELECT, INSERT, UPDATE, DELETE, ALTER, DROP

Sequence SELECT, USAGE

upvoted 1 times

**alelejaja** 7 months ago

**Selected Answer:** CD

UPDATE: tables, views. REFERENCES: tables, views ALTER: tables, sequences. SELECT: tables, views, sequences. DELETE: tables, views. INSERT: tables, views.

Doc: [https://docs.oracle.com/en/database/oracle/oracle-database/23/sqlrf/GRANT.html#GUID-20B4E2C0-A7F8-4BC8-A5E8-BE61BDC41AC3\\_\\_BGBCIIIEG](https://docs.oracle.com/en/database/oracle/oracle-database/23/sqlrf/GRANT.html#GUID-20B4E2C0-A7F8-4BC8-A5E8-BE61BDC41AC3__BGBCIIIEG)

upvoted 2 times

**wi11111** 1 year, 4 months ago

**Selected Answer:** CD

C and D

upvoted 1 times

🗨️ 👤 **TheOracleWasTaken** 1 year, 5 months ago

C and D because you can grant References on a view. Tried out! :)

upvoted 1 times

🗨️ 👤 **CMjer** 1 year, 7 months ago

**Selected Answer: CE**

CE,

according to this site:

<https://docs.oracle.com/database/121/TTSQL/privileges.htm#TTSQL338>

A. INSERT can be granted only on tables and sequences. (FALSE, table + synonym)

B. DELETE can be granted on tables, views, and sequences. (FALSE, only table)

C. SELECT can be granted on tables, views, and sequences. (TRUE)

D. ALTER can be granted only on tables and sequences. (FALSE, ALTER is system priv. not object)

E. REFERENCES can be granted only on tables. (TRUE, tables + materialised views, which are different than 'normal' views)

upvoted 3 times

🗨️ 👤 **CyberP** 1 year, 7 months ago

But answer E mentioned only on Table so E is clearly wrong :D

It didnt mention and "materialised views" it said Tables "only" which is wrong.

So i think C & D is correct.

upvoted 2 times

Examine the description of the BOOKS table:

Name	Null?	Type
TRANSACTION_ID	NOT NULL	VARCHAR2 (6)
TITLE		VARCHAR2 (40)
AMOUNT		NUMBER (10, 2)
CUSTOMER_ID		VARCHAR2 (6)

The table has 100 rows.

Examine this sequence of statements issued in a new session:

```
INSERT INTO books VALUES ('ADV112', 'Adventures of Tom Sawyer', NULL, NULL);
```

```
SAVEPOINT a;
```

```
DELETE FROM books;
```

```
ROLLBACK TO SAVEPOINT a;
```


```
ROLLBACK;
```

Which two statements are true? (Choose two.)

- A. The second ROLLBACK command replays the delete.
- B. The first ROLLBACK command restores the 101 rows that were deleted and commits the inserted row.
- C. The first ROLLBACK command restores the 101 rows that were deleted, leaving the inserted row still to be committed.
- D. The second ROLLBACK command undoes the insert.
- E. The second ROLLBACK command does nothing.

**Suggested Answer:** CD

Community vote distribution



**ragdoll\_raincoat** 5 months, 2 weeks ago

**Selected Answer: CD**

A. The second ROLLBACK command replays the delete.

This is false. The second rollback undoes the INSERT operation but does not affect the DELETE. The first rollback already restored the rows deleted before the savepoint.

B. The first ROLLBACK command restores the 101 rows that were deleted and commits the inserted row.

This is false. The first rollback restores the deleted rows but does not commit the inserted row. The insert is still in an uncommitted state until the second rollback is issued.

C. The first ROLLBACK command restores the 101 rows that were deleted, leaving the inserted row still to be committed.

This is true. The first rollback restores the deleted rows (100 rows), and the inserted row is still uncommitted, meaning it will be undone by the second rollback.

D. The second ROLLBACK command undoes the insert.

This is true. The second rollback undoes the insert, as it was the most recent uncommitted operation at the time.

E. The second ROLLBACK command does nothing.

This is false. The second rollback does undo the inserted row ('ADV112', 'Adventures of Tom Sawyer'), so it does something.

upvoted 1 times

**alelejaja** 7 months ago

**Selected Answer: CD**

CD are correct

upvoted 1 times

**yaya32** 11 months, 3 weeks ago

**Selected Answer: DE**

For me DE

upvoted 1 times

  **TheOracleWasTaken** 1 year, 4 months ago

**Selected Answer: CD**

Seems right!

upvoted 1 times

  **Orxan\_H** 1 year, 8 months ago

**Selected Answer: CD**

Rollback without savepoint:

- end of the transaction
- delete savepoints
- undoes all changes in the transactions

upvoted 3 times

Which two statements are true about an Oracle database? (Choose two.)

- A. A table can have multiple primary keys.
- B. A column definition can specify multiple data types.
- C. A table can have multiple foreign keys.
- D. A VARCHAR2 column without data has a NULL value.
- E. A NUMBER column without data has a zero value.

**Suggested Answer:** CD

Community vote distribution

CD (100%)

🗨️ 👤 **Mahendra\_19** 11 months ago

Table can have multiple foreign Key:

```
CREATE TABLE ENCOUNTER (  
  ENCT_ID VARCHAR2(25) PRIMARY KEY,  
  ENCT_DATE DATE NOT NULL,  
  PT_MRN NUMBER NOT NULL,  
  ENCT_BP VARCHAR2(10) NOT NULL,  
  ENCT_WT NUMBER(3,2) NOT NULL,  
  ENCT_TEMP NUMBER(3) NOT NULL,  
  ACCT_ID VARCHAR2(25) NOT NULL,  
  PHX_ID VARCHAR2(25) NOT NULL,  
  CONSTRAINT FK_PATIENT FOREIGN KEY (PT_MRN) REFERENCES PATIENT(PT_MRN)  
  CONSTRAINT FK_ACCOUNT FOREIGN KEY (ACCT_ID) REFERENCES ACCOUNT(ACCT_ID)  
  CONSTRAINT FK_PTHISTORY FOREIGN KEY (PHX_ID) REFERENCES PTHISTORY(PHX_ID));  
upvoted 1 times
```

🗨️ 👤 **Kamaljeet001** 1 year, 4 months ago

CD is correct answer. PFB discussion link for same question

<https://www.examttopics.com/discussions/oracle/view/21762-exam-1z0-071-topic-2-question-64-discussion/>

upvoted 1 times

🗨️ 👤 **TheOracleWasTaken** 1 year, 4 months ago

**Selected Answer: CD**

One pk and yeah if you dont insert anything in a row with a with a varchar datatype the value is null and not an empty string which i guess is what they're trying to trick you into?

upvoted 1 times

Examine the data in the EMP table:

ENO	ENAME	SAL	DEPTNO
1001	John	12000	10
1002	Sam	40000	20
1003	Daniel	12000	20
1004	Andrea	5000	10

You execute this query:

```
SELECT deptno AS "Department", AVG(sal) AS AverageSalary, MAX(sal) AS "Max Salary"
FROM emp
WHERE sal >= 12000
GROUP BY "Department"
ORDER BY AverageSalary;
```

Why does an error occur?

- A. An alias name must not contain space characters.
- B. An alias name must always be specified in quotes.
- C. An alias name must not be used in an ORDER BY clause.
- D. An alias name must not be used in a GROUP BY clause.

**Suggested Answer:** D

Community vote distribution

D (100%)

🗳️ 👤 **Orxan\_H** Highly Voted 1 year, 8 months ago

**Selected Answer:** D

alias does not work with Group by and Having

upvoted 6 times

🗳️ 👤 **shotcom** 1 year, 3 months ago

from 23c we can use alias names in the group by and having clause

<https://oracle-base.com/articles/23c/group-by-and-having-clause-using-column-alias-or-column-position-23c>

upvoted 2 times

🗳️ 👤 **NSilva** 8 months, 3 weeks ago

Thanks!

I started stuying recently and I wasnt able to see any Alias related error.

upvoted 1 times

🗳️ 👤 **ogi33** 8 months ago

Exam has been validated for product version Oracle Database 11.2.0.1.0 and up to 19c

upvoted 3 times

Which two actions can you perform with object privileges? (Choose two.)


- A. Create roles.
- B. Create FOREIGN KEY constraints that reference tables in other schemas.
- C. Delete rows from tables in any schema except SYS.
- D. Set default and temporary tablespaces for a user.
- E. Execute a procedure or function in another schema.

**Suggested Answer: AE**

Reference:

<http://www.cruzroja.es/help/wvtdbobjp.htm>

Community vote distribution



**heoj10272** Highly Voted 2 years, 8 months ago

BE

A : False. Creating role is system privilege.

B : True. Use REFERENCE privilege.

C : False. Use DELETE ANY TABLE system privilege.

D : False?

E : True. Use EXECUTE privilege.

upvoted 16 times

**alic\_alex** Highly Voted 2 years, 3 months ago

**Selected Answer: BE**

B, E - correct

see description in the table 7-2

<https://docs.oracle.com/database/121/TTSQ/privileges.htm#TTSQ1341>

upvoted 7 times

**bhp12** 2 years, 2 months ago

As per your source:

Privilege: DELETE; Description:

Enables a user to delete from a table..

Privilege: REFERENCES; Enables a user to create a foreign key dependency on a table or materialized view

So, BC

upvoted 4 times

**826cb95** Most Recent 2 months, 2 weeks ago

**Selected Answer: BE**

E. Execute a procedure or function in another schema. ✓

If you're granted the EXECUTE privilege on a procedure/function in another schema, you can run it.

upvoted 1 times

**maksud020** 7 months ago

**Selected Answer: BC**

ANSWER BC

upvoted 1 times

**acd3ea6** 9 months, 2 weeks ago

SELECT \* FROM DBA\_SYS\_PRIVS;

SELECT \* FROM user\_tab\_privs;

upvoted 1 times

🗨️ 👤 **yolanda\_suny** 10 months, 2 weeks ago

**Selected Answer: BE**

B. Correct. To create foreign key constraints in one schema that reference tables in another schema, you must have appropriate object permissions on the target table, such as the REFERENCES permission. This is a typical application of object permissions.

E Correct. Executing a procedure or function in another schema requires having the EXECUTE permission on that procedure or function, which is also a type of object permission.

upvoted 1 times

🗨️ 👤 **ArslanAltaf** 1 year ago

`BE is correct

upvoted 1 times

🗨️ 👤 **hmatinnn** 1 year, 4 months ago

**Selected Answer: BE**

should be bE

upvoted 1 times

🗨️ 👤 **hmatinnn** 1 year, 6 months ago

B and E. A can not be correct cause it is system privilege.

upvoted 1 times

🗨️ 👤 **ESZ** 1 year, 9 months ago

**Selected Answer: BE**

BE for me, please update the answers

upvoted 3 times

🗨️ 👤 **HassanMkhlalati** 1 year, 9 months ago

**Selected Answer: BE**

You can delete from a table under sys schema if object priv is granted

upvoted 2 times

🗨️ 👤 **zouve** 2 years ago

System privileges allow a user to perform a particular database operation or class of database operations. For example, to create a table, the user needs the create table privilege.

Objects have privileges associated with them, such as insert, update and delete a table.

upvoted 1 times

🗨️ 👤 **zouve** 1 year, 11 months ago

BE for me

upvoted 2 times

🗨️ 👤 **jfc1** 2 years, 5 months ago

**Selected Answer: BC**

A is wrong because : Create Role is a system privilege

upvoted 6 times

🗨️ 👤 **dexdinh91** 2 years, 5 months ago

**Selected Answer: BC**

ADE is system privileges

upvoted 3 times

🗨️ 👤 **paddy95** 2 years, 7 months ago

A You must have the CREATE ROLE system privilege.

B Object privileges REFERENCES object type Table or materialized view

Enables a user to create a foreign key dependency on a table or materialized view.

The REFERENCES privilege on a parent table implicitly grants SELECT privilege on the parent table




C System privileges DELETE ANY TABLE Enables a user to delete from any table in the database.

D I think it's System privileges

E Object privileges: Object type PL/SQL package, procedure or function;

Enables a user to execute a PL/SQL package, procedure or function directly.

upvoted 2 times

  **iuliana23** 2 years, 8 months ago

**Selected Answer: BE**

Delete is an object privilege, same as execute, I think B and E are correct, please correct me if I'm wrong

upvoted 3 times

  **Viviana3184** 2 years, 8 months ago

DELETE ANY TABLE : Enables a user to delete from any table in the database . -> this is a System privileges that is different from the object privilege DELETE that enables a user to delete from a table.

Than the correct answer are BE

upvoted 2 times

No user-defined locks are used in your database.

Which three are true about Transaction Control Language (TCL)? (Choose three.)

- A. ROLLBACK without the TO SAVEPOINT clause undoes all the transaction's changes, releases its locks, and erases all its savepoints.
- B. ROLLBACK without the TO SAVEPOINT clause undoes all the transaction's changes but does not release its locks.
- C. ROLLBACK without the TO SAVEPOINT clause undoes all the transaction's changes but does not erase its savepoints.
- D. ROLLBACK TO SAVEPOINT undoes the transaction's changes made since the named savepoint and then ends the transaction.
- E. COMMIT ends the transaction and makes all its changes permanent.
- F. COMMIT erases all the transaction's savepoints and releases its locks.

**Suggested Answer:** AEF

*Community vote distribution*

AEF (100%)

🗳️ 👤 **Orxan\_H** 8 months, 2 weeks ago

**Selected Answer:** AEF

AEF is correct

upvoted 1 times

🗳️ 👤 **dexdinh91** 11 months, 3 weeks ago

**Selected Answer:** AEF

AEF is correct answer

upvoted 1 times

Examine the description of the EMPLOYEES table:

Name	Null?	Type
EMPLOYEE_ID	NOT NULL	NUMBER(38)
DEPARTMENT_ID	NOT NULL	NUMBER(38)
MANAGER_ID		NUMBER(38)

Which two queries return rows for employees whose manager works in a different department? (Choose two.)

A.

```
SELECT emp.*
FROM employees emp
WHERE manager_id NOT IN (
  SELECT mgr.employee_id
  FROM employees mgr
  WHERE emp.department_id <> mgr.department_id
) ;
```

B.

```
SELECT emp.*
FROM employees emp
JOIN employees mgr
  ON emp.manager_id = mgr.employee_id
AND emp.department_id <> mgr.department_id;
```

C.

```
SELECT emp.*
FROM employees emp
LEFT JOIN employees mgr
  ON emp.manager_id = mgr.employee_id
AND emp.department_id <> mgr.department_id;
```


D.

```
SELECT emp.*
FROM employees emp
WHERE NOT EXISTS (
  SELECT NULL
  FROM employees mgr
  WHERE emp.manager_id = mgr.employee_id
  AND emp.department_id <> mgr.department_id
) ;
```

E.

```
SELECT emp.*
FROM employees emp
RIGHT JOIN employees mgr
  ON emp.manager_id = mgr.employee_id
AND emp.department_id <> mgr.department_id
WHERE emp.employee_id IS NOT NULL;
```

Suggested Answer: BE

 **NI\_maruf** 1 year, 11 months ago


Can anyone please explain why option E is the right answer here? why not C?

upvoted 1 times

 **ogi33** 11 months, 2 weeks ago

C may contain employees without a manager, that is, the boss of the bosses ;)

upvoted 1 times

 **Lydia1054** 1 year, 11 months ago

To the E part. AND is part of JOIN and it is RIGHT JOIN, so in the left table (emp) only the result meet the condition emp.manager\_id = mgr.employee\_id and mgr.department\_id <> emp.department\_id will show up in the left table; otherwise will be NULL. Therefore, adding WHERE to filter the NULL out of the result.

upvoted 1 times

  **Lydia1054** 1 year, 11 months ago

C has to change AND to WHERE then the result will be correct.

Use AND then the emp.department\_id <> mgr.department\_id is part of the JOIN and it will show all the rows in the left table, it not filter out any rows.

upvoted 1 times

  **Shrimathi** 1 year, 11 months ago

```
create table emp(emp_id number,dept_id number,manager_id number);
```

```
insert all
```

```
into emp values(1,5,50)
```

```
into emp values(2,5,50)
```

```
into emp values(3,5,5)
```

```
into emp values(50,5,null)
```

```
into emp values(5,4,null)
```

```
select * from dual;
```

```
commit;
```

```
select e.* from emp e
```

```
right join emp m
```

```
on e.manager_id=m.emp_id
```

```
and e.dept_id<>m.dept_id
```

```
where e.emp_id is not null;
```

output:

```
3 5 5
```

E is correct, since it is mentioned e.employee\_id is not null

Right out join it will return matched row of both table and unmatched row that is all row from the emp m table. e.emp\_id is not null this condition retrieves only the employee whose manager works in different department.


upvoted 3 times

Which three are true about dropping columns from a table? (Choose three.)

- A. A column must be set as unused before it is dropped from a table.
- B. A primary key column cannot be dropped.
- C. Multiple columns can be dropped simultaneously using the ALTER TABLE command.
- D. A column can be removed only if it contains no data.
- E. A column that is referenced by another column in any other table cannot be dropped.
- F. A column drop is implicitly committed.

**Suggested Answer:** CEF

Community vote distribution



🗳️ **maksud020** 7 months ago

**Selected Answer:** CE

answer CEF

upvoted 1 times

🗳️ **alelejaja** 1 year, 1 month ago

**Selected Answer:** CF

You can drop a column that is referenced by another column with CASCADE CONSTRAINTS clause (E), so I only find here 2 correct answers (CF)

upvoted 2 times

🗳️ **yaya32** 1 year, 5 months ago

**Selected Answer:** CEF

CEF for me

upvoted 1 times

🗳️ **jm9999** 1 year, 9 months ago

For E, using cascade constraints, I was able to delete a primary key column that was referenced by a foreign key. C and F seem like the only truly correct options.

upvoted 3 times

🗳️ **TheOracleWasTaken** 1 year, 11 months ago

**Selected Answer:** CEF

If use have on delete then E is indeed right

upvoted 1 times

🗳️ **J4vi** 2 years, 1 month ago

**Selected Answer:** CF

CF are the most correct, the other options have some special cases

upvoted 2 times

🗳️ **dexdinh91** 2 years, 5 months ago

E is partial correct because you can use ON DELETE clause

upvoted 4 times

🗳️ **KizTine** 1 year, 8 months ago

Yes, it is possible to drop this type of column with :

ALTER TABLE tableName DROP COLUMN colName CASCADE CONSTRAINT;

upvoted 2 times

Which three statements are true about views in an Oracle Database? (Choose three.)

- A. A SELECT statement cannot contain a WHERE clause when querying a view containing a WHERE clause in its defining query.
- B. Views have no segment.
- C. Views have no object number.
- D. Views can join tables only if they belong to the same schema.
- E. A view can be created that refers to a non-existent table in its defining query.
- F. Rows inserted into a table using a view are retained in the table if the view is dropped.

**Suggested Answer:** BEF

Community vote distribution

BEF (100%)

🗲 👤 **ShahedOdeh** 9 months, 2 weeks ago

the correct answer is B and e and F  
upvoted 1 times

🗲 👤 **cadcadley** 1 year, 7 months ago

what that is mean no segment about view  
upvoted 1 times

🗲 👤 **ArslanAltaf** 1 year ago

segment mean physical storage ...  
upvoted 1 times

🗲 👤 **KizTine** 1 year, 8 months ago

I don't understand the answer E.  
How can views refers to non-existent table ?  
upvoted 4 times

🗲 👤 **lucemqy** 1 year, 7 months ago

You can create view to non-existent table or table with invalid columns by using FORCE option in the CREATE VIEW command. When you query the data from the view you will get an error  
upvoted 4 times

🗲 👤 **TheOracleWasTaken** 1 year, 11 months ago

**Selected Answer: BEF**

select object\_id, object\_name  
from user\_objects

In this table you will also find views so they do have an object number. Views have no segment since they don't take up physical space  
upvoted 4 times

You start a session and execute these commands successfully:

```
CREATE GLOBAL TEMPORARY TABLE invoices_gtt (
  customer_id INTEGER,
  invoice_total NUMBER(10,2)
) ON COMMIT PRESERVE ROWS;
```

```
INSERT INTO invoices_gtt VALUES (1, 100);
```


```
COMMIT;
```

Which two are true? (Choose two.)

- A. To drop the table in this session, you must first truncate it.
- B. Other sessions can view the committed row.
- C. You can add a column to the table in this session.
- D. You can add a foreign key to the table.
- E. When you terminate your session, the row will be deleted.

**Suggested Answer:** CE

Community vote distribution



**Kashka** Highly Voted 2 years, 9 months ago

1. CREATE GLOBAL TEMPORARY TABLE my\_temp\_table (
 id NUMBER,
 description VARCHAR2(20)
 )
 ON COMMIT PRESERVE ROWS
 2. INSERT INTO my\_temp\_table VALUES (1, 'ONE');
 3. Commit;
 4. ALTER TABLE my\_temp\_table
 ADD ID\_1 VARCHAR(20); --> ORA-14450: attempt to access a transactional temp table already in use
 5. drop table my\_temp\_table; --> ORA-14452: attempt to create, alter or drop an index on temporary table already in use
 6. truncate table my\_temp\_table;
 7. drop table my\_temp\_table;
 upvoted 13 times

**Artur1991** 2 years, 9 months ago

thats correct  
upvoted 1 times

**jsalas** 2 years, 9 months ago

Based on your explanation, A is a valid answer?  
upvoted 1 times

**Massy** Most Recent 2 weeks, 4 days ago

Selected Answer: CE

How can someone answer A and take this exam?  
upvoted 1 times

**8c2fb57** 1 month ago

Selected Answer: CE

ANSWER CE  
upvoted 1 times

**826cb95** 2 months, 2 weeks ago

Selected Answer: CE

Here is why A is wrong:

✗ A. To drop the table in this session, you must first truncate it.

False.

Truncating is not required before dropping a table in Oracle.

You can drop a temporary table at any time (assuming permissions).

✓ No dependency on truncation.

upvoted 1 times

🗨️ **maksud020** 7 months ago

**Selected Answer: CE**

ANSWER CE

upvoted 1 times

🗨️ **yolanda\_suny** 10 months, 2 weeks ago

**Selected Answer: CE**

A. Incorrect. Dropping a table is unrelated to whether it has been truncated. If you have sufficient privileges, you can drop the table directly without needing to truncate it first.

B. Incorrect. In a global temporary table, data is isolated per session, and other sessions cannot see the data in the current session, even if it has been committed.

C. Correct. If you have sufficient privileges, you can modify the table structure in the current session, including adding new columns.

D. Incorrect. Generally, you cannot add a foreign key to a global temporary table because a foreign key requires referencing a permanent table, and the rows in a global temporary table are cleared at the end of the session.

E. Correct. With a global temporary table using ON COMMIT PRESERVE ROWS, all data inserted during the session will be automatically cleared when the session ends.

upvoted 1 times

🗨️ **alelejaja** 1 year, 1 month ago

**Selected Answer: CD**

At least in 23c C and E are correct.

I haven't faced any error when adding a column, so C correct

I could drop the table without truncating it first -- A incorrect

upvoted 1 times

🗨️ **Fredderik91** 1 year, 1 month ago

**Selected Answer: AE**

A and E, test in 19c.

C in the session error ORA 14450 attempt to access a transactional temp table already in use, FIRST end session and add column, but no in the same session

upvoted 1 times

🗨️ **ac77458** 1 year, 2 months ago

**Selected Answer: AE**

You get the following error if you don't truncate table invoices\_gtt

drop table invoices\_gtt

Error report:

SQL Error: ORA-14452: attempt to create, alter or drop an index on temporary table already in use

14452. 00000 - "attempt to create, alter or drop an index on temporary table already in use"

\*Cause: An attempt was made to create, alter or drop an index on temporary table which is already in use.

\*Action: All the sessions using the session-specific temporary table have

to truncate table and all the transactions using transaction

specific temporary table have to end their transactions.

upvoted 1 times

🗨️ **hmatinnn** 1 year, 5 months ago

**Selected Answer: AE**

Should be A and E tried it out

upvoted 1 times



🗄️ 👤 **yanoolthecool** 1 year, 6 months ago

**Selected Answer: CE**

why is C wrong? google tells me its possible

'A temporary table can be altered in the same way as a permanent base table although there is no official support to toggle the behavior of the ON COMMIT clause.'

-some article

upvoted 1 times

🗄️ 👤 **TheOracleWasTaken** 1 year, 11 months ago

**Selected Answer: AE**

You cant add an fk to a temporary table. Can't refer to it and cannot refer to another table from the temporary table. Tried it out!

upvoted 2 times

🗄️ 👤 **kuff** 2 years, 1 month ago

**Selected Answer: AE**

AE is the answer

upvoted 1 times

🗄️ 👤 **jackaongao** 2 years, 3 months ago

**Selected Answer: AE**

A is correct ,you can try execute the sql ,and you would get the err msg:ORA-03290

E is correct,when you terminate your session,the row will be deleted,you can try it.

upvoted 1 times

🗄️ 👤 **paddy95** 2 years, 7 months ago

DDL operation on global temporary tables

It is not possible to perform a DDL operation (except TRUNCATE) on an existing global temporary table if one or more sessions are currently bound to that table.

upvoted 2 times

🗄️ 👤 **Viviana3184** 2 years, 8 months ago

for me , correct are

A. To drop the table in this session, you must first truncate it.

E. When you terminate your session, the row will be deleted.

upvoted 2 times

🗄️ 👤 **iuliana23** 2 years, 8 months ago

I think A and E are correct. You have a row inserted, so first you have to truncate the table in order to drop it. When you end the session, the table will be dropped, what it means that the row will be also deleted. Please correct me if I am wrong.

upvoted 1 times

Examine this statement:


```
SELECT last_name
FROM employees
ORDER BY CASE WHEN salary = (SELECT MAX(salary) FROM employees)
              THEN 'A'
              ELSE last_name
              END, last_name DESC;
```

Which two statements are true? (Choose two.)

- A. The names of employees earning the maximum salary will appear first in an unspecified order.
- B. All remaining employee names will appear in descending order.
- C. All remaining employee names will appear in an unspecified order.
- D. All remaining employee names will appear in ascending order.
- E. The names of employees earning the maximum salary will appear first in ascending order.
- F. The names of employees earning the maximum salary will appear first in descending order.

**Suggested Answer:** DF

Community vote distribution



**yanoolthecool** Highly Voted 1 year, 6 months ago

**Selected Answer:** AB

- A. The names of employees earning the maximum salary will appear first in an unspecified order: Because if there more than one employ with the same salary the order is unspecified cuz both will use 'A' letter.
- B. All remaining employee names will appear in descending order. Because there is last\_name DESC

I should just google the answers instead of depending on the ones here...  
upvoted 6 times

**Eren1998** Most Recent 3 months, 4 weeks ago

**Selected Answer:** DE

I tested, D & F are right. Order by 'A' means nothing, are not sorted. all of salary > avg salary rows that are not sorted at first, then they are sorted by 'last\_name DESC;' at the end of query.  
upvoted 1 times

**acting\_dowry\_chill** 4 months, 1 week ago

**Selected Answer:** DF

see answer by ninjax\_m further down:

- D. All remaining employee names will appear in ascending order.  
Ascending is default ordering. case - else outputs employee name that should be sorted ascending.
  - F. The names of employees earning the maximum salary will appear first in descending order.  
They appear first because of 'A', then after that they are sorted by name, but in descending order (2nd order by clause).
- upvoted 1 times

**Thameur01** 5 months, 1 week ago

**Selected Answer:** BF

- B. All remaining employee names will appear in descending order.  
Correct: The last\_name DESC clause ensures this.
  - F. The names of employees earning the maximum salary will appear first in descending order.  
Correct: The last\_name DESC clause ensures this.
- upvoted 1 times

**safasayed19** 8 months, 2 weeks ago

A and B IS CORRECT ANSWER

upvoted 1 times

🗨️ 👤 **Psili** 9 months ago

Chat GPT says A, B

The two correct answers are:

A. The names of employees earning the maximum salary will appear first in an unspecified order.

In the CASE statement, when the employee's salary matches the maximum salary (MAX(salary)), the value 'A' is assigned. Since 'A' comes before any other string alphabetically, employees with the maximum salary will be listed first. However, the order of employees with the maximum salary is unspecified unless there is a secondary sort defined for this group (which is not in this query).

B. All remaining employee names will appear in descending order.

After the maximum salary employees are displayed, all other employees will be ordered by their last names in descending order due to the clause ORDER BY last\_name DESC.

upvoted 1 times

🗨️ 👤 **Nennuzzo** 1 year, 6 months ago

**Selected Answer: DF**

DF is correct

upvoted 2 times

🗨️ 👤 **Shrimathi** 1 year, 11 months ago

DF is the correct

create table emp(last\_name varchar2(100), salary number);

insert all

into emp values('Shri',5000)

into emp values('van',5000)

into emp values('ben',15000)

into emp values('zoo',15000)

into emp values('cat',4000)

select \* from dual;

output:

zoo

ben

Shri

cat

van

upvoted 4 times

🗨️ 👤 **Shrimathi** 1 year, 11 months ago

zoo,ben contains highest salary , it appears in descending order.

And remaining orders appear in an unspecified order

upvoted 1 times

🗨️ 👤 **Shrimathi** 1 year, 11 months ago

zoo,ben contains highest salary , it appears in descending order.

And remaining orders appear in ascending order

upvoted 2 times

🗨️ 👤 **yanoolthecool** 1 year, 11 months ago

I was also a bit confused as they appeared unordered, but looking closer they are, note that the 'S' is capital, which has lower ANSI value.

upvoted 2 times

🗨️ 👤 **samuele1981** 1 year, 11 months ago

**Selected Answer: AB**

A. The names of employees earning the maximum salary will appear first in an unspecified order: Because if there more than one employ whit the same salary the order is unspecified.

B. All remaining employee names will appear in descending order. Because there is last\_name DESC

upvoted 2 times

🗨️ 👤 **JuciDelevati** 2 years, 1 month ago

**Selected Answer: DF**

DF is correct

upvoted 2 times

🗨️ 👤 **Nelita** 2 years, 1 month ago

**Selected Answer: DE**

D/E ARE CORRECT.

upvoted 1 times

🗨️ 👤 **alic\_alex** 2 years, 3 months ago

**Selected Answer: DF**

DF is correct, I have checked in DB

upvoted 2 times

🗨️ 👤 **cadcadley** 2 years, 5 months ago

DF is correct

upvoted 1 times

🗨️ 👤 **dexdinh91** 2 years, 5 months ago

**Selected Answer: DF**

DF is correct

upvoted 2 times

🗨️ 👤 **dexdinh91** 2 years, 5 months ago

<https://www.educba.com/sql-order-by-case/>

upvoted 1 times

🗨️ 👤 **Ramona13** 2 years, 7 months ago

BE is correct

upvoted 2 times

🗨️ 👤 **iuliana23** 2 years, 8 months ago

Can someone explain me why D and F are correct?

upvoted 1 times

🗨️ 👤 **karols32** 2 years, 7 months ago

I tested it and these answers are correct.

Employees which have not a highest salary appear in ascending order and it's true because ELSE starts working and it sorts by last\_name ASC(as default for sorting without giving any postfix).

But the second answer, F one is interesting because where it appear, depends from a character after THEN. We have 'A' now and employees with highest salary comes first but when I replace it by 'B' employees appear from 5th row, by 'C' from 14th row.

It sorts by last\_name DESC, take it from end of query'.

upvoted 3 times

🗨️ 👤 **ninjax\_m** 2 years, 4 months ago

D. All remaining employee names will appear in ascending order.

Ascending is default ordering. case - else outputs employee name that should be sorted ascending.

F. The names of employees earning the maximum salary will appear first in descending order.

They appear first because of 'A' , then after that they are sorted by name, but in descending order (2nd order by clause).


upvoted 3 times

Which two are true about external tables that use the ORACLE\_DATAPUMP access driver? (Choose two.)

- A. When creating an external table, data can be selected only from a table whose rows are stored in database blocks.
- B. Creating an external table creates a directory object.
- C. When creating an external table, data can be selected from another external table or from a table whose rows are stored in database blocks.
- D. Creating an external table creates a dump file that can be used by an external table in the same or a different database.
- E. Creating an external table creates a dump file that can be used only by an external table in the same database.

**Suggested Answer:** CD

Community vote distribution



🗳️ 👤 **Psili** 9 months ago

ChatGPT says also C and D:

True Statements:

C. When creating an external table, data can be selected from another external table or from a table whose rows are stored in database blocks.

This statement is correct because external tables can be created based on data in other external tables, as well as regular database tables. This flexibility allows for various data processing scenarios where data can be accessed from multiple sources.

D. Creating an external table creates a dump file that can be used by an external table in the same or a different database.

This statement is also true. The ORACLE\_DATAPUMP access driver allows the creation of dump files that can be utilized across different databases. This enables data movement and integration between Oracle databases effectively.

upvoted 1 times

🗳️ 👤 **MIGHTYSTUNNERS** 1 year, 1 month ago

**Selected Answer: BD**

BD IS CORRET FOR SURE

upvoted 2 times

🗳️ 👤 **Psili** 9 months ago

ChatGPT answer about B:

B. Creating an external table creates a directory object.

While external tables do require a directory object for file storage references, creating the external table itself does not inherently create a directory object; the directory must exist beforehand.

upvoted 2 times

🗳️ 👤 **yanoolthecool** 1 year ago

uhh, no.

'You must create the directory object before you create the external table; otherwise, an error is generated'

source: <https://docs.oracle.com/en/database/oracle/oracle-database/23/sutil/oracle-external-tables-concepts.html#GUID-ACF1D3AA-1D61-4682-AEC5-42C944756E12>

upvoted 1 times

🗳️ 👤 **TheOracleWasTaken** 1 year, 11 months ago

**Selected Answer: CD**

<https://www.examttopics.com/discussions/oracle/view/22606-exam-1z0-071-topic-2-question-60-discussion/>

upvoted 3 times

Examine the description of the EMPLOYEES table:

Name	Null?	Type
EMPLOYEE_ID	NOT NULL	NUMBER(4)
EMPLOYEE_NAME	NOT NULL	VARCHAR2(100)
SALARY	NOT NULL	NUMBER(6,2)
DEPARTMENT_ID	NOT NULL	NUMBER(4)

Which statement will fail?

A.

```
SELECT department_id, COUNT(*)
  FROM employees
 WHERE department_id <> 90
    AND COUNT(*) >= 3;
 GROUP BY department_id
```

B.

```
SELECT department_id, COUNT(*)
  FROM employees
 HAVING department_id <> 90
    AND COUNT(*) >= 3;
 GROUP BY department_id
```

C.

```
SELECT department_id, COUNT(*)
  FROM employees
 WHERE department_id <> 90
 HAVING COUNT(*) >= 3;
 GROUP BY department_id
```

D.

```
SELECT department_id, COUNT(*)
  FROM employees
 WHERE department_id <> 90
 GROUP BY department_id
 HAVING COUNT(*) >= 3;
```

**Suggested Answer: A**

 **darkma**  1 year, 3 months ago

The first query is syntactically incorrect. You cannot use the COUNT function in the WHERE clause.

The second query is also incorrect. The HAVING clause is used to filter the results of an aggregation, and it should be used after the GROUP BY clause, not after the WHERE clause.

The third query is also incorrect for the same reason as the second one. The HAVING clause should come after the GROUP BY clause.

The fourth query is correct and will work as expected. It filters rows where the department\_id is not equal to 90 using the WHERE clause, groups the remaining rows by department\_id, and then applies the HAVING clause to count the rows within each group and filter out groups where the count is greater than or equal to 3.

upvoted 7 times

 **JayaprasanthGurunathan** 6 months, 1 week ago

D is the answer. This query is correct. The WHERE clause correctly filters rows before aggregation (i.e., excludes rows where department\_id = 90), and the HAVING clause correctly filters the grouped results (i.e., only returns groups where the count of rows is greater than or equal to 3).

Conclusion: This query will not fail.

upvoted 1 times

 **NSilva** 8 months ago

I tried it out in Oracle's Apex (Db v19C), using a data set from the internet.

SELECT id

FROM TITANIC  
HAVING ID = 1  
GROUP BY ID;

It does not affect the result if you use HAVING and GROUP BY in a different position from one to another.

upvoted 2 times

🗳️ 👤 **JayaprasanthGurunathan** Most Recent 5 months, 1 week ago

A will fail because of the use of COUNT(\*) in the WHERE clause. In SQL, aggregate functions like COUNT(\*) should be used in the HAVING clause, not the WHERE clause. The WHERE clause is applied before the grouping (aggregation), so you can't filter on the result of COUNT(\*) until after the data has been grouped.

upvoted 1 times

🗳️ 👤 **alelejaja** 7 months ago

A is the correct answer

upvoted 1 times

🗳️ 👤 **m\_darwish** 7 months, 1 week ago

Ignore semicolon , A will fail .

upvoted 1 times

🗳️ 👤 **MIGHTYSTUNNERS** 8 months ago

C is wrong Because HAVING Clause come under group by clause

upvoted 1 times

🗳️ 👤 **zouve** 1 year, 5 months ago

A

if you correct the semicolon to each and every query A the one that is failing with ORA-00934: group function is not allowed here

upvoted 2 times

🗳️ 👤 **fgramos** 1 year, 6 months ago

D is the right answer. Works fine and AB and C are incorrect by semicolon

upvoted 2 times

🗳️ 👤 **yanoolthecool** 1 year, 5 months ago

Agreed, that is if we assume the question is 'which statement will NOT fail'

upvoted 1 times

🗳️ 👤 **karols32** 1 year, 7 months ago

Don't pay attention about semicolon and the answer is A.

Order of HAVING and GROUP BY doesn't matter.

WHERE is not required in statement.

upvoted 2 times

🗳️ 👤 **Orxan\_H** 1 year, 8 months ago

A B are incorrect answers

upvoted 1 times

🗳️ 👤 **Ronit\_8055** 1 year, 11 months ago

ABC are incorrect for semicolon and even for HAVING should come after GROUP BY clause.

upvoted 1 times

🗳️ 👤 **TheOracleWasTaken** 1 year, 4 months ago

I think it works regardless, probably not correct but it works

upvoted 1 times

🗳️ 👤 **JCKOMI** 2 years ago

A IS NOT correct for semicolon and also count is not applied to where clause

upvoted 1 times

🗳️ 👤 **MZEnuovazelandia** 2 years, 3 months ago

It's a mess! the first three are incorrect for the semicolon

upvoted 4 times



Examine the data in the NEW\_EMPLOYEES table:

EMPLOYEE_ID	NAME	DEPARTMENT_ID	MANAGER_ID	JOB_ID	SALARY
101	David	20	120	SA_REP	14000
102	Sam	10	105	CLERK	12500
103	Andrew	20	120	FIN_ADMIN	14200
104	Adrian	30	108	MAR_CLERK	12500
105	Maria	30	108	FIN_ADMIN	15000
106	Tracy	40	110	AD_ASST	13000
108	Kate	30	110	FIN_DIR	16500
110	Anne	40	120	EX_DIR	18000
120	Fran	20	110	SQ_DIR	16500

Examine the data in the EMPLOYEES table:

EMPLOYEE_ID	NAME	JOB_ID	SALARY
101	David	CLERK	14000
102	Sam	SA_REP	11500
104	Adrian	MAR_CLERK	12500
108	Kate	FIN_DIR	16500
110	Annie	EX_DIR	18000

You want to:

1. Update existing employee details in the EMPLOYEES table with data from the NEW\_EMPLOYEES table.
2. Add new employee details from the NEW\_EMPLOYEES table to the EMPLOYEES table.

Which statement will do this?

A.

```
MERGE INTO employees e
USING new_employees ne
  ON (e.employee_id = ne.employee_id)
  WHEN MATCHED THEN UPDATE SET e.name = ne.name, e.job_id = ne.job_id, e.salary =
ne.salary
  WHEN NOT MATCHED THEN INSERT VALUES (ne.employee_id, ne.name, ne.job_id, ne.salary);
```

B.

```
MERGE INTO employees e
USING new_employees ne
  ON (e.employee_id = ne.employee_id)
  WHEN FOUND THEN UPDATE SET e.name = ne.name, e.job_id = ne.job_id, e.salary =
ne.salary
  WHEN NOT FOUND THEN INSERT VALUES (ne.employee_id, ne.name, ne.job_id, ne.salary);
```

C.

```
MERGE INTO employees e
USING new_employees ne
WHERE e.employee_id = ne.employee_id
  WHEN MATCHED THEN UPDATE SET e.name = ne.name, e.job_id = ne.job_id, e.salary =
ne.salary
  WHEN NOT MATCHED THEN INSERT VALUES (ne.employee_id, ne.name, ne.job_id, ne.salary);
```

D.

```
MERGE INTO employees e
USING new_employees ne
WHERE e.employee_id = ne.employee_id
  WHEN FOUND THEN UPDATE SET e.name = ne.name, e.job_id = ne.job_id, e.salary =
ne.salary
  WHEN NOT FOUND THEN INSERT VALUES (ne.employee_id, ne.name, ne.job_id, ne.salary);
```

Suggested Answer: A

 **TheOracleWasTaken** Highly Voted 11 months ago

B is not correct because "when found then" isn't valid syntax

C doesn't have an on clause

D also has a "when found then" which isn't valid

upvoted 7 times



Examine the description of the EMPLOYEES table:

Name	Null?	Type
EMP_NO	NOT NULL	NUMBER(4)
LAST_NAME		VARCHAR2(10)
HIRE_DATE		DATE
SALARY		NUMBER(6,2)

For each employee in department 90 you want to display:

1. their last name
2. the number of complete weeks they have been employed

The output must be sorted by the number of weeks, starting with the longest serving employee first.

Which statement will accomplish this?

A.

```
SELECT last_name, ROUND((SYSDATE - hire_date) / 7) AS tenure
FROM employees
WHERE department_id = 90
ORDER BY tenure DESC;
```

B.

```
SELECT last_name, TRUNC((SYSDATE - hire_date) / 7) AS tenure
FROM employees
WHERE department_id = 90
ORDER BY tenure DESC;
```

C.

```
SELECT last_name, ROUND((SYSDATE - hire_date) / 7) AS tenure
FROM employees
WHERE department_id = 90
ORDER BY tenure;
```

D.

```
SELECT last_name, TRUNC((SYSDATE - hire_date) / 7) AS tenure
FROM employees
WHERE department_id = 90
ORDER BY tenure;
```

**Suggested Answer: C**

 **gabolka**  2 years, 9 months ago

its asking for complete weeks so its B

upvoted 33 times

 **holdfaststrongly**  2 years, 9 months ago

B.

Rounding up gives then employees extra days for their "complete" weeks. B would make more sense.

upvoted 11 times

 **826cb95**  2 months, 1 week ago

B, Truncate not round for complete week.

upvoted 2 times

 **highwater** 3 months ago

The correct answer is B.

C doesn't have sense

upvoted 1 times

 **AWSH** 7 months, 2 weeks ago

B. its asking for completed weeks (TRUNC not ROUND) and order by DESC.

upvoted 2 times

 **speedy\_vin** 8 months ago

Moderator, please change the correct answer to B. Currently it is C which is incorrect as the order needs to be descending.

upvoted 2 times

🗨️ 👤 **Psili** 9 months ago

There is no department in the table..

I also vote for B because of complete weeks, starting from longest period of time

upvoted 1 times

🗨️ 👤 **ArslanAltaf** 1 year ago

The answer should be changed to A or B because both are correct. By no means C is correct, it violates the last requirement LONGEST SERVING ON TOP.

upvoted 1 times

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

b is correct

upvoted 3 times

🗨️ 👤 **ArslanAltaf** 1 year, 2 months ago

A & B, Both are correct. More precise is B.

The requirement "output must be sorted by the number of weeks, starting with the longest" means the sorting order MUST be DESC on Tenure (High --> Low)

Since in C & D, the order by tenure will adopt ASC by default, it will be from low -> High so lowest tenure will be on top of output and highest tenure at bottom.

upvoted 1 times

🗨️ 👤 **Lee\_jong\_suk** 1 year, 4 months ago

B is the correct answer

upvoted 3 times

🗨️ 👤 **Arth1989** 1 year, 5 months ago

Are you kidding me?

Why paying for incorrect answers?

Obviously the answer is B!

upvoted 7 times

🗨️ 👤 **hmatinnn** 1 year, 6 months ago

Answer B. Complete weeks with trunc

upvoted 3 times

🗨️ 👤 **lucemqy** 1 year, 7 months ago

B would be more accurate since it does not give extra days by rounding

upvoted 2 times

🗨️ 👤 **DadiJoe** 1 year, 9 months ago

how to filter department = 90 when theres no Department or Dept\_ID in the table...

All options must be wrong.

upvoted 6 times

🗨️ 👤 **HassanMkhlalati** 1 year, 9 months ago

Answer is B, complete weeks.

upvoted 1 times

🗨️ 👤 **MariGare** 1 year, 10 months ago

I would say A. If we have 5 working days in a week, then Round will give us the correct result. And as we were asked to show employees, who serves longer, first, then we should order by desc. So, I do not understand why correct answer is C ...

upvoted 1 times

Examine the description of the PRODUCT\_DETAILS table:

Name	Null?	Type
PRODUCT_ID	NOT NULL	NUMBER(2)
PRODUCT_NAME	NOT NULL	VARCHAR2(25)
PRODUCT_PRICE		NUMBER(8,2)
EXPIRY_DATE		DATE

Which two statements are true? (Choose two.)

- A. EXPIRY\_DATE contains the SYSDATE by default if no date is assigned to it.
- B. PRODUCT\_PRICE can be used in an arithmetic expression even if it has no value stored in it.
- C. PRODUCT\_NAME cannot contain duplicate values.
- D. EXPIRY\_DATE cannot be used in arithmetic expressions.
- E. PRODUCT\_PRICE contains the value zero by default if no value is assigned to it.
- F. PRODUCT\_ID can be assigned the PRIMARY KEY constraint.

**Suggested Answer:** BF

Community vote distribution

BF (88%)

13%

826cb95 2 months, 1 week ago

**Selected Answer: BF**

B. PRODUCT\_PRICE can be used in an arithmetic expression even if it has no value stored in it.

Why? Because while the result is NULL, the expression itself is valid SQL and won't cause an error.

upvoted 1 times

braintop 7 months, 1 week ago

**Selected Answer: BD**

EXPIRY\_DATE cannot be used in ALL arithmetic operators

As there is no information on whether PRODUCT\_ID contains duplicate values, we cannot determine whether PK can be assigned on it.

upvoted 1 times

lucemqy 1 year, 7 months ago

**Selected Answer: BF**

You can add NULL to a number it will result in NULL

upvoted 2 times

HassanMkhlalati 1 year, 9 months ago

**Selected Answer: BF**

BF correct

upvoted 1 times

TheOracleWasTaken 1 year, 11 months ago

**Selected Answer: BD**

We dont know if there's any duplicate values for the id column so how do we know if it can be a PK

upvoted 1 times

NSilva 1 year, 2 months ago

D. EXPIRY\_DATE cannot be used in arithmetic expressions.

False. You can use Arithmetic expressions in Dates, just not all of them.

SELECT

DATE'2023-05-20' - 1 AS expire\_date,

DATE'2023-05-20' + 1 AS extra\_day

FROM DUAL;

I agree that the primary key could contain duplicated values, but is the one that seems most feasible out of the remaining options.

upvoted 1 times

🗨️ **yanoolthecool** 1 year, 11 months ago

**Selected Answer: BF**

B's correct, so is F

upvoted 1 times

🗨️ **CMjer** 2 years, 1 month ago

**Selected Answer: BF**

BF correct, see discussion on same question: <https://www.examttopics.com/discussions/oracle/view/8221-exam-1z0-071-topic-1-question-229-discussion/>

upvoted 1 times

🗨️ **Sbablo** 2 years, 1 month ago

B is wrong, absence of data means null. An arithmetic operation with a null value always yields null

upvoted 2 times

🗨️ **Marko\_1991** 2 years, 1 month ago

what about nvl?

upvoted 1 times

🗨️ **HassanMkhlalati** 1 year, 9 months ago

But it doesn't send an error message

upvoted 1 times

🗨️ **zouve** 1 year, 11 months ago

BF

PRODUCT\_PRICE can be used in an arithmetic expression even if it has no value stored in it, but the result will be NULL

upvoted 1 times

🗨️ **dexdinh91** 2 years, 5 months ago

**Selected Answer: BF**

BF are correct

upvoted 2 times

🗨️ **SantiBZ\_07032022\_1744** 2 years, 5 months ago

Expiry\_DATE cannot be used in arithmetic expressions -> FALSE.

upvoted 1 times

🗨️ **KizTine** 1 year, 8 months ago

yes it's possible. For example :

SELECT expiry\_date + 3 FROM product\_details; (to add 3 days)

SELECT expiry\_date + interval '1' month FROM product\_details; (to add 1 month)

upvoted 1 times

🗨️ **JCKOMI** 2 years, 6 months ago

why expiry date can not be used in arithmetic expressions ?

upvoted 1 times

🗨️ **karols32** 2 years, 2 months ago

EXPIRY\_DATE cannot be used in arithmetic expressions. - FALSE

Because on date data type you can use a "MINUS" arithmetic expressions.

upvoted 1 times

🗨️ **TheOracleWasTaken** 1 year, 11 months ago

But can you multiply a date or divide it?

upvoted 1 times

Examine the description of the EMPLOYEES table:

Name	Null?	Type
EMPLOYEE_ID	NOT NULL	NUMBER(3)
FIRST_NAME		VARCHAR2(15)
LAST_NAME	NOT NULL	VARCHAR2(15)
SALARY		NUMBER(6,2)

Which two queries will result in an error? (Choose two.)

A.

```
SELECT first_name last_name
FROM employees;
```

B.

```
SELECT first_name, last name
FROM employees;
```

C.

```
SELECT last_name, 12 * salary AS annual_salary
FROM employees
WHERE annual_salary > 100000
ORDER BY 12 * salary;
```

D.

```
SELECT last_name, 12 * salary AS annual_salary
FROM employees
WHERE 12 * salary > 100000
ORDER BY 12 * salary;
```

E.

```
SELECT last_name, 12 * salary AS annual_salary
FROM employees
WHERE annual_salary > 100000
ORDER BY annual_salary;
```

F.

```
SELECT last_name, 12 * salary AS annual_salary
FROM employees
WHERE 12 * salary > 100000
ORDER BY annual_salary;
```

Suggested Answer: CE

 **williamsmith95** Highly Voted 7 months ago

Correct answer: C and E

upvoted 13 times

 **Kashka** Highly Voted 2 years, 3 months ago

Cannot use aliases in WHERE clause:

ORA-00904: "ANNUAL\_SALARY": invalid identifier

upvoted 5 times

 **wojtekb** Most Recent 3 months ago

A - logical typo but still will work

B - Won't work because if \_ is missing oracle does not recognize column (invalid identifier)

C - Won't work because you can't use aliases in where clause

D - Works

E - The same as C

F - Works

Answers : A, C, E

Tested on Oracle live

upvoted 1 times

🗨️ 👤 **wojtekb** 3 months ago

\*Meant ofc BCE  
upvoted 2 times

🗨️ 👤 **TheOracleWasTaken** 1 year, 4 months ago

C and E because of the alias used in the where clause  
upvoted 4 times

🗨️ 👤 **zouve** 1 year, 5 months ago

AB -> typos  
CE -> fails  
upvoted 5 times

🗨️ 👤 **alelejaja** 7 months ago

A could be using last\_name as an alias too  
upvoted 2 times

🗨️ 👤 **J4vi** 1 year, 7 months ago

BCE fail. Tested all.  
upvoted 4 times

🗨️ 👤 **shotcom** 1 year, 8 months ago

a, b is typo error hence answer is CE because we cant use column aliase in the where clause  
Answer CE  
upvoted 2 times

🗨️ 👤 **alelejaja** 7 months ago

A could be using last\_name as an alias of first\_name  
upvoted 2 times

🗨️ 👤 **ninjax\_m** 1 year, 10 months ago

a,b,c,e will fail. Pick 2 :)  
upvoted 2 times

🗨️ 👤 **alelejaja** 7 months ago

A doesn't fail  
upvoted 2 times

🗨️ 👤 **NajiJabr** 1 year, 11 months ago

A -- no comma separator between column names

B- wrong column name last name it should be Last\_Name  
upvoted 2 times

🗨️ 👤 **cadcadley** 1 year, 11 months ago

CE is correct  
upvoted 1 times

🗨️ 👤 **sqlwilly** 2 years ago

what about B. select first\_name, last name  
from employees; - Will result in ORA-00904 invalid Identifier  
upvoted 4 times

You create a table named 123.  
Which statement runs successfully?

- A. SELECT \* FROM TABLE(123);
- B. SELECT \* FROM "123";
- C. SELECT \* FROM '123';
- D. SELECT \* FROM '123';

**Suggested Answer: B**

*Community vote distribution*

B (100%)

🗲️ 👤 **Monesha\_Ganesan** 1 year ago

**Selected Answer: B**

You have to use double quotes to refer to the table names  
upvoted 2 times

🗲️ 👤 **lucemqy** 1 year, 1 month ago

**Selected Answer: B**

You have to use double quotes to refer to the table names since you can't create a table named 123 using regular means  
upvoted 1 times

🗲️ 👤 **Tyler4747** 1 year, 1 month ago

Why wont Option D work when it works for me.  
upvoted 1 times

🗲️ 👤 **Roneros** 1 year, 3 months ago

**Selected Answer: B**

Object Quoted Names.  
<https://docs.oracle.com/en/database/oracle/oracle-database/21/sqlrf/Database-Object-Names-and-Qualifiers.html#GUID-75337742-67FD-4EC0-985F-741C93D918DA>  
upvoted 1 times

🗲️ 👤 **Izzicertificacion** 1 year, 4 months ago


las 4 me salen ERROR  
upvoted 1 times

Which two statements are true regarding indexes? (Choose two.)

- A. An update to a table can result in updates to any or all of the table's indexes.
- B. An update to a table can result in no updates to any of the table's indexes.
- C. A UNIQUE index can be altered to be non-unique.
- D. When a table is dropped and is moved to the RECYCLE BIN, all indexes built on that table are permanently dropped.
- E. A table belonging to one user cannot have an index that belongs to a different user.

**Suggested Answer: BD**

Community vote distribution



**holdfaststrongly** Highly Voted 2 years, 9 months ago

A and B can co-exist. Answer should be ABD.

upvoted 6 times

**Kuraudio** Highly Voted 2 years, 6 months ago

D is wrong: [https://docs.oracle.com/cd/B19306\\_01/server.102/b14200/statements\\_9012.htm](https://docs.oracle.com/cd/B19306_01/server.102/b14200/statements_9012.htm)

When you drop a table, any indexes on the table are dropped and put into the recycle bin along with the table. If subsequent space pressures arise, then the database reclaims space from the recycle bin by first purging indexes. In this case, when you flash back the table, you may not get back all of the indexes that were defined on the table.

upvoted 6 times

**Thameur01** Most Recent 5 months, 1 week ago

Selected Answer: AB

A,B and D

upvoted 1 times

**Psili** 9 months ago

ChatGPT find E also correct:

E. A table belonging to one user cannot have an index that belongs to a different user.

Indexes must belong to the same schema as the table. This means that if a table is owned by one user (in one schema), an index on that table must be created within the same schema. The system won't allow you to create an index on a table that belongs to one user while placing the index under another user's schema.

upvoted 1 times

**alelejaja** 1 year, 1 month ago

Selected Answer: AB

D can not be correct. Indexes are sent to the recycle bin as well

upvoted 1 times

**lucemqy** 1 year, 7 months ago

D is definitely not correct since index are moved to recycle bin not permanently removed

upvoted 2 times

**KizTine** 1 year, 8 months ago

Selected Answer: AB

D is wrong.

When a table is moved to the recycle bin, indexes on that table are also moved to the recycle bin and can be restored at the same time as the table.

upvoted 1 times

**HassanMkhlalati** 1 year, 9 months ago

Selected Answer: AB

AB exactly

upvoted 1 times



🗨️ 👤 **TheOracleWasTaken** 1 year, 11 months ago

**Selected Answer: AB**

a and b. I tried E out and you can create an index on someone else's table  
upvoted 3 times

🗨️ 👤 **zouve** 2 years ago

**Selected Answer: AB**

for me ABE  
upvoted 2 times

🗨️ 👤 **Poke90** 2 years, 3 months ago

**Selected Answer: AB**

D is wrong, indexes can be restored  
upvoted 3 times

🗨️ 👤 **Eltelwany** 2 years, 3 months ago

A is correct as well, it should be ABD.  
upvoted 1 times

🗨️ 👤 **Eltelwany** 2 years, 3 months ago

I also think D is not accurate. Because when a table is dropped and moved to the recycle bin (if the recycle bin is enabled), any associated indexes on that table are also moved to the recycle bin.

The recycle bin in Oracle holds dropped objects for a period of time, during which they can be restored if needed. However, the objects in the recycle bin are not permanently dropped until the bin is purged. If an index is dropped along with a table and is moved to the recycle bin, it is not permanently dropped until the recycle bin is purged.

upvoted 3 times

🗨️ 👤 **dexdinh91** 2 years, 5 months ago

**Selected Answer: BD**

BD are correct  
upvoted 1 times

🗨️ 👤 **Averagenoob** 2 years, 10 months ago

ABD it should be  
upvoted 1 times

🗨️ 👤 **Kashka** 2 years, 9 months ago

Hi Averagenoob.

You can flashback the table and restore the indexes, it just will be a different name, will it not?

upvoted 5 times

Which two are true about queries using set operators (UNION, UNION ALL, INTERSECT and MINUS)? (Choose two.)


- A. The name of each column in the first SELECT list must match the name of the corresponding column in each subsequent SELECT list.
- B. None of the set operators can be used when selecting CLOB columns.
- C. There must be an equal number of columns in each SELECT list.
- D. Each SELECT statement in the query can have an ORDER BY clause.
- E. The FOR UPDATE clause cannot be specified.

**Suggested Answer:** CE

Reference:

[https://docs.oracle.com/cd/B19306\\_01/server.102/b14200/queries004.htm](https://docs.oracle.com/cd/B19306_01/server.102/b14200/queries004.htm)

Community vote distribution



**jm9999** Highly Voted 1 year, 9 months ago

**Selected Answer:** CE

Union all works with clob.

upvoted 6 times

**jfc1** Highly Voted 2 years, 5 months ago

**Selected Answer:** BE

Set operations cannot be performed on BLOB, CLOB, BFILE, VARRAY, or nested table columns.

BCE is the Answer

upvoted 6 times

**ArslanAltai** Most Recent 1 year, 2 months ago

B,C & E All correct

upvoted 3 times

**ShahedOdeh** 9 months, 2 weeks ago

strongly agree

upvoted 2 times

**safasayed19** 8 months, 2 weeks ago

YES..B C AND E ALL ARE CORRECT

upvoted 1 times

**TATSSSS** 1 year, 3 months ago

CB

The set operators are not valid on columns of type BLOB, CLOB, BFILE, VARRAY, or nested table.

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/The-UNION-ALL-INTERSECT-MINUS-Operators.html#GUID-B64FE747-586E-4513-945F-80CB197125EE>

upvoted 1 times

**hmatinnn** 1 year, 5 months ago

**Selected Answer:** CE

Should be C and E. Tried it out !

upvoted 3 times

**Sangiii** 1 year, 10 months ago

B: this option is wrong.

we can use UNION ALL with 2 tables having columns with CLOB datatype. But we cannot use UNION, INTERSECT and MINUS SET operators as it performs removal of duplicate and sorting implicitly and could not work with CLOB.

upvoted 1 times

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

Hi, but in official doc from oracle, i found:

The set operators are subject to the following restrictions:

The set operators are not valid on columns of type BLOB, CLOB, BFILE, VARRAY, or nested table.

The UNION, INTERSECT, EXCEPT, and MINUS operators are not valid on LONG columns.

If the select list preceding the set operator contains an expression, then you must provide a column alias for the expression in order to refer to it in the order\_by\_clause.

You cannot also specify the for\_update\_clause with the set operators.

You cannot specify the order\_by\_clause in the subquery of these operators.

You cannot use these operators in SELECT statements containing TABLE collection expressions.

upvoted 3 times

🗨️ 👤 **MariGare** 1 year, 10 months ago

B should be included into the TRUE answer:

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/The-UNION-ALL-INTERSECT-MINUS-Operators.html#GUID-B64FE747-586E-4513-945F-80CB197125EE>

The set operators are subject to the following restrictions:

The set operators are not valid on columns of type BLOB, CLOB, BFILE, VARRAY, or nested table.

The UNION, INTERSECT, and MINUS operators are not valid on LONG columns.

If the select list preceding the set operator contains an expression, then you must provide a column alias for the expression in order to refer to it in the order\_by\_clause.

You cannot also specify the for\_update\_clause with the set operators.

You cannot specify the order\_by\_clause in the subquery of these operators.

You cannot use these operators in SELECT statements containing TABLE collection expressions.

upvoted 1 times

🗨️ 👤 **Shrimathi** 1 year, 11 months ago

CE is the answer

None of the set operator can be used with CLOB -False

The set operator which is excluding duplicate rows (intersect,minus,union) will fail to compare the LOB data type.

Union all operator doesn't throw any error if it contains LOB datatype

upvoted 2 times

🗨️ 👤 **johannfurma** 1 year, 6 months ago

The docu states:

"The set operators are not valid on columns of type BLOB, CLOB, BFILE, VARRAY, or nested table."

upvoted 2 times

🗨️ 👤 **gyanji** 2 years, 7 months ago

<https://docs.oracle.com/database/121/SQLRF/queries004.htm#SQLRF52341>

upvoted 1 times

🗨️ 👤 **holdfaststrongly** 2 years, 9 months ago

BCE per link.

upvoted 2 times

🗨️ 👤 **Averagenoob** 2 years, 9 months ago

BCE it is for sure

upvoted 2 times

🗨️ 👤 **Artur1991** 2 years, 9 months ago

Checking provided link, valid are BCE

upvoted 1 times

🗨️ 👤 **nhquoc1990** 2 years, 9 months ago

Union All can use with BLOB. checked

upvoted 5 times

BOOK\_SEQ is an existing sequence in your schema.

Which two CREATE TABLE commands are valid? (Choose two.)

A.

```
CREATE TABLE bookings (
  bk_id      NUMBER(4) DEFAULT book_seq.NEXTVAL PRIMARY KEY,
  start_date DATE      DEFAULT SYSDATE,
  end_date   DATE      DEFAULT SYSDATE NOT NULL);
```

B.

```
CREATE TABLE bookings (
  bk_id      NUMBER(4)
  start_date DATE      DEFAULT SYSDATE,
  end_date   DATE      DEFAULT (end_date >= start_date);
```

C. C.

```
CREATE TABLE bookings (
  bk_id      NUMBER(4) NOT NULL DEFAULT book_seq.CURRVAL,
  start_date DATE      NOT NULL,
  end_date   DATE      DEFAULT SYSDATE);
```


D.

```
CREATE TABLE bookings (
  bk_id      NUMBER(4) NOT NULL PRIMARY KEY,
  start_date DATE      NOT NULL,
  end_date   DATE      DEFAULT SYSDATE);
```

E.

```
CREATE TABLE bookings (
  bk_id      NUMBER(4) DEFAULT book_seq.CURRVAL,
  start_date DATE      DEFAULT SYSDATE,
  end_date   DATE      DEFAULT start_date);
```

**Suggested Answer: AD**

 **ArslanAltaf** 8 months, 2 weeks ago

for Book.seq currval- is there a logical error despite the statement does execute?

upvoted 1 times

 **hmatinnn** 10 months, 3 weeks ago

How not null and primary key can be set on the same column as primary key = not null + unique key ? I checked it is possible but technically it should not work in that way

upvoted 1 times

 **boddoju** 2 years ago

Can any one explain A,D are the correct answers

upvoted 1 times

 **anzac** 1 year, 12 months ago

A is correct because it have valid order of statements:

DEFAULT \_\_\_\_ NOT NULL

From this reason C is not valid because it have

NOT NULL DEFAULT \_\_\_\_ which is wrong order.

B and E are not valid because the are mentioning other columns in default statement.

upvoted 6 times

 **holdfaststrongly** 2 years, 3 months ago

CURRVAL might have a null value so option C may not work.

upvoted 1 times

 **karols32** 2 years, 1 month ago



so what not an option E?

upvoted 1 times

 **karols32** 2 years, 1 month ago

nvm, column cannot be a default value in next column.

upvoted 3 times

  **Orxan\_H** 1 year, 8 months ago  
in which situations curval may be Null?  
upvoted 1 times

Which three statements are true about multiple row subqueries? (Choose three.)



- A. Two or more values are always returned from the subquery.
- B. They can contain HAVING clauses.
- C. They can contain GROUP BY clauses.
- D. They can return multiple columns.
- E. They cannot contain a subquery.

**Suggested Answer:** BCD

Community vote distribution

BCD (67%)

ABC (33%)

  **bfb7c7d** 9 months ago

multiple row subqueries return multiple row with one column . then how D could be the answer.



upvoted 1 times

  **ArslanAltaf** 1 year, 2 months ago

A is wrong due to word ALWAYS. it is not necessary multiple row sub query always return multiple rows, it can return one row also depends on DB.

BCD is correct

upvoted 4 times



  **CMjer** 2 years, 1 month ago

**Selected Answer:** BCD

BCD,

<https://www.w3resource.com/sql/subqueries/multiple-row-column-subqueries.php>



upvoted 2 times

  **astride** 2 years, 2 months ago

**Selected Answer:** ABC

A Must to be also true because it severals rows which are values

upvoted 2 times

  **kaplanek** 2 years, 2 months ago

keyword ALWAYS makes it untrue, subquery can return 1 value

upvoted 7 times

  **hoangHai\_Pro** 1 year, 3 months ago

if 1 values (1 row) => single row subqueries

upvoted 1 times

Which three actions can you perform on an existing table containing data? (Choose three.)

- A. Increase the width of a numeric column.
- B. Add a new column as the table's first column.
- C. Define a default value that is automatically inserted into a column containing nulls.
- D. Change a DATE column containing data to a NUMBER data type.
- E. Change the default value of a column.
- F. Add a new NOT NULL column with a DEFAULT value.

**Suggested Answer:** AEF

Community vote distribution

AEF (67%)

ADE (33%)

🗳️ 👤 **ArslanAltaf** 1 year ago

AEF is 100% correct.  
upvoted 1 times

🗳️ 👤 **yaya32** 1 year, 5 months ago

**Selected Answer: AEF**

AEF for me  
upvoted 2 times

🗳️ 👤 **yanoolthecool** 1 year, 6 months ago

**Selected Answer: ADE**

why not ADE?  
will E work? when we add a new column with default value, will it use that default value automatically for all the existing rows? im not sure, I know that it will do so for new ones though.  
as for D, shouldnt it be possible? as dates are stored as numeric value in the DB from what I know..  
upvoted 1 times

🗳️ 👤 **shotcom** 2 years, 2 months ago

AEF  
B is wrong an additional column will appear at the end  
C is wrong default value will affect rows that will be inserted after creating default constraint  
upvoted 4 times

Which two statements are true about selecting related rows from two tables based on an Entity Relationship Diagram (ERD)? (Choose two.)

- A. Rows from unrelated tables cannot be joined.
- B. Relating data from a table with data from the same table is implemented with a self join.
- C. Implementing a relationship between two tables might require joining additional tables.
- D. Every relationship between the two tables must be implemented in a join condition.
- E. An inner join relates rows within the same table.

**Suggested Answer:** BC

🗨️ 👤 **rani\_12** 10 months, 3 weeks ago

why D is not correct ?

upvoted 1 times

🗨️ 👤 **Patrick\_Alain** 1 year, 10 months ago

I thought that the option "A" for this question would apply the same reasoning as the question 50, where the option that said "PRODUCT\_PRICE can be used in an arithmetic expression even if it has no value stored in it." was justified by saying "it doesn't make sense, it's not advisable, but it can be done".

upvoted 3 times

🗨️ 👤 **Sagarkpvs** 2 years, 3 months ago

Inner join can be done on different tables . Same table is called self join

upvoted 4 times

🗨️ 👤 **Poke90** 2 years, 3 months ago

How is E wrong?

upvoted 1 times

🗨️ 👤 **karols32** 2 years, 2 months ago

You can join tables in where. `SELECT * FROM FIRST_TABLE F, SECOND_TABLE S WHERE F.ID = S.ID`. This example uses a where clause, not a join.

upvoted 1 times



Which three statements about roles are true? (Choose three.)

- A. Roles are assigned to users using the ALTER USER statement.
- B. Privileges are assigned to a role using the GRANT statement.
- C. A role is a named group of related privileges that can only be assigned to a user.
- D. A single user can be assigned multiple roles.
- E. Privileges are assigned to a role using the ALTER ROLE statement.
- F. Roles are assigned to roles using the ALTER ROLE statement.
- G. A single role can be assigned to multiple users.

**Suggested Answer:** BDG

*Community vote distribution*

BDG (100%)

🗉 👤 **TheOracleWasTaken** 11 months ago

**Selected Answer:** BDG

Not C. Look out for \*only\*  
upvoted 3 times

🗉 👤 **TheOracleWasTaken** 11 months ago

you can also assign roles to roles  
upvoted 1 times

🗉 👤 **Kuraudio** 1 year, 6 months ago

**Selected Answer:** BDG

BDG is correct  
upvoted 1 times

🗉 👤 **Artur1991** 1 year, 9 months ago

BDG are correct  
upvoted 3 times


The INVOICE table has a QTY\_SOLD column of data type NUMBER and an INVOICE\_DATE column of data type DATE. NLS\_DATE\_FORMAT is set to DD-MON-RR.

Which two are true about data type conversions involving these columns in query expressions? (Choose two.)

- A. invoice\_date = '15-march-2019' : uses implicit conversion
- B. qty\_sold BETWEEN '101' AND '110' : uses implicit conversion
- C. invoice\_date > '01-02-2019' : uses implicit conversion
- D. qty\_sold = '0554982' : requires explicit conversion
- E. CONCAT (qty\_sold, invoice\_date) : requires explicit conversion

**Suggested Answer: AB**

*Community vote distribution*



**Deni61** 1 week, 6 days ago

**Selected Answer: BE**

It could not be A, because string doesn't match NLS format (RR not RRRR for year).

E is correct in my opinion because CONCAT can work only with strings, you need to use TO\_CHAR for both arguments (date and number).

upvoted 1 times

**ArslanAltaf** 1 year ago

ABC are correct

upvoted 3 times

**Bertrand\_dok** 1 year ago

**Selected Answer: BE**

A and C can't work because it doesn't match the NLS date format, D works fine by implicit conversion , so it's BE

upvoted 2 times

**yanoolthecool** 1 year ago

E says 'explicit' did you see any conversion function?

A works, as MON is for the full month name and RR.. just google it.

upvoted 1 times

**799f3be** 1 year, 2 months ago

why is E wrong?

upvoted 1 times

**a947739** 1 year ago

SELECT CONCAT(sal,hiredate) from emp; try this , it doesn't need explicit conversion

upvoted 1 times

**Nennuzzo** 1 year, 6 months ago

**Selected Answer: AB**

AB are correct, I tested it.

upvoted 3 times

**Monesha\_Ganesan** 1 year, 6 months ago

**Selected Answer: CD**



C. invoice\_date > '01-02-2019': This statement uses implicit conversion. The date string is implicitly converted to a DATE type using the default date format.

D. qty\_sold = '0554982': This statement requires explicit conversion. You need to explicitly convert the string to a NUMBER to compare it with the QTY\_SOLD column.

upvoted 1 times

**MoonLight** 2 years, 2 months ago

Can anyone please explain the correct answers? not quite getting the sense... thank you.  
upvoted 1 times

  **kuff** 2 years, 1 month ago

- A. Implicit conversion can be done as string matches the `NSL_DATE_FORMAT`.
- B. Implicit conversion to number
- C. Can't use implicit conversion as string doesn't match `NSL_DATE_FORMAT`
- D. Doesn't require explicit conversion
- E. Doesn't require explicit conversion

upvoted 5 times

  **lucemqy** 1 year, 7 months ago

but the month is spelled out isn't `MON` the three letter abbreviation of the month?

upvoted 2 times


Which three statements are true about inner and outer joins? (Choose three.)

- A. A full outer join returns matched and unmatched rows.
- B. Outer joins can be used when there are multiple join conditions on two tables.
- C. A full outer join must use Oracle syntax.
- D. Outer joins can only be used between two tables per query.
- E. A left or right outer join returns only unmatched rows.
- F. An inner join returns matched rows.

**Suggested Answer:** ABF

Community vote distribution

ABF (100%)

  **826cb95** 2 months, 1 week ago

**Selected Answer:** ABF

✓ B. Outer joins can be used when there are multiple join conditions on two tables.

→ True – You can absolutely have multiple conditions in an outer join using AND, and it works across columns.

A & F are obvious.

upvoted 1 times

  **lucemqy** 1 year, 1 month ago

**Selected Answer:** ABF

Full Outer join returns rows from both tables regardless of match. Inner join only return matched rows

upvoted 1 times

  **zouve** 1 year, 6 months ago

**Selected Answer:** ABF

From chagpt

A. Option A is true. A full outer join returns both matched rows and unmatched rows from the joined tables. Matched rows are those that satisfy the join condition, while unmatched rows are those that do not have a matching row in the other table.

B. Option B is true. Outer joins can be used when there are multiple join conditions on two tables. This allows for more complex join conditions and provides flexibility in joining tables based on multiple criteria.

F. Option F is true. An inner join returns only the matched rows from the joined tables. It filters out the unmatched rows, resulting in a result set that contains only the rows that satisfy the join condition.

upvoted 1 times

  **KizTine** 1 year, 2 months ago

I can't understand why answer B is right.

For example, the query below (with OUTER JOIN) works but don't use multiple join conditions :

```
SELECT A.col FROM tableA LEFT OUTER JOIN tableB ON tableA.col = tableB.col WHERE tableB.col IS NULL;
```

upvoted 2 times

  **TofOra** 10 months, 4 weeks ago

The question says "can" and not "must"

upvoted 1 times

  **Eltelwany** 1 year, 9 months ago



What exactly is the Oracle syntax here?

upvoted 1 times

  **TofOra** 10 months, 4 weeks ago

When you put the join condition in the WHERE instead of the JOIN

upvoted 1 times

  **MoreC** 1 year, 9 months ago

QUESTION

UNMATCHED COLUMN ARE THEY NOT PRODUCED BY LEFT AND RIGHT OUTER JOINS

upvoted 1 times

  **Poke90** 1 year, 9 months ago

E states that ONLY unmatched rows are found, and matched rows are not in the output. False

upvoted 1 times

Which statement will execute successfully?

A.

```
SELECT 1, 2 FROM DUAL
UNION
SELECT 3, 4 FROM DUAL
ORDER BY 1, 2;
```

B.

```
SELECT 1 FROM DUAL
UNION
SELECT 2 FROM DUAL
ORDER BY 1, 2;
```

C.

```
SELECT 3 FROM DUAL
UNION
SELECT 4 FROM DUAL
ORDER BY 3;
```

D.

```
SELECT 1, 2 FROM DUAL
UNION
SELECT 3, 4 FROM DUAL
ORDER BY 3, 4;
```

Suggested Answer: A

  **super04** 10 months, 1 week ago

1. The ORDER BY clause uses the column of the first SELECT query.
2. By default, the first column of the first SELECT query is used to sort the output in ascending order.
3. The ORDER BY clause accepts the column name or an alias.
4. You can use ORDER BY clause only once in a compound query.
5. The ORDER BY clause does not recognize the column names of the second SELECT query. To avoid confusion over column names, it is common practice to ORDER BY column positions.

upvoted 2 times

  **zouve** 1 year ago

A for sure

upvoted 1 times

  **Chandru007** 1 year, 6 months ago

@cspro2410

A) If you're running both select without order by means it will union and return only 2 columns and after giving order by 1,2 it will run successfully. This is because 1,2 represents two columns.

B) In here it will union and it return one column only but in order by its given 1,2 so will fail

C) Here it will union and it return one column but order by is happening on 3rd column

D) Here union and it return two columns but order by happening on 3rd,4th column

so a is correct

upvoted 4 times

  **cspro2410** 1 year, 6 months ago

why C is incorrect ?

upvoted 1 times

  **ManAtWar** 1 year, 5 months ago

C is not correct since the second element tries to "order by 3". This means that the parser is trying to order by the 3rd element in the select clause. However, there's only 2 elements.

upvoted 1 times

  **zouve** 11 months, 2 weeks ago

you mean 1 element

it is sorting only from the elements from the 1st select

upvoted 1 times

Examine the description of the EMPLOYEES table:

Name	Null?	Type
EMPLOYEE_ID	NOT NULL	NUMBER(6)
EMPLOYEE_NAME	NOT NULL	VARCHAR2(20)
SALARY	NOT NULL	NUMBER
DEPARTMENT_ID	NOT NULL	NUMBER(4)

Which two queries return all rows for employees whose salary is greater than the average salary in their department? (Choose two.)

A.

```
SELECT *
  FROM employees
 WHERE salary > AVG(salary) OVER (PARTITION BY department_id);
```

B.

```
SELECT *
  FROM employees e1
 WHERE salary > (
    SELECT AVG(salary)
      FROM employees e2
    WHERE e1.department_id = e2.department_id
  );
```

C.

```
SELECT *
  FROM employees
 WHERE salary > (
    SELECT AVG(salary)
      FROM employees
    GROUP BY department_id
  );
```

D.

```
SELECT *
  FROM employees
 WHERE salary > ANY (
    SELECT AVG(salary)
      FROM employees
    GROUP BY department_id
  );
```


E.

```
SELECT *
  FROM (
    SELECT e.*, AVG(salary) OVER (PARTITION BY department_id) avg_sal
      FROM employees e
  )
 WHERE salary > avg_sal;
```


Suggested Answer: BE

 **Thameur01** 5 months, 1 week ago

A,B and E by chatGpt  
upvoted 1 times

 **ArslanAltaf** 8 months, 2 weeks ago

B & D is correct only.  
E is not  
upvoted 1 times

 **ArslanAltaf** 8 months, 2 weeks ago


Correction - B, D and E all worked. All tested  
upvoted 1 times



  **ogi33** 8 months ago

D execute but gives no corect data

upvoted 5 times

  **Poke90** 1 year, 9 months ago

IN E, the column alias avg\_salary from the sub query is used in the WHERE clause of the main query. Is that allowed?

upvoted 2 times

  **Eltelwany** 1 year, 9 months ago

Yes, it is allowed. It follows from the order of the execution. Since the engine would execute the subquery first, it already knows what the alias refers to.

upvoted 3 times

Which three statements are true about the Oracle join and ANSI join syntax? (Choose three.)

- A. The Oracle join syntax supports creation of a Cartesian product of two tables.
- B. The Oracle join syntax only supports right outer joins.
- C. The SQL:1999 compliant ANSI join syntax supports creation of a Cartesian product of two tables.
- D. The Oracle join syntax performs less well than the SQL:1999 compliant ANSI join syntax.
- E. The Oracle join syntax supports natural joins.
- F. The Oracle join syntax performs better than the SQL:1999 compliant ANSI join syntax.
- G. The SQL:1999 compliant ANSI join syntax supports natural joins.

**Suggested Answer:** ACG

Community vote distribution

ACG (100%)

🗳️ 👤 **highwater** 3 months ago

**Selected Answer:** ACG

ACG are correct

upvoted 1 times

🗳️ 👤 **bca123** 5 months, 3 weeks ago

**Selected Answer:** ACF

HERE QUES 3 AND 64 ARE SAME CONCEPT, U PROVIDED answer for question 3 is oracle joins is better than ansi join, now you are saying different annwer

upvoted 1 times

🗳️ 👤 **osiedlowylis** 9 months, 3 weeks ago

ACG.

E is incorrect, because the natural join exists only in SQL:1999 (ANSI)

upvoted 1 times

🗳️ 👤 **Kuraudio** 1 year ago

**Selected Answer:** ACG

ACG is correct

upvoted 1 times

🗳️ 👤 **holdfaststrongly** 1 year, 3 months ago

E is also correct. The newer Oracle syntax does support Natural Join. <https://docs.oracle.com/javadb/10.8.3.0/ref/rrefsqljnaturaljoin.html>

upvoted 1 times

Which two are true about the NVL, NVL2, and COALESCE functions? (Choose two.)


- A. NVL must have expressions of the same data type.
- B. NVL can have any number of expressions in the list.
- C. NVL2 can have any number of expressions in the list.
- D. COALESCE stops evaluating the list of expressions when it finds the first non-null value.
- E. The first expression in NVL2 is never returned.
- F. COALESCE stops evaluating the list of expressions when it finds the first null value.

**Suggested Answer:** DE

Reference:

<https://www.interviewsansar.com/difference-between-nvl-nvl2-nullif-and-coalesce-functions/>

Community vote distribution



826cb95 2 months, 1 week ago

**Selected Answer: DE**

D. COALESCE(expr1, expr2, ..., exprN) returns the first non-null expression from the list and stops evaluating after that.

E. In NVL2(expr1, expr2, expr3), if expr1 is not null, then expr2 is returned, otherwise expr3. So, the first expression can influence the result, but isn't returned.

upvoted 1 times

Thameur01 5 months ago

**Selected Answer: DE**

A can perform implicit expression : wrong

D: First expression of NVL2 is never returned, it's only used for evaluation: true

upvoted 1 times

odi\_USER 5 months, 3 weeks ago

**Selected Answer: DE**

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/NVL.html>

upvoted 1 times

ArslanAltatf 8 months, 2 weeks ago

A is not necessarily be true. DB applies implicit conversation which sometimes will not work. <https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/NVL.html>

upvoted 1 times

lucemqy 1 year, 1 month ago

**Selected Answer: DE**

First expression of NVL2 is never returned, it's only used for evaluation.

upvoted 3 times

HassanMkhlalati 1 year, 1 month ago

**Selected Answer: DE**

A: NVL, does implicit conversion. Could have expressions of different datatype

upvoted 4 times

Sangiii 1 year, 4 months ago

A: NVL must have expressions of same datatype

my inputs- here in option used must have hence it is not right

SELECT NVL(ENAME, 0) FROM EMP; -- Here Ename is Varchar2 and we are passing 0 Number data type as second parameter. In this scenario, NVL working for 2 different datatypes.

upvoted 1 times

🗳️ 👤 **HassanMkhalati** 1 year, 1 month ago

Implicit conversion is always successful from number to varchar, not the opposite.

upvoted 1 times

🗳️ 👤 **anushamathew12071997** 1 year, 5 months ago

DE is correct the data types of the nvl are possible to explicitly convert the data type

upvoted 2 times

🗳️ 👤 **anushamathew12071997** 1 year, 5 months ago

D IS CORRECT

COALESCE () BECAUSE Return first not null expression in the expression list.

upvoted 1 times

🗳️ 👤 **anushamathew12071997** 1 year, 5 months ago

E IS CORRECT

The first expression in NVL2 is never returned, it is only used to determine whether expression2 must be returned, or expression3.

upvoted 1 times

🗳️ 👤 **zouve** 1 year, 6 months ago

**Selected Answer: AD**

E is not right because

NVL2 (expr1, expr2, expr3)

expr1 is the source value or expression that may contain null

expr2 is the value returned if expr1 is not null

expr3 is the value returned if expr1 is null

F is not right because

The COALESCE() function returns the first non-null value in a list.

upvoted 1 times

🗳️ 👤 **ESZ** 1 year, 6 months ago

**Selected Answer: AD**

AD is correct, if not then please correct me.

upvoted 1 times

🗳️ 👤 **CMjer** 1 year, 7 months ago

**Selected Answer: DE**

DE is correct:

D. COALESCE stops evaluating the list of expressions when it finds the first non-null value. TRUE

<https://www.oracletutorial.com/oracle-comparison-functions/oracle-coalesce/>

E. The first expression in NVL2 is never returned. TRUE

If expr1 is not null, then NVL2 returns expr2. If expr1 is null, then NVL2 returns expr3. (<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/NVL2.html>)

upvoted 2 times

🗳️ 👤 **ESZ** 1 year, 8 months ago

**Selected Answer: AD**

A, D is correct

upvoted 1 times

🗳️ 👤 **Eltelwany** 1 year, 9 months ago

It should be AD.

A: it is true that oracle can make implicit conversions but still the NVL function requires that the two exprs be of the same data type. Hence A is correct

E is wrong because if we specified that we want the same exp to be returned when it's not null it would return its first exp. Tried this in Toad and got 'h1'

select nvl2('h1', 'h1', 4) from dual;

upvoted 2 times

🗳️ 👤 **CyberP** 1 year, 7 months ago

nvl2(expression1,expression2,expression3)

so even expression2 = expression1 we will not say expression1 is returned

A: we can implicitly use different type of data with nvl

select nvl(1,'1') from dual; it will work so A is incorrect

D,E is correct.

upvoted 3 times

  **lorenzo87** 1 year, 11 months ago

it's also true the A, select nvl(1,'pippo') from dual returns ora-01722

upvoted 1 times

  **MorticiaAAddams** 1 year, 11 months ago

Yes, but some values can be impicit converted into anotherdata type like nvl(1,'1').

upvoted 3 times

Examine this statement:

```
SELECT 1 AS id, 'John' AS first_name
FROM DUAL
UNION
SELECT 1, 'John' AS name
FROM DUAL
ORDER BY 1;
```

What is returned upon execution?

- A. an error
- B. 2 rows
- C. 0 rows
- D. 1 row

**Suggested Answer:** D

Community vote distribution

D (100%)

acd3ea6 9 months, 1 week ago

Selected Answer: D

D because it is union  
upvoted 1 times

lucemqy 1 year, 7 months ago

Selected Answer: D

Union will get rid of duplicate so only 1 row remaining  
upvoted 1 times

HassanMkhlalati 1 year, 9 months ago

Selected Answer: D

It is not union all to return duplicated rows  
upvoted 1 times

zouve 2 years ago

D checked

select 1 as id, 'john' as first\_name from dual  
union  
select 1, 'john' as name from dual order by 1;  
upvoted 1 times

AzaRza 2 years, 5 months ago

Because both rows are the same, the UNION result gives distinct answers  
upvoted 3 times

boddoju 2 years, 6 months ago

Option-B is the correct choose  
upvoted 2 times

boddoju 2 years, 6 months ago

my mistake D is the correct answer  
upvoted 2 times

Chandru007 2 years, 6 months ago

yes union all only gives duplicate so it will be 2 rows  
but union will give 1 row eliminating the duplicates  
upvoted 3 times

Examine this statement:

```
SELECT 1 AS id, 'John' AS first_name, NULL AS commission
FROM DUAL
INTERSECT
SELECT 1, 'John', NULL
FROM DUAL
ORDER BY 3;
```

What is returned upon execution?

- A. an error
- B. 2 rows
- C. 0 rows
- D. 1 row

**Suggested Answer:** D

Community vote distribution

D (100%)

🗨️ **Massy** 2 weeks, 3 days ago

**Selected Answer: D**

Oracle documentation says:

"In set operations, two NULLs are considered equal when determining duplicates."

upvoted 1 times

🗨️ **826cb95** 2 months, 1 week ago

**Selected Answer: A**

Important SQL Rule:

In Oracle SQL, NULL is not equal to another NULL. That means:

sql

Copy

Edit

NULL = NULL → FALSE (or unknown, technically)

The INTERSECT operator returns only the distinct rows that are equal in both sets.

Since NULL ≠ NULL, even if the rows look the same, Oracle treats them as not equal thus ERROR!

upvoted 1 times

🗨️ **lucemqy** 1 year, 1 month ago

**Selected Answer: D**

INTERSECT will only keep rows that exists for both tables so only 1 row remaining

upvoted 1 times

🗨️ **metallars** 1 year, 2 months ago

**Selected Answer: D**

Checked

upvoted 1 times

🗨️ **zouve** 1 year, 6 months ago

D tested

select 1 as id, 'john' as first\_name, NULL as commission from dual

intersect

select 1, 'john' as name, null from dual order by 3;

upvoted 2 times

🗨️ **J4vi** 1 year, 7 months ago

**Selected Answer: D**

It returns 1 row, because the values are the same. I do tested as well  
upvoted 2 times

🗨️ 👤 **Rik92** 1 year, 10 months ago

tested on 19c, returns 0 rows.  
upvoted 1 times

🗨️ 👤 **DiDivision** 1 year, 10 months ago

it does return 1 rows, it is fast to test.  
upvoted 2 times

🗨️ 👤 **9ramone7** 8 months, 3 weeks ago

tested on livesql.oracle.com and it does produce 1 row  
upvoted 1 times



Which two statements execute successfully? (Choose two.)

A.

```
SELECT TO_DATE('2019-DEC-25 15:30', 'YYYY-MON-DD HH24:MI',  
'NLS_DATE_LANGUAGE = AMERICAN')  
FROM DUAL;
```

B.

```
SELECT TO_CHAR(TO_DATE('2019-DEC-25 03:30', 'YYYY-MON-DD  
HH12:MI'))  
FROM DUAL;
```

C.

```
SELECT TO_DATE(TO_CHAR('2019-DEC-25 03:30', 'YYYY-MON-DD  
HH12:MI'))  
FROM DUAL;
```



D.

```
SELECT TO_CHAR('2019-DEC-25 15:30', 'YYYY-MON-DD HH24:MI')  
FROM DUAL;
```

E.

```
SELECT TO_CHAR('2019-DEC-25 15:30', 'YYYY-MON-DD HH24:MI',  
'NLS_DATE_LANGUAGE = AMERICAN')  
FROM DUAL;
```

Suggested Answer: AB

  **ShahedOdeh** 9 months, 2 weeks ago



the answer is A and B tested

upvoted 1 times

  **DrunkenMonkey** 1 year, 8 months ago

how did u extract only text from this text image

upvoted 1 times

  **yaya32** 1 year, 5 months ago



<https://www.editpad.org/tool/extract-text-from-image>

upvoted 1 times

  **zouve** 2 years ago

AB tested

upvoted 2 times

  **Poke90** 2 years, 3 months ago

why is D wrong?

upvoted 1 times

  **Eltelwany** 2 years, 3 months ago

It gives this error ORA-01722: invalid number. Also, there is no point in converting this to char, It can not parse it correctly. This needs to be converted to date first.


upvoted 2 times

An Oracle Database session has an uncommitted transaction in progress which updated 5000 rows in a table. In which three situations does the transaction complete thereby committing the updates? (Choose three.)

- A. when a CREATE TABLE AS SELECT statement is issued in the same session but fails with a syntax error
- B. when a DBA issues a successful SHUTDOWN TRANSACTIONAL statement and the user then issues a COMMIT
- C. when the session logs out successfully
- D. when a CREATE INDEX statement is executed successfully in the same session
- E. when a DBA issues a successful SHUTDOWN IMMEDIATE statement and the user then issues a COMMIT
- F. when a COMMIT statement is issued by the same user from another session in the same database instance

**Suggested Answer:** ACD

Community vote distribution



🗳️ 👤 **HassanMkhlalati** 7 months, 3 weeks ago

**Selected Answer:** BCD

A is not successful DML  
upvoted 2 times

🗳️ 👤 **jm9999** 8 months, 3 weeks ago

**Selected Answer:** BCD

For A to be true the statement would have to be syntactically valid.  
upvoted 2 times

🗳️ 👤 **ramonciks** 9 months, 1 week ago

Should be BCD:  
A - FALSE

- it is syntax error therefore does not implicitly commit: Oracle Database issues an implicit COMMIT under the following circumstances:

Before any syntactically valid data definition language (DDL) statement, even if the statement results in an error

After any data definition language (DDL) statement that completes without an error <https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/COMMIT.html#GUID-6CD5C9A7-54B9-4FA2-BA3C-D6B4492B9EE2>

upvoted 1 times

🗳️ 👤 **HassanMkhlalati** 9 months, 2 weeks ago

**Selected Answer:** BCD

A is wrong, ddl is still in compilation mode, so no ddl considered.  
upvoted 1 times

🗳️ 👤 **Shrimathi** 11 months ago

A - false since it is syntax error (checked ) if it is run time error then it commits the transaction.

B is true The SHUTDOWN TRANSACTIONAL waits for all uncommitted transactions to complete before shutting down the database instance. This saves the work for all users without requesting them to log off.

BCD would be the right answer  
upvoted 1 times

🗳️ 👤 **TheOracleWasTaken** 11 months ago

A works. It commits before and after every ddl statement so A will implicitly commit despite failing due to a syntax error  
upvoted 1 times

🗳️ 👤 **yanoolthecool** 1 year, 1 month ago

**Selected Answer:** ACD

Chat GPT says A is correct and B is wrong.

as SHUTDOWN TRANSACTIONAL functionality is that it waits all transactions to end then it shuts down the database, but he says since it was

'successful' then that means that the database already was shut down before the user committed.

I think there is some logical problem with the answer.. its confusing, how is it successful but there was still a transaction in progress?

upvoted 1 times

🗨️ 👤 **Oratns** 1 year, 3 months ago

ALTER TABLE and CREATE INDEX , which do cause a commit.

Oracle performs an implicit commit after any SQL DDL (Data Definition Language) statement (even if this DDL statement fails).

This includes table / index creation. So no open transaction should exist.

upvoted 1 times

🗨️ 👤 **BARNABA23** 1 year, 5 months ago

I CHECKED making an experiment and A IS CORRECT answer.

Looks like just CREATE TABLE AS creates a COMMIT (error was placed later - I used unexisting table as a source)

upvoted 2 times

🗨️ 👤 **dexdinh91** 1 year, 5 months ago

Don't know why A is correct answer

upvoted 1 times

🗨️ 👤 **jr\_C** 1 year, 3 months ago

Because before any DDL instructions, there is an implicit COMMIT

upvoted 2 times

🗨️ 👤 **heoj10272** 1 year, 7 months ago

**Selected Answer: BCD**

BCD is correct.

upvoted 3 times

🗨️ 👤 **iuliana23** 1 year, 8 months ago

**Selected Answer: BCD**

B should be correct, because the session will be killed after the last transaction will be committed or the user will use rollback... A is incorrect because the statement failed with a syntax error

upvoted 2 times

🗨️ 👤 **Peter\_Chen** 1 year, 9 months ago

why is not B?

C used SQLPLUS is right but use sql developer ....Maybe not

upvoted 3 times


Which two are true about using constraints? (Choose two.)

- A. NOT NULL can be specified at the column and at the table level.
- B. A table can have only one PRIMARY KEY and one FOREIGN KEY constraint.
- C. A FOREIGN KEY column in a child table and the referenced PRIMARY KEY column in the parent table must have the same names.
- D. PRIMARY KEY and FOREIGN KEY constraints can be specified at the column and at the table level.
- E. A table can have multiple PRIMARY KEY and multiple FOREIGN KEY constraints.
- F. A table can have only one PRIMARY KEY but may have multiple FOREIGN KEY constraints.

**Suggested Answer:** DF

Community vote distribution

DF (100%)

 **lucemqy** 7 months, 2 weeks ago

**Selected Answer:** DF

Not NULL constraints can only be set at column level not table level. Each table have only have one primary key constraint with multiple foreign key constraints


upvoted 1 times

 **HassanMkhlalati** 7 months, 3 weeks ago

**Selected Answer:** DF

D and F are true


upvoted 1 times

 **Roneros** 9 months, 3 weeks ago

**Selected Answer:** DF


<https://docs.oracle.com/javadb/10.8.3.0/ref/rrefsqlj13590.html>

upvoted 2 times

 **Eltelwany** 1 year, 3 months ago


What does D mean?

upvoted 1 times

 **J4vi** 1 year, 1 month ago

Is the way you specify the constraints syntactically. There are two ways... by column or by table (at the end of the statement). PK and FK can be defined in both.

upvoted 2 times

 **leongkaon** 1 year, 6 months ago

why A not correct?

upvoted 2 times

 **Sidoogle** 1 year, 5 months ago

Its exception that not null constrain is not allowed to be defined at table level.

upvoted 3 times

Examine this statement:

```
CREATE TABLE orders
(serial_no    NUMBER UNIQUE,
order_id     NUMBER PRIMARY KEY,
order_date   DATE NOT NULL,
status       VARCHAR2(10) CHECK (status IN ('CREDIT', 'CASH')),
product_id   NUMBER REFERENCES products (product_id),
order_total  NUMBER);
```

On which two columns of the table will an index be created automatically? (Choose two.)

- A. ORDER\_ID
- B. ORDER\_TOTAL
- C. ORDER\_DATE
- D. PRODUCT\_ID
- E. STATUS
- F. SERIAL\_NO

**Suggested Answer:** AF

Community vote distribution

AF (100%)

🗳️ 👤 **Wadyba** 10 months, 4 weeks ago

Selected Answer: AF

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: AF**

UNIQUE and PRIMARY key constraints will have unique index created automatically

upvoted 3 times

Examine this partial query:

```
SELECT ch.channel_type, t.month, co.country_code, SUM(s.amount_sold) SALES
FROM sales s, times t, channels ch, countries co
WHERE s.time_id = t.time_id
AND s.country_id = co.country_id
AND s.channel_id = ch.channel_id
AND ch.channel_type IN ('Direct Sales', 'Internet')
AND t.month IN ('2000-09', '2000-10')
AND co.country_code IN ('GB', 'US')
```

Examine this output:

CHANNEL_TYPE	MONTH	CO	SALES
Internet	2000-09	GB	16569
Internet	2000-09	US	124224
Internet	2000-09		140793
Internet	2000-10	GB	14539
Internet	2000-10	US	137054
Internet	2000-10		151593
Internet			292387
Direct Sales	2000-09	GB	85223
Direct Sales	2000-09	US	638201
Direct Sales	2000-09		723424
Direct Sales	2000-10	GB	91925
Direct Sales	2000-10	US	682297
Direct Sales	2000-10		774222
Direct Sales			1497646

Which GROUP BY clause must be added so the query returns the results shown?

- A. GROUP BY ch.channel\_type, ROLLUP(t.month, co.country\_code);
- B. GROUP BY ch.channel\_type, t.month, ROLLUP(co.country\_code);
- C. GROUP BY CUBE(ch.channel\_type, t.month, co.country\_code);
- D. GROUP BY ch.channel\_type, t.month, co.country\_code;

**Suggested Answer: A**

Community vote distribution

A (75%)

D (25%)

 **acd3ea6** 9 months ago

**Selected Answer: A**

```
create table channels_order (
channel_type varchar2(50),
month date,
code varchar2(5),
sales number);
```

```
insert into channels_order values('internet','2009-09','GB',16569);
insert into channels_order values('internet','2009-09','US',124224);
insert into channels_order values('internet','2009-10','GB',14539);
insert into channels_order values('internet','2009-10','US',137054);
```

```
insert into channels_order values('direct sales','2009-09','GB',85223);
insert into channels_order values('direct sales','2009-09','US',638201);
insert into channels_order values('direct sales','2009-10','GB',91925);
insert into channels_order values('direct sales','2009-10','US',682297);
```

```
select channel_type,month,code,sum(sales) sums from channels_order ch
group by channel_type,rollup(month,code);
```

upvoted 1 times

🗨️ 👤 **Lee\_jong\_suk** 1 year, 4 months ago

D is correct, can't say about rollup

upvoted 1 times

🗨️ 👤 **yaya32** 1 year, 5 months ago

**Selected Answer: A**

A for me

upvoted 1 times

🗨️ 👤 **Mcromeo** 1 year, 6 months ago

**Selected Answer: D**

D is the correct answer

upvoted 1 times

🗨️ 👤 **Roneros** 1 year, 9 months ago

**Selected Answer: A**

Ok, out of curiosity I added the first group to see the result...  $16569+124224+14539+137054 = 292386$ . A difference of 1?? Why? Could it have been a sum of undisplayed decimal values?

upvoted 1 times

🗨️ 👤 **asdf123asdf** 1 year, 10 months ago

**Selected Answer: A**

As per my understanding from below link I think A is correct.

[https://www.oracletutorial.com/oracle-basics/oracle-](https://www.oracletutorial.com/oracle-basics/oracle-rollup/#:~:text=The%20ROLLUP%20works%20as%20follows%3A%20First%2C%20calculate,to%20left.%203%20Finally%2C%20calculate%20the%20grand)

[rollup/#:~:text=The%20ROLLUP%20works%20as%20follows%3A%20First%2C%20calculate,to%20left.%203%20Finally%2C%20calculate%20the%20grand](https://www.oracletutorial.com/oracle-basics/oracle-rollup/#:~:text=The%20ROLLUP%20works%20as%20follows%3A%20First%2C%20calculate,to%20left.%203%20Finally%2C%20calculate%20the%20grand)

upvoted 1 times

🗨️ 👤 **fgramos** 2 years ago

D is the correct answer

upvoted 2 times

Examine the description of the EMPLOYEES table:

Name	Null?	Type
EMPLOYEE_ID	NOT NULL	NUMBER (3)
FIRST_NAME		VARCHAR2 (15)
LAST_NAME	NOT NULL	VARCHAR2 (15)
SALARY		NUMBER (6, 2)

Which statement will execute successfully, returning distinct employees with non-null first names?


- A. SELECT first\_name, DISTINCT last\_name FROM employees WHERE first\_name <> NULL;
- B. SELECT first\_name, DISTINCT last\_name FROM employees WHERE first\_name IS NOT NULL;
- C. SELECT DISTINCT \* FROM employees WHERE first\_name IS NOT NULL;
- D. SELECT DISTINCT \* FROM employees WHERE first\_name <> NULL;

**Suggested Answer: C**

Community vote distribution

C (80%)


A (20%)

 **lucemqy** 7 months, 2 weeks ago

**Selected Answer: C**

Not valid to compare NULL to any value must use IS NOT NULL or IS NULL


upvoted 2 times

 **metallars** 8 months, 1 week ago

**Selected Answer: C**

C is OK, last comment was a mistake.


upvoted 2 times

 **metallars** 8 months, 1 week ago

**Selected Answer: A**

AD are the correct answers

upvoted 1 times

 **Mcromeo** 6 months, 3 weeks ago

do you think you can get rows with this expression first\_name <> NULL ?

upvoted 1 times

 **CyberP** 1 year, 1 month ago

Distinct should come first after select, also we use is not null

so it is C.

upvoted 2 times



Examine the description of the BRICKS table:

Name	Null?	Type
BRICK_ID		NUMBER(38)
SHAPE		VARCHAR2(30)
COLOR		VARCHAR2(30)
WEIGHT		NUMBER

Examine the description of the BRICKS\_STAGE table:

Name	Null?	Type
WEIGHT		NUMBER
SHAPE		VARCHAR2(30)
COLOR		VARCHAR2(30)

Which two queries execute successfully? (Choose two.)

A.

```
SELECT brick_id, shape FROM bricks
MINUS
SELECT weight, color FROM bricks_stage;
```

B.

```
SELECT * FROM bricks
MINUS
SELECT * FROM bricks_stage;
```

C.

```
SELECT shape, color FROM bricks
MINUS
SELECT weight, color FROM bricks_stage;
```

D.

```
SELECT shape, color FROM bricks
MINUS
SELECT color, shape FROM bricks_stage;
```

E.

```
SELECT shape, color, weight FROM bricks
MINUS
SELECT * FROM bricks_stage;
```

**Suggested Answer: BE**

 **Artur1991**  2 years, 9 months ago

AD are the correct answers  
upvoted 29 times

 **paddy95**  2 years, 7 months ago

A D

```
create table briks ( brick_id number(38)
, shape varchar2(30)
, color varchar2(30)
, weight number );
```

```
create table bricks_stage( weight number
, shape varchar2(30)
, color varchar2(30));
```

--A +

```
select brick_id , shape from briks
minus
select weight, color from bricks_stage;
```

--B -

```

select * from briks
minus
select * from bricks_stage;

-- C -
select shape, color from briks
minus
select weight, color from bricks_stage;



-- D +
select shape, color from briks
minus
select color, shape from bricks_stage;



```



```



--E
select shape, color, weight from briks
minus
select * from bricks_stage;
upvoted 7 times



```




 **zouve** 2 years ago  
 tested twice AD  
 upvoted 6 times




 **Thameur01** Most Recent 5 months, 1 week ago  
 A and D are the correct options  
 upvoted 1 times




 **mage2611** 12 months ago  
 Answer is A, D is not correct because of the ordering of the columns (datatype mismatch)  
 upvoted 1 times




 **Ayman\_Khalifa** 11 months, 2 weeks ago  
 shape and color both have the same datatype, therefore, ordering doesn't matter  
 upvoted 1 times




 **ArslanAltaf** 1 year, 2 months ago  
 AD is correct. The Data type and number of columns have to be same.  
 upvoted 1 times




 **yaya32** 1 year, 5 months ago  
 AD for me  
 upvoted 2 times




 **Arth1989** 1 year, 5 months ago  
 AD, number of columns and order types should match  
 upvoted 2 times


 **yanoolthecool** 1 year, 6 months ago  
 if D works, why wont E?  
 upvoted 1 times


 **yanoolthecool** 1 year, 6 months ago  
 oh nvm, cuz D have varchar and varchar for both, so the order wont matter. E's wrong cuz it will match a number with varchar, so yea..  
 upvoted 1 times


 **Raja29** 1 year, 7 months ago  
 AD are the correct answers  
 upvoted 1 times


 **lucemqy** 1 year, 7 months ago  
 AD should be the correct answer since the data types match with same number of columns  
 upvoted 1 times


 **metallars** 1 year, 8 months ago

AD are the correct answers

upvoted 1 times

🗨️ 👤 **HassanMkhlalati** 1 year, 9 months ago

AD are the correct

upvoted 2 times

🗨️ 👤 **Arth1989** 1 year, 10 months ago

It is no brainer. AD!

upvoted 2 times

🗨️ 👤 **super04** 1 year, 10 months ago

AD is Correct answer.

Use Minus operator to return all distinct rows selected by the first query, but not present in the second query result set.

In Minus operator, the number of columns must be the same and data type of columns being selected by the SELECT statements in queries must belong to the same data type group in all the SELECT statements used in the query. The names of the columns, however, need not be identical.

upvoted 2 times

🗨️ 👤 **Sushmasri** 1 year, 11 months ago

A,D are correct

upvoted 3 times

🗨️ 👤 **Nelita** 2 years ago

AD ARE CORRECT. The others show errors.

upvoted 1 times

🗨️ 👤 **Orxan\_H** 2 years, 2 months ago

why C inccorect?

upvoted 2 times

🗨️ 👤 **fgramos** 2 years ago

Because Shape (1st column, 1st select) is VARCHAR2(20) and the first column of the second query is Weight, type number

upvoted 2 times

Table EMPLOYEES contains columns including EMPLOYEE\_ID, JOB\_ID and SALARY.

Only the EMPLOYEE\_ID column is indexed.

Rows exist for employees 100 and 200.

Examine this statement:

```
UPDATE employees
  SET (job_id, salary) =
      (SELECT job_id, salary
       FROM employees
        WHERE employee_id = 200)
 WHERE employee_id = 100;
```


Which two statements are true? (Choose two.)

- A. Employees 100 and 200 will have the same SALARY as before the update command.
- B. Employee 100 will have JOB\_ID set to the same value as the JOB\_ID of employee 200.
- C. Employee 200 will have JOB\_ID set to the same value as the JOB\_ID of employee 100.
- D. Employees 100 and 200 will have the same JOB\_ID as before the update command.
- E. Employee 100 will have SALARY set to the same value as the SALARY of employee 200.
- F. Employee 200 will have SALARY set to the same value as the SALARY of employee 100.

**Suggested Answer:** BE

Community vote distribution

BE (100%)

 **lucemqy** 7 months, 2 weeks ago

**Selected Answer:** BE

Inner query is run first then update query is ran  
upvoted 1 times

 **HassanMkhlalati** 9 months, 2 weeks ago

**Selected Answer:** BE

BE Correct  
upvoted 1 times

 **Nelita** 1 year ago

**Selected Answer:** BE

BE ARE CORRECT.  
upvoted 1 times

Examine these two queries and their output:

```
SELECT deptno, dname FROM dept;
```

DEPTNO	DNAME
10	ACCOUNTING
20	RESEARCH
30	SALES
40	OPERATIONS

```
SELECT ename, job, deptno FROM emp ORDER BY deptno;
```

ENAME	JOB	DEPTNO
CLARK	MANAGER	10
KING	PRESIDENT	10
MILLER	CLERK	10
JONES	MANAGER	20
FORD	ANALYST	20
ADAMS	CLERK	20
SMITH	CLERK	20
SCOTT	ANALYST	20
WARD	SALESMAN	30
TURNER	SALESMAN	30
ALLEN	SALESMAN	30
JAMES	CLERK	30
BLAKE	MANAGER	30
MARTIN	SALESMAN	30

Now examine this query:

```
SELECT ename, dname
FROM emp CROSS JOIN dept
WHERE job = 'MANAGER'
AND dept.deptno IN (10, 20);
```

How many rows will be displayed?

- A. 64
- B. 6
- C. 3
- D. 12

**Suggested Answer: B**

Community vote distribution

B (100%)

🗳️ **Artur1991** Highly Voted 2 years, 9 months ago

B is correct answer

From first table it will take 2 records, from second table it will take 3 records. After cross join we will have 6 records (2x3).

upvoted 6 times

🗳️ **Wadyba** 10 months, 4 weeks ago

but the 3rd manager has a deptno = 30, so I think it should be dropped.

upvoted 1 times

🗳️ **pukovnik** 10 months, 1 week ago

deptno in where clause does not relate to deptno column i emp table but dept table. So basically it's taking all managers from emp table(three of them), departments 10 and 20 from dept table and making cartesian product out of those.

upvoted 1 times

🗳️ **bca123** Most Recent 6 months, 1 week ago

**Selected Answer: A**

Ans is 4 cause each manager from from emp table (clark and jones) is matched with both dept table (10,20)--- $2 \times 2 = 4$   
upvoted 1 times

🗨️ 👤 **Massy** 2 weeks, 3 days ago  
what about blake?  
upvoted 1 times

🗨️ 👤 **Nelita** 2 years ago  
**Selected Answer: B**  
B IS CORRECT.  
upvoted 3 times

🗨️ 👤 **cadcadley** 2 years, 5 months ago  
B is correct because in emp table he take all manager rows then in dept he take 10 and 20  
will be  $3 \times 2 = 6$   
upvoted 2 times

🗨️ 👤 **dexdinh91** 2 years, 5 months ago  
**Selected Answer: B**  
 $3 \times 2 = 6$  so B  
upvoted 2 times

You want to return the current date and time from the user session, with a data type of `TIMESTAMP WITH TIME ZONE`.

Which function will do this?

- A. `SYSDATE`
- B. `CURRENT_TIMESTAMP`
- C. `LOCALTIMESTAMP`
- D. `CURRENT_DATE`

**Suggested Answer:** B

Reference:

[https://docs.oracle.com/cd/E11882\\_01/server.112/e10729/ch4datetime.htm#NLSPG004](https://docs.oracle.com/cd/E11882_01/server.112/e10729/ch4datetime.htm#NLSPG004)

*Community vote distribution*

B (100%)

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

**Selected Answer: B**

`CURRENT_TIMESTAMP` returns a timestamp data type with time zone

upvoted 1 times

🗨️ 👤 **Nelita** 1 year ago

**Selected Answer: B**

Correct is B

upvoted 1 times

🗨️ 👤 **CyberP** 1 year, 1 month ago

`LOCALTIMESTAMP` returns a `TIMESTAMP` value while `CURRENT_TIMESTAMP` returns a `TIMESTAMP WITH TIME ZONE` value

B is correct

upvoted 1 times

You have been tasked to create a table for a banking application.

One of the columns must meet three requirements:


- 1) Be stored in a format supporting date arithmetic without using conversion functions
- 2) Store a loan period of up to 10 years
- 3) Be used for calculating interest for the number of days the loan remains unpaid

Which data type should you use?

- A. INTERVAL YEAR TO MONTH
- B. TIMESTAMP WITH TIMEZONE
- C. INTERVAL DAY TO SECOND
- D. TIMESTAMP WITH LOCAL TIMEZONE
- E. TIMESTAMP

**Suggested Answer: C**

Community vote distribution



🗳️ 👤 **Thameur01** 6 months ago

**Selected Answer: C**

C. The data type of each column returned by the second query must be implicitly convertible to the data type of the corresponding column returned by the first query.

The data types must be implicitly convertible between corresponding columns. For example:

A NUMBER column in the first query can correspond to a VARCHAR column in the second query, provided Oracle can implicitly convert the VARCHAR to a NUMBER.

E. The number, but not names, of columns must be identical for all select statements in the query.

All queries in a UNION, INTERSECT, or MINUS statement must return the same number of columns, but the column names do not have to match.  
upvoted 1 times

🗳️ 👤 **kay000001** 7 months, 3 weeks ago

**Selected Answer: C**

C is correct.

upvoted 1 times

🗳️ 👤 **Misi\_Oracle** 1 year, 1 month ago

C is correct because it fulfill all the condition.

- 1) store without conversion
- 2) store up to 99 years and 11 month
- 3) can be use for calculation

upvoted 1 times

🗳️ 👤 **archit4321** 1 year, 2 months ago

**Selected Answer: A**

The INTERVAL DAY TO SECOND data type is used to store a period of time in terms of days, hours, minutes, and seconds. It does not support storing a loan period of up to 10 years, as it does not have a year component.

To meet the requirements of storing a loan period of up to 10 years and supporting date arithmetic without using conversion functions, you should use the INTERVAL YEAR TO MONTH data type instead. This data type allows you to store a period of time in terms of years and months, which aligns with the requirement of storing a loan period.

upvoted 1 times

🗳️ 👤 **oblitius** 1 year, 8 months ago

**Selected Answer: C**

ChatGPT came back with C (INTERVAL DAY TO SECOND):

-This data type stores a period of time in terms of days, hours, minutes, and seconds.

-It supports date arithmetic and can precisely represent the number of days a loan remains unpaid, making it suitable for calculating daily interest.



-This seems like the most appropriate choice for the given requirements.

Based on the analysis:



The most appropriate choice is C. INTERVAL DAY TO SECOND because it fulfills all the requirements.

Options A, B, D, and E are not suitable because they either don't consider the number of days specifically (A) or represent specific points in time rather than durations (B, D, E).

-----



Please correct me if I'm wrong.

upvoted 1 times

  **ESZ** 1 year, 8 months ago



ChatGPT comes back with INTERVAL YEAR TO MONTH

upvoted 1 times

  **kay000001** 8 months, 2 weeks ago

Thank you for your thorough research.

upvoted 1 times

  **obsy** 1 year, 11 months ago

**Selected Answer: B**

B is correct

upvoted 1 times

  **Izzicertificacion** 1 year, 11 months ago

¿porque es correcto?

upvoted 1 times

  **dilshod** 2 years ago

C is correct

upvoted 1 times

  **TheOracleWasTaken** 1 year, 11 months ago

but why?

upvoted 1 times

Which two are true about a SQL statement using SET operators such as UNION? (Choose two.)

- A. The data type group of each column returned by the second query must match the data type group of the corresponding column returned by the first query.
- B. The names and number of columns must be identical for all select statements in the query.
- C. The data type of each column returned by the second query must be implicitly convertible to the data type of the corresponding column returned by the first query.
- D. The data type of each column returned by the second query must exactly match the data type of the corresponding column returned by the first query.
- E. The number, but not names, of columns must be identical for all select statements in the query.

**Suggested Answer: AE**

Community vote distribution

AE (100%)

🗳️ **826cb95** 2 months ago

**Selected Answer: CE**

Not A. because "Data type group" isn't a strict SQL requirement – Oracle uses implicit conversion, not strict groups.  
upvoted 2 times

🗳️ **ShahedOdeh** 9 months, 2 weeks ago

A , C and E are correct  
upvoted 1 times

🗳️ **Fargo01** 1 year, 3 months ago

CE for sure 100%  
upvoted 4 times

🗳️ **zzz02** 9 months, 1 week ago

Agree on that one  
upvoted 1 times

🗳️ **yaya32** 1 year, 5 months ago

**Selected Answer: AE**

AE for me - I don't think they need to match exactly.  
upvoted 2 times

🗳️ **lucemqy** 1 year, 7 months ago

**Selected Answer: AE**

The data types does not need to match exactly  
upvoted 1 times

🗳️ **obsy** 1 year, 11 months ago

**Selected Answer: AE**

AE Correct  
upvoted 1 times

🗳️ **dilshod** 2 years ago

CE is 100% percent is correct  
upvoted 4 times

🗳️ **Nelita** 2 years ago

**Selected Answer: AE**

AE are correct.  
upvoted 1 times

🗳️ **Viviana3184** 2 years, 6 months ago



c is incorrect ->The set operators are not valid on columns of type BLOB, CLOB, BFILE, VARRAY, or nested table.

i think D E are correct

The UNION, INTERSECT, and MINUS operators are not valid on LONG columns.

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/The-UNION-ALL-INTERSECT-MINUS-Operators.html#GUID-B64FE747-586E-4513-945F-80CB197125EE>

upvoted 2 times

  **Viviana3184** 2 years, 6 months ago

sorry, AE are correct, not DE


upvoted 2 times

Which two are true about queries using set operators such as UNION? (Choose two.)

- A. In a query containing multiple set operators, INTERSECT always takes precedence over UNION and UNION ALL.
- B. An expression in the first SELECT list must have a column alias for the expression.
- C. All set operators are valid on columns of all data types.
- D. CHAR columns of different lengths used with a set operator return a VARCHAR2 whose length equals the longest char value.
- E. Queries using set operators do not perform implicit conversion across data type groups (e.g. character, numeric).

**Suggested Answer: BC**

Community vote distribution



**Thameur01** 6 months ago

**Selected Answer: AD**

A. In a query containing multiple set operators, INTERSECT always takes precedence over UNION and UNION ALL.

Correct:

In queries with multiple set operators, the precedence order is:

INTERSECT

UNION / UNION ALL

MINUS

If parentheses are not used to explicitly define precedence, INTERSECT is evaluated before UNION or UNION ALL.

D. CHAR columns of different lengths used with a set operator return a VARCHAR2 whose length equals the longest char value.

Correct:

When CHAR columns of varying lengths are combined using a set operator, Oracle returns the result as a VARCHAR2 column whose length matches the longest CHAR value in the result set.

upvoted 2 times

**acd3ea6** 9 months ago

**Selected Answer: DE**

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/The-UNION-ALL-INTERSECT-MINUS-Operators.html#GUID-B64FE747-586E-4513-945F-80CB197125EE>

upvoted 3 times

**Bissto** 10 months ago

**Selected Answer: AD**

A. In a query containing multiple set operators, INTERSECT always takes precedence over UNION and UNION ALL.

When multiple set operators are used in a query, INTERSECT has higher precedence than UNION and UNION ALL. This means that the INTERSECT operation will be evaluated before the other set operators, unless parentheses are used to explicitly define the order of operations.

D. CHAR columns of different lengths used with a set operator return a VARCHAR2 whose length equals the longest char value.

When CHAR columns of different lengths are combined using set operators, the result is a VARCHAR2 column with a length equal to the longest CHAR value in the result set. This is due to how Oracle handles the return type when different CHAR lengths are combined.

upvoted 1 times

**yaya32** 1 year, 5 months ago

**Selected Answer: DE**

DE are correct

upvoted 3 times

**lucemqy** 1 year, 7 months ago

**Selected Answer: DE**

DE are the correct answer

upvoted 2 times

🗄️ 👤 **HassanMkhlalati** 1 year, 9 months ago

**Selected Answer: DE**

DE are correct

upvoted 1 times

🗄️ 👤 **kuff** 2 years, 1 month ago

**Selected Answer: DE**

DE are correct

upvoted 1 times

🗄️ 👤 **Orxan\_H** 2 years, 2 months ago

**Selected Answer: DE**

D E are correct.

source: [https://docs.oracle.com/cd/B19306\\_01/server.102/b14200/queries004.htm](https://docs.oracle.com/cd/B19306_01/server.102/b14200/queries004.htm)

upvoted 1 times

🗄️ 👤 **Sagarkpvs** 2 years, 3 months ago

DE , tested that

upvoted 1 times

🗄️ 👤 **Poke90** 2 years, 3 months ago

**Selected Answer: DE**

Definitely not BC

upvoted 1 times

🗄️ 👤 **jfc1** 2 years, 5 months ago

**Selected Answer: DE**

B is incorrect because it is not mandatory to have an alias for each column in the first SELECT

```
select first_name || "||" last_name, employee_id
```

```
from hr.employees
```

```
union
```

```
select first_name, employee_id
```

```
from HR.EMPLOYEES
```

```
order by 1,2;
```

C is incorrect because the datatypes : BLOC,CLOB are not valid the set Operators.

DE is the correct answer.

upvoted 1 times

🗄️ 👤 **dexdinh91** 2 years, 5 months ago

**Selected Answer: DE**

DE are correct

upvoted 1 times

🗄️ 👤 **Viviana3184** 2 years, 6 months ago

DE are correct

upvoted 1 times

🗄️ 👤 **anzac** 2 years, 6 months ago

According to:

[https://docs.oracle.com/cd/B19306\\_01/server.102/b14200/queries004.htm](https://docs.oracle.com/cd/B19306_01/server.102/b14200/queries004.htm)

"If both queries select values of datatype CHAR of equal length, then the returned values have datatype CHAR of that length. If the queries select values of CHAR with different lengths, then the returned value is VARCHAR2 with the length of the larger CHAR value."

AND

"In queries using set operators, Oracle does not perform implicit conversion across datatype groups. "

So DE.

upvoted 4 times

🗄️ 👤 **Renad\_16** 2 years, 6 months ago

I think it's BD?

upvoted 1 times

  **Orxan\_H** 2 years, 2 months ago

B may be correct if in the end will be "order by" clause. Therefore DE correct  
upvoted 2 times

Examine this business rule:

Each student can work on multiple projects and each project can have multiple students.

You must design an Entity Relationship (ER) model for optimal data storage and allow for generating reports in this format:

STUDENT\_ID FIRST\_NAME LAST\_NAME PROJECT\_ID PROJECT\_NAME PROJECT\_TASK



Which two statements are true? (Choose two.)

- A. PROJECT\_ID must be the primary key in the PROJECTS entity and foreign key in the STUDENTS entity.
- B. STUDENT\_ID must be the primary key in the STUDENTS entity and foreign key in the projects entity.
- C. An associative table must be created with a composite key of STUDENT\_ID and PROJECT\_ID, which is the foreign key linked to the students and projects entities.
- D. The ER must have a many-to-many relationship between the STUDENTS and PROJECTS entities that must be resolved into one-to-many relationships.
- E. The ER must have a one-to-many relationship between the STUDENTS and PROJECTS entities.

**Suggested Answer:** CD

Community vote distribution

CD (100%)

  **CMjer** 1 year, 7 months ago

**Selected Answer:** CD

<https://www.examttopics.com/discussions/oracle/view/15644-exam-1z0-071-topic-1-question-17-discussion/>  
upvoted 1 times

  **Orxan\_H** 1 year, 8 months ago

Why B incorrect?  
upvoted 1 times

  **Tof0ra** 10 months, 4 weeks ago

If student\_id is a FK in the projects, only one student can "participate" in the project. Actually you could have it, as long as you also have the many-to-many, but is bad practice.  
upvoted 1 times

Which three are key components of an Entity Relationship Model? (Choose three.)

- A. an activity
- B. a table
- C. a relationship
- D. an attribute
- E. a unique identifier
- F. an entity

**Suggested Answer:** CDF

*Community vote distribution*

CDF (100%)

🗳️ 👤 **ShahedOdeh** 9 months, 2 weeks ago

C , D and F are correct  
upvoted 1 times

🗳️ 👤 **obsy** 1 year, 11 months ago

CDF CORRECT  
upvoted 1 times

🗳️ 👤 **Nelita** 2 years ago

**Selected Answer: CDF**  
CDF ARE CORRECT  
upvoted 1 times

🗳️ 👤 **Orxan\_H** 2 years, 2 months ago

**Selected Answer: CDF**  
CDF are correct  
<https://moqups.com/templates/diagrams-flowcharts/erd/>  
upvoted 2 times



Examine the data in the ORDERS table:

ORDER_ID	ORDER_DATE
1	<null>
2	<null>
3	01-JAN-2019
4	01-FEB-2019
5	01-MAR-2019

Examine the data in the INVOICES table:

INVOICE_ID	ORDER_ID	ORDER_DATE
1	1	<null>
2	2	01-JAN-2019
3	3	<null>
4	4	01-FEB-2019
5	5	01-APR-2019

Examine this query:


```
SELECT order_id, order_date FROM orders
INTERSECT
SELECT order_id, order_date FROM invoices;
```

- A. 2
- B. 1
- C. 3
- D. 5 01-MAR-2019
- E. 3 01-JAN-2015
- F. 4 01-FEB-2015

**Suggested Answer:** AE

Community vote distribution

BF (100%)

 **abdullah\_barham** Highly Voted 2 years, 6 months ago

**Selected Answer:** BF

drop table orders;

create table orders(order\_id number , order\_date date);

insert into orders values(1,null);

insert into orders values(2,null);

insert into orders values(3,'1-jan-2019');

insert into orders values(4,'1-feb-2019');

insert into orders values(5,'1-mar-2019');

create table invoices(

invoice\_id number, order\_id number, order\_date date);

insert into invoices values(1,1,null);

insert into invoices values(2,2,'01-jan-2019');

insert into invoices values(3,3,null);

insert into invoices values(4,4,'01-feb-2019');

insert into invoices values(5,5,'01-APR-2019');

```
SELECT ORDER_ID , ORDER_DATE FROM ORDERS
INTERSECT
SELECT ORDER_ID , ORDER_DATE FROM INVOICES;
```

ANSWER IS B,F.

upvoted 11 times

🗳️ 👤 **Mcromeo** 1 year, 6 months ago

Based on your own example, it return 2 rows. so B is not correct

upvoted 1 times

🗳️ 👤 **ShahedOdeh** 9 months, 2 weeks ago

options are not showing the number of rows , they are showing the actual result .

So , the correct answers are : B and F

upvoted 1 times

🗳️ 👤 **Rik92** Highly Voted 🏆 2 years, 4 months ago

This needs an update badly. There is no question, just some tables and a query showing on this site..

upvoted 10 times

🗳️ 👤 **luismonge** Most Recent 🕒 1 month, 2 weeks ago

Selected Answer: BF

1

4 01-FEB-19

upvoted 1 times

🗳️ 👤 **ArslanAltaf** 1 year ago

I think the question is what the result of intersect

The result will only include rows where both the order\_id and order\_date match between the ORDERS and INVOICES tables.

So, order ID 1& 4 has exact match ----- but the 4 1-Feb-2019 is not in result (if you dont consider the Typo in option F)

B seems to be the answer in given circumstances

upvoted 1 times

🗳️ 👤 **Fargo01** 1 year, 3 months ago

Best answer is B. 01-FEB-2019

upvoted 1 times

🗳️ 👤 **uselessscript** 1 year, 7 months ago

Just finished this exam. Question: Which three results will be returned?

upvoted 1 times

🗳️ 👤 **metallars** 1 year, 8 months ago

Selected Answer: BF

La respuesta es la B y la F

upvoted 1 times

🗳️ 👤 **Sangiii** 1 year, 10 months ago

B and F is answer- below is the result we will get. F option has typo 01-FEB-2019

orderid order\_date

1 (null)

4 01-02-19

upvoted 1 times

🗳️ 👤 **AishMani20** 2 years ago

Question is not there

upvoted 1 times

🗳️ 👤 **zouve** 2 years ago

In Oracle, the INTERSECT operator is used to combine the result sets of two or more SELECT statements and returns only the rows that are common to all result sets.

upvoted 1 times

🗨️ 👤 **Nelita** 2 years ago

**Selected Answer: BF**

BF ARE CORRECT

upvoted 1 times

🗨️ 👤 **kuff** 2 years, 1 month ago

**Selected Answer: BF**

BF is the answer

upvoted 1 times

🗨️ 👤 **Ricx24** 2 years, 4 months ago

I already took the exam ant these question ask for 3 answers

B, E, F

upvoted 1 times

🗨️ 👤 **Shakhzod1999** 2 years, 4 months ago

can you write to me on telegram @Melsson

upvoted 1 times

🗨️ 👤 **SantiBZ\_07032022\_1744** 2 years, 4 months ago

Tesded:

1 null

4 01/02/19 00:00:00

upvoted 3 times

🗨️ 👤 **jfc1** 2 years, 5 months ago

BF is the answer

upvoted 1 times

🗨️ 👤 **cadcadley** 2 years, 5 months ago

bf only if F is 01-FEB-2019 and not 01-FEB-2015 i am sorry then is B only

upvoted 2 times

🗨️ 👤 **cadcadley** 2 years, 5 months ago

is t 100% bf went we do created everything

upvoted 1 times

Which two will execute successfully? (Choose two.)

- A. SELECT COALESCE(0, SYSDATE) FROM DUAL;
- B. SELECT NVL('DATE', SYSDATZ) FROM DUAL;
- C. SELECT COALESCE('DATE', SYSDATE) FROM DUAL;
- D. SELECT NVL('DATE', 200) FROM (SELECT NULL AS "DATE" FROM DUAL);
- E. SELECT COALESCE('DATE', SYSDATE) FRM (SELECT NULL AS "DATE" FROM DUAL);

**Suggested Answer: BE**

Community vote distribution

BD (100%)

🗳️ 👤 **62d4f30** 9 months, 2 weeks ago

why isnt C correct?

upvoted 1 times

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

BD is correct provided there is a type on Question B, sysdate

upvoted 1 times

🗳️ 👤 **Omar116** 1 year, 8 months ago

**Selected Answer: BD**

checked

upvoted 3 times

🗳️ 👤 **Sangiii** 1 year, 10 months ago

B and D is correct answers.

upvoted 2 times

🗳️ 👤 **zouve** 2 years ago

**Selected Answer: BD**

checked

for E -> ORA-00932: inconsistent datatypes: expected CHAR got DATE

upvoted 1 times

🗳️ 👤 **Nelita** 2 years ago

**Selected Answer: BD**

BD are correct

upvoted 2 times

🗳️ 👤 **kuff** 2 years, 1 month ago

**Selected Answer: BD**

BD is the answer

upvoted 2 times

🗳️ 👤 **kaplanek** 2 years, 2 months ago

**Selected Answer: BD**

verified, only BD is correct, E will throw an error ORA-00932

upvoted 1 times

🗳️ 👤 **jfc1** 2 years, 5 months ago

E is wrong

ORA-00932: inconsistent datatypes: expected CHAR got DATE



BD is the correct one

upvoted 1 times

🗳️ 👤 **cadcadley** 2 years, 5 months ago

BD IS 100% IF YOU GO TO CHEK

upvoted 1 times

  **Viviana3184** 2 years, 6 months ago

B D if B is SELECT NVL('DATE', SYSDate) FROM DUAL;

upvoted 1 times

  **Babi265** 2 years, 6 months ago

Checked should be B D. E is wrong "inconsistent data type"

upvoted 1 times

Which three statements are true about a self join? (Choose three.)

- A. It must be an equijoin.
- B. The ON clause must be used.
- C. It must be an inner join.
- D. It can be an outer join.
- E. The ON clause can be used.
- F. The query must use two different aliases for the table.

**Suggested Answer:** DEF

Community vote distribution

DEF (67%)

BCF (33%)

826cb95 2 months ago

Selected Answer: DEF

D. It can be an outer join.

True: A self join can indeed be an outer join. For example, you can use a left outer join or a right outer join when joining a table to itself, allowing you to retain unmatched rows from one of the instances of the table.

upvoted 1 times

kay000001 8 months, 1 week ago

Selected Answer: BDF

Answer is B, D, F.

It cannot be E. 'On' clause Must be used.

upvoted 1 times

ShahedOdeh 9 months, 2 weeks ago

D , E and F

upvoted 1 times

yaya32 1 year, 5 months ago

Selected Answer: DEF

DEF for me

upvoted 2 times

yanoolthecool 1 year, 6 months ago

DEF, D says CAN

upvoted 1 times

Nico\_Nikko 1 year, 3 months ago

It can never be an outer JOIN though, because you are JOINing Data with itself it will always be as if you would use INTERSECT as all Data of 1 Table exists in the other. Therefore it must be an inner JOIN.

Also regarding F, you can just give one of the tables an alias, the other one does not need one, while both have different names then, they technically do not have different aliases, which leads me to choose C and E.

The rest seems wrong to me.

upvoted 1 times

Leo1212 1 year, 6 months ago

Selected Answer: BCF

not sure why ain't like this

upvoted 1 times

You execute this query:

```
SELECT TO_CHAR(NEXT_DAY(LAST_DAY(SYSDATE), 'MON'), 'dd "Monday for" fmMonth rrrr')  
FROM DUAL;
```

What is the result?

- A. It executes successfully but does not return any result.
- B. It returns the date for the first Monday of the next month.
- C. It returns the date for the last Monday of the current month.
- D. It generates an error.

**Suggested Answer: B**

Community vote distribution

B (67%)

D (33%)

🗨️ 👤 **Bissto** 10 months ago

**Selected Answer: B**

B. It returns the date for the first Monday of the next month.

Explanation:

LAST\_DAY(SYSDATE): This function returns the last day of the current month.

NEXT\_DAY(LAST\_DAY(SYSDATE), 'MON'): This function returns the next Monday after the date provided (which in this case is the last day of the current month). Since the last day of the current month is used, NEXT\_DAY will return the first Monday of the next month.

TO\_CHAR(..., 'dd "Monday for" fmMonth rrrr'): This formats the date to display the day (dd), followed by the string "Monday for", followed by the full month name (with fm removing any leading spaces), and finally the year (rrrr).

Thus, the query returns the date for the first Monday of the next month formatted as specified.

upvoted 2 times

🗨️ 👤 **ArslanAltat** 1 year, 2 months ago

B correct.

D is wrong. output is "06 Monday for May 2024"

upvoted 1 times

🗨️ 👤 **deepzsiva** 1 year, 5 months ago

**Selected Answer: B**

Answer is B. Tested.

upvoted 3 times

🗨️ 👤 **zangado** 1 year, 6 months ago

**Selected Answer: D**

There should be 'Monday' instead of 'MON'

upvoted 2 times

🗨️ 👤 **ogi33** 1 year, 2 months ago

The argument char must be a day of the week in the date language of your session, either the full name or the abbreviation

[https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/NEXT\\_DAY.html#GUID-01B2CC7A-1A64-4A74-918E-26158C9096F6](https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/NEXT_DAY.html#GUID-01B2CC7A-1A64-4A74-918E-26158C9096F6)

upvoted 1 times

🗨️ 👤 **zibby92** 1 year, 6 months ago

<https://www.examtopycs.com/discussions/oracle/view/27077-exam-1z0-082-topic-1-question-86-discussion/>

It works fine

upvoted 2 times





Which two statements are true about the WHERE and HAVING clauses in a SELECT statement? (Choose two.)

- A. WHERE and HAVING clauses can be used in the same statement only if applied to different table columns.
- B. The WHERE clause can be used to exclude rows after dividing them into groups.
- C. The HAVING clause can be used with aggregating functions in subqueries.
- D. The WHERE clause can be used to exclude rows before dividing them into groups.
- E. Aggregating functions and columns used in HAVING clauses must be specified in the SELECT list of a query.

**Suggested Answer:** CD

Community vote distribution

CD (100%)

🗳️ 👤 **826cb95** 2 months ago

**Selected Answer:** CD

C & D tested fine.

upvoted 1 times

🗳️ 👤 **RIDA\_007** 7 months, 1 week ago

C is incorrect because subqueries using aggregate functions typically do not use HAVING clauses; rather, HAVING is used in the outer query to filter the results of aggregates

upvoted 1 times

🗳️ 👤 **yaya32** 11 months, 2 weeks ago

**Selected Answer:** CD

CD for me

upvoted 2 times


Which two are true about global temporary tables? (Choose two.)

- A. Indexes can be created on them.
- B. Backup and recovery operations are available for these tables.
- C. Their data is always stored in the default temporary tablespace of the user who created them.
- D. If the ON COMMIT clause is transaction-specific, all rows in the table are deleted after each COMMIT OR ROLLBACK.
- E. They can be created only by a user with the DBA role, but can be accessed by all users who can create a session.
- F. If the ON COMMIT clause is session-specific, the table is dropped when the session is terminated.

**Suggested Answer:** AC

Community vote distribution

AD (100%)

 **lucemqy** 7 months, 2 weeks ago

**Selected Answer:** AD

AD is correct


upvoted 2 times

 **DobrySomsiad** 9 months, 1 week ago

**Selected Answer:** AD

AD is correct

upvoted 2 times

 **Hari1233535** 9 months, 1 week ago

AD is correct


upvoted 2 times

 **zouve** 1 year ago

Option A is correct. Indexes can be created on global temporary tables, allowing for efficient data retrieval and manipulation.

Option D is correct. If the ON COMMIT clause is specified as transaction-specific (ON COMMIT DELETE ROWS), all rows in the global temporary table are automatically deleted after each COMMIT or ROLLBACK statement. This ensures that the table is cleared for the next transaction.

upvoted 1 times

 **CMjer** 1 year, 1 month ago


**Selected Answer:** AD

A D are correct:

D. If the ON COMMIT clause is transaction-specific, all rows in the table are deleted after each COMMIT OR ROLLBACK.

"all rows in the table are deleted" = TRUNCATE

upvoted 2 times

 **Kamiyl** 1 year, 3 months ago

C isn't correct:

By default, Oracle stores the data of the global temporary table in the default temporary tablespace of the table's owner.

But you can explicitly assign another tablespace to a global temporary table during table creation using the TABLESPACE clause

<https://www.oracletutorial.com/oracle-basics/oracle-global-temporary-table/>

upvoted 2 times

 **Sidoogle** 1 year, 4 months ago

AC is correct

D is wrong, because :

The ON COMMIT clause specifies whether data in the table is transaction-specific or session-specific:

The ON COMMIT DELETE ROWS clause specifies that the global temporary table is transaction-specific. It means that ORACLE TRUNCATE THE TABLE after each commit.

The ON COMMIT PRESERVE ROWS clause specifies that the global temporary table is session-specific, meaning that ORACLE TRUNCATE THE TABLE when you terminate the session, not when you commit a transaction.

Delete is DML, truncate is DDL

upvoted 1 times

🗨️ 👤 **Averagenoob** 1 year, 4 months ago

I think A,C

upvoted 1 times

🗨️ 👤 **SantiBZ\_07032022\_1744** 1 year, 5 months ago

I Think A, D.

D -> If the ON COMMIT is transaction specific -> COMMIT, ROLLBACK delete rows.

NOT C-> Allocation of temporal segments are done on USER TEMP TABLESPACE, at time of use table, no always in the TEMP TABLESPACE of the OWNER.

upvoted 3 times

🗨️ 👤 **Viviana3184** 1 year, 6 months ago

I think are A and C

upvoted 1 times

Which three are true about privileges? (Choose three.)

- A. A combination of object and system privileges can be granted to a role.
- B. All types of schema objects have associated object privileges.
- C. Schema owners can grant object privileges on objects in their schema to any other user or role.
- D. Only users with the DBA role can create roles.
- E. Only users with the GRANT ANY PRIVILEGE privilege can grant and revoke system privileges from other users.
- F. Object privileges granted on a table automatically apply to all synonyms for that table.

**Suggested Answer:** ACF

Community vote distribution

ACF (100%)

🗳️ 👤 **kay000001** 8 months, 2 weeks ago

**Selected Answer:** ACF

F: Is correct.

upvoted 1 times

🗳️ 👤 **a947739** 11 months ago

ACE

F is incorrect

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer:** ACF

ACF correct

upvoted 1 times

🗳️ 👤 **TheOracleWasTaken** 1 year, 11 months ago

why not B?

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 7 months ago

[https://docs.oracle.com/cd/A58617\\_01/server.804/a58227/ch18.htm](https://docs.oracle.com/cd/A58617_01/server.804/a58227/ch18.htm)

Some schema objects do not have associated object privileges; their use is controlled with system privileges

upvoted 1 times

🗳️ 👤 **Sbablo** 2 years, 1 month ago

why is E wrong ?

upvoted 1 times

🗳️ 👤 **zouve** 1 year, 11 months ago

i think the "only" is killing it

upvoted 2 times

🗳️ 👤 **astride** 2 years, 2 months ago

Why C is not correct?

upvoted 1 times

🗳️ 👤 **kaplanek** 2 years, 2 months ago

A C F are correct

upvoted 1 times

Examine the description of the EMPLOYEES table:


Name	Null?	Type
EMPLOYEE_ID	NOT NULL	NUMBER(6)
FIRST_NAME		VARCHAR2(20)
LAST_NAME	NOT NULL	VARCHAR2(25)
SALARY		NUMBER(8,2)
HIREDATE		DATE
DEPARTMENT_ID		NUMBER(4)

Which two statements will insert a row into the EMPLOYEES table? (Choose two.)

- A. INSERT INTO employees VALUES (101, 'John', 'Smith', 12000, SYSDATE);
- B. INSERT INTO employees VALUES (101, 'John', 'Smith', 10, 12000, SYSDATE);
- C. INSERT INTO employees (employee\_id, salary, first\_name, hiredate, last\_name) VALUES (101, 12100, 'John', SYSDATE, 'Smith');
- D. INSERT INTO employees (employee\_id, first\_name, last\_name, salary, hiresate) VALUES ( (SELECT 101, 'John', 'Smith', 12000, SYSDATE FROM dual) );
- E. INSERT INTO employees SELECT 101, 'John', 'Smith', 12000, (SELECT SYSDATE FROM dual), 10 FROM dual;
- F. INSERT INTO employees VALUES (101, 'John', '', 12000, SYSDATE, 10);

**Suggested Answer:** CF

Community vote distribution



**Babi265** Highly Voted 2 years, 6 months ago

i am finding 3 that are correct. CEF  
upvoted 7 times

**Renad\_16** Highly Voted 2 years, 6 months ago

I think it's CE ?  
upvoted 5 times

**826cb95** Most Recent 2 months ago

**Selected Answer:** CE

✓ C.

sql

Copy

Edit

INSERT INTO employees

(employee\_id, salary, first\_name, hiredate, last\_name)

VALUES (101, 12100, 'John', SYSDATE, 'Smith');

This works: all required columns (EMPLOYEE\_ID, LAST\_NAME) are provided.

The column order doesn't have to match the table, as long as it's declared.

DEPARTMENT\_ID is nullable, so it's fine to omit.

✓ E.

sql

Copy

Edit

INSERT INTO employees

SELECT 101, 'John', 'Smith', 12000, (SELECT SYSDATE FROM dual), 10 FROM dual;

upvoted 1 times

🗨️ 👤 **Thameur01** 6 months ago

**Selected Answer: CE**

C: Explicitly matches columns with their values.

E: Uses a SELECT query to populate all columns correctly.

F: Provides all values in the correct order.

upvoted 1 times

🗨️ 👤 **ArslanAltaf** 1 year, 2 months ago

CF is correct.

in F, if have space between ' ' for last Name, it will be accepted. if there is no space, then not null constraint will apply.

upvoted 1 times

🗨️ 👤 **usarovsherzod68** 1 year, 3 months ago

A, C for me. E is wrong because SELECT 101 is not correct.

upvoted 1 times

🗨️ 👤 **yaya32** 1 year, 5 months ago

**Selected Answer: EF**

EF is correct for me as ' ' is not null...

upvoted 1 times

🗨️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: CE**

CE definitely works tested

upvoted 1 times

🗨️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: CD**

CD also works

upvoted 1 times

🗨️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: CF**

CF Should be the correct answer

upvoted 1 times

🗨️ 👤 **jm9999** 1 year, 9 months ago

C, E worked. F gave me ora-01400

upvoted 1 times

🗨️ 👤 **TheOracleWasTaken** 1 year, 11 months ago

**Selected Answer: CF**

E works too

upvoted 1 times

🗨️ 👤 **zouve** 2 years ago

CE tested

upvoted 1 times

🗨️ 👤 **zouve** 2 years ago

if there is a space inside ' ' on F is working

upvoted 1 times

🗨️ 👤 **Nelita** 2 years ago

**Selected Answer: CE**

CEF are correct. I checked in data base.

upvoted 2 times

🗨️ 👤 **CyberP** 2 years, 1 month ago

Why A is wrong?

upvoted 1 times

🗨️ 👤 **CyberP** 2 years, 1 month ago

Maybe AC are correct?

upvoted 3 times

🗨️ 👤 **TheOracleWasTaken** 1 year, 11 months ago

when you specify the values keyword you must insert a value for every column in the table. If it permits nulls then just insert null  
upvoted 1 times

🗨️ 👤 **shotcom** 1 year, 9 months ago

<https://www.techonthenet.com/oracle/insert.php>

according to that link You can omit a column from the Oracle INSERT statement if the column allows NULL values. so answer A is correct  
upvoted 2 times

🗨️ 👤 **J4vi** 2 years, 1 month ago

**Selected Answer: CE**

C E are correct. F would be ok with an space between ' '.

If you copy and paste the exact alternative from here, it works, but I imagine that the exam doesn't have that space and that's why there are only two correct options.

upvoted 2 times

🗨️ 👤 **kuff** 2 years, 1 month ago

**Selected Answer: CE**

CE is the answer

upvoted 1 times

Examine this command:

TRUNCATE TABLE test;

Table truncated.

Which two are true? (Choose two.)

- A. The structure of the TEST table is removed.
- B. All the indexes on the TEST table are dropped.
- C. All the constraints on the TEST table are dropped.
- D. Removed rows can not be recovered using the ROLLBACK command.
- E. All the rows in the TEST table are removed.

**Suggested Answer:** AC

Community vote distribution

DE (100%)

🗳️ 👤 **c1a0049** 9 months, 4 weeks ago

**Selected Answer:** DE

D, E - TRUNCATE only affects rows(data), not table definition.

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer:** DE

DE should be the correct answer

upvoted 3 times

🗳️ 👤 **metallars** 1 year, 8 months ago

**Selected Answer:** DE

D y E son correctas.

upvoted 2 times

🗳️ 👤 **HassanMkhlalati** 1 year, 9 months ago

**Selected Answer:** DE

The structure of the TEST table is still available. It is dropped only when drop the table

upvoted 1 times

🗳️ 👤 **AishMani20** 2 years ago

I think instead of TRUNCATE command DROP should be use to get option A and C as answers

upvoted 1 times

🗳️ 👤 **Nelita** 2 years ago

**Selected Answer:** DE

DE ARE CORRECT

upvoted 2 times

🗳️ 👤 **kuff** 2 years, 1 month ago

**Selected Answer:** DE

DE is the answer

upvoted 1 times

🗳️ 👤 **Klawal** 2 years, 2 months ago

**Selected Answer:** DE

D and E are the correct answers.

upvoted 1 times



🗒️ 👤 **pmeyer** 2 years, 2 months ago

Selected Answer: DE

"D" and "E" are correct.

upvoted 1 times

🗒️ 👤 **Orxan\_H** 2 years, 2 months ago

Selected Answer: DE

DE are correct

upvoted 2 times

🗒️ 👤 **tonl20ll** 2 years, 4 months ago

Selected Answer: DE

DE are correct

upvoted 2 times

🗒️ 👤 **Rik92** 2 years, 4 months ago

Selected Answer: DE

DE are correct.

upvoted 2 times

🗒️ 👤 **jfc1** 2 years, 5 months ago

Selected Answer: DE

[https://docs.oracle.com/database/121/SQLRF/statements\\_10007.htm#SQLRF01707](https://docs.oracle.com/database/121/SQLRF/statements_10007.htm#SQLRF01707)

DE is correct answer

upvoted 2 times

🗒️ 👤 **Beesho** 2 years, 5 months ago

DE is correct

A. The structure of the TEST table is removed. wrong, only content will get removed

B. All the indexes on the TEST table are dropped. indexes will stay

C. All the constraints on the TEST table are dropped. wrong

D. Removed rows can not be recovered using the ROLLBACK command. true

E. All the rows in the TEST table are removed. true

upvoted 2 times

🗒️ 👤 **dexdinh91** 2 years, 5 months ago

Selected Answer: DE

DE are correct

upvoted 2 times

🗒️ 👤 **Averagenoob** 2 years, 5 months ago

DE Correct Answer

upvoted 1 times

🗒️ 👤 **Viviana3184** 2 years, 6 months ago

DE is correct

upvoted 1 times

You issued this command:


```
DROP TABLE hr.employees;
```

Which three statements are true? (Choose three.)

- A. Sequences used to populate columns in the HR.EMPLOYEES table are dropped.
- B. Synonyms for HR.EMPLOYEES are dropped.
- C. Views referencing HR.EMPLOYEES are dropped.
- D. All constraints defined on HR.EMPLOYEES are dropped.
- E. The HR.EMPLOYEES table may be moved to the recycle bin.
- F. All indexes defined on HR.EMPLOYEES are dropped

**Suggested Answer:** DEF

Community vote distribution



🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer: DEF**

DEF is the right answer.

upvoted 1 times

🗳️ 👤 **826cb95** 2 months ago

**Selected Answer: BDF**

Synonyms for HR.EMPLOYEES are dropped.

When you drop a table, any synonyms that point to that table are automatically dropped because the underlying object no longer exists.

upvoted 1 times

🗳️ 👤 **yaya32** 11 months, 2 weeks ago

**Selected Answer: DEF**

DEF for me

upvoted 3 times

🗳️ 👤 **Mcromeo** 1 year ago

**Selected Answer: DEF**

DEF is correct

upvoted 1 times

🗳️ 👤 **zangado** 1 year ago

**Selected Answer: DEF**

I think the same as lucemqy and jm9999

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 1 month ago

**Selected Answer: DEF**

Views and synonyms are not dropped but invalid

upvoted 1 times

🗳️ 👤 **jm9999** 1 year, 2 months ago



**Selected Answer: EF**

For me, only EF proved to be true. The constraint was still listed in user\_constraints after issuing the drop. I think ABC are definitely false and you have to pick 3 so ....



upvoted 1 times

🗳️ 👤 **captainkul** 1 year, 3 months ago

D,E,F are correct  
upvoted 1 times

  **jm9999** 1 year, 3 months ago

DEF, I think synonyms and views are invalidated but not dropped  
upvoted 2 times

  **Sushmasri** 1 year, 5 months ago

**Selected Answer: BCD**

B,C,D are correct  
upvoted 2 times

Examine this statement:

```
SELECT cust_id, cust_last_name "Last Name"  
  FROM customers  
 WHERE country_id = 10  
 UNION  
 SELECT cust_id CUST_NO, cust_last_name  
  FROM customers  
 WHERE country_id = 30
```

Identify three ORDER BY clauses, any one of which will complete the query successfully. (Choose three.)

- A. ORDER BY CUST\_NO
- B. ORDER BY 2, cust\_id
- C. ORDER BY 2, 1
- D. ORDER BY "Last Name"
- E. ORDER BY "CUST\_NO"

**Suggested Answer:** BCD

Community vote distribution

BCD (100%)

826cb95 2 months ago

Selected Answer: BCD

A & E doesn't work because

he ORDER BY clause needs to refer to something that exists in both parts of the union (either a column position or an alias that applies to both parts).

upvoted 1 times

lucemqy 7 months, 2 weeks ago

Selected Answer: BCD

You can only order by column names from the first select query

upvoted 2 times

super04 10 months, 3 weeks ago

Correct answer is BCD.

The ORDER BY clause does not recognize the column names of the second SELECT query. So,

upvoted 1 times


Which two statements are true about views? (Choose two.)

- A. The WITH CHECK clause prevents certain rows from being updated or inserted in the underlying table through the view.
- B. The WITH CHECK clause prevents certain rows from being displayed when querying the view.
- C. Views can be indexed.
- D. Views can be updated without the need to re-grant privileges on the view.
- E. Tables in the defining query of a view must always exist in order to create the view.

**Suggested Answer:** AD

Community vote distribution



AD (100%)

  **826cb95** 2 months ago

**Selected Answer:** AD


Why E is wrong? Because - \* Per official Oracle documentation, you can use 'FORCE' keyword to create view over table that doesn't exists when needed.

<https://www.oracletutorial.com/oracle-view/oracle-create-view/>  
upvoted 1 times



  **lucemqy** 7 months, 2 weeks ago

**Selected Answer:** AD

AD is correct  
upvoted 1 times

  **ESZ** 1 year, 3 months ago

C and E also correct?  
upvoted 1 times

  **CMjer** 1 year, 1 month ago

C - Views can be indexed. FALSE

Only materialised views can be indexed:

<https://asktom.oracle.com/pls/apex/asktom.search?tag=indexes-on-view>

E - Tables in the defining query of a view must always exist in order to create the view. FALSE:

Usually, you create a new view based on existing tables. However, sometimes, you may want to create a view based on the tables that you will create later or the tables that you don't have sufficient privileges to access at the time of creating the view. In these cases, you can use the FORCE option.

<https://www.oracletutorial.com/oracle-view/oracle-create-view/>  
upvoted 1 times

Examine the description of the EMPLOYEES table:

Name	Null?	Type
EMPLOYEE_ID	NOT NULL	NUMBER(3)
FIRST_NAME		VARCHAR2(15)
LAST_NAME	NOT NULL	VARCHAR2(15)
SALARY		NUMBER(6,2)

Which two statements will run successfully? (Choose two.)

- A. SELECT 'The first\_name is " || first\_name || "' FROM employees;
- B. SELECT 'The first\_name is "" || first\_name || ""' FROM employees;
- C. SELECT 'The first\_name is ' || first\_name || "' FROM employees;
- D. SELECT 'The first\_name is ' || first\_name || "' FROM employees;
- E. SELECT 'The first\_name is \" || first\_name || \"' FROM employees;

**Suggested Answer: AC**

Community vote distribution

BD (100%)


 **RIDA\_007** 7 months, 1 week ago

D is incorrect, there is a missing single quote before the (||) operator  
upvoted 1 times

 **archit4321** 8 months ago

**Selected Answer: BD**

- A. SELECT 'The first\_name is " || first\_name || "' FROM (select 'user' first\_name from dual); --O/P = ORA-00923: FROM keyword not found where expected
  - B. SELECT 'The first\_name is "" || first\_name || ""' FROM (select 'user' first\_name from dual); --O/P = The first\_name is 'user'
  - C. SELECT 'The first\_name is ' || first\_name || "' FROM (select 'user' first\_name from dual); --O/P = ORA-00923: FROM keyword not found where expected
  - D. SELECT 'The first\_name is ' || first\_name || "' FROM (select 'user' first\_name from dual); --O/P = The first\_name is user
  - E. SELECT 'The first\_name is \" || first\_name || \"' FROM (select 'user' first\_name from dual); --O/P = ORA-24450: Cannot pre-process OCI statement
- upvoted 1 times

 **yaya32** 11 months, 2 weeks ago

**Selected Answer: BD**

BD for me  
upvoted 1 times

 **lucemqy** 1 year, 1 month ago

**Selected Answer: BD**

BD is correct checked  
upvoted 1 times

 **Sushmasri** 1 year, 5 months ago

Can someone explain why B,D is correct ?  
upvoted 3 times

 **jelenadinere\_1** 1 year, 5 months ago

it depends on quotation mark, since we can not copy from this question normally you have to guess where are double quotes and where are single ones, but if it like that:

```
SELECT 'The first_name is ' || first_name || '' FROM temp; (column name || first_name || '')
SELECT 'The first_name is ' || first_name || ' ' FROM temp; (column name || first_name || ')
SELECT 'The first_name is ' || first_name || ' ' FROM temp; ( The first_name is Jhon )
```

then 3 answers are right: BCD

upvoted 1 times

🗨️ **zouve** 1 year, 6 months ago

**Selected Answer: BD**

BD checked

upvoted 1 times

🗨️ **CyberP** 1 year, 7 months ago

Why its BD ? explain the reason logically & don't tell me I tested it.

In the exam there will be no tool to test anything if something changed so we need to understand the logic please.

upvoted 1 times

🗨️ **kuff** 1 year, 7 months ago

**Selected Answer: BD**

BD is the answer

upvoted 1 times

🗨️ **pmeyer** 1 year, 8 months ago

**Selected Answer: BD**

Who chose the correct answer?

upvoted 2 times

🗨️ **MoonLight** 1 year, 8 months ago

I tried... none works.... anybody can explain even just a bit?

upvoted 2 times

🗨️ **Poke90** 1 year, 9 months ago

**Selected Answer: BD**

tested it

upvoted 1 times

🗨️ **Orakol** 1 year, 10 months ago

\*Tested in SQL\* - BD are correct.

upvoted 1 times

🗨️ **jfc1** 1 year, 11 months ago

**Selected Answer: BD**

BD is the correct answer

upvoted 2 times

🗨️ **dexdinh91** 1 year, 11 months ago

**Selected Answer: BD**

BD are correct

upvoted 2 times

🗨️ **cadcadley** 1 year, 11 months ago

BD is correct regarding to try in SQL

upvoted 2 times

🗨️ **Viviana3184** 2 years ago

BD are correct

upvoted 2 times

Which two are true about unused columns? (Choose two.)

- A. Setting an indexed column to unused results in an error.
- B. You can query the data dictionary to see the names of unused columns.
- C. You can specify multiple column names in an ALTER TABLE...SET UNUSED statement.
- D. If you set all the columns of a table to unused, the table is automatically dropped.
- E. CASCADE CONSTRAINTS must be specified when setting a column to unused if that column is referenced in a constraint on another column.

**Suggested Answer: CE**

Community vote distribution

CE (100%)

🗳️ 👤 **ShahedOdeh** 9 months, 2 weeks ago

C and D

upvoted 1 times

🗳️ 👤 **ShahedOdeh** 9 months, 2 weeks ago

B and C

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: CE**

CE is correct you can see number of unused columns but not their names in the data dictionary

upvoted 1 times

🗳️ 👤 **ShahedOdeh** 9 months, 2 weeks ago

no we have the USER\_TAB\_COLUMNS dictionary that is showing the unused columns

upvoted 1 times

🗳️ 👤 **musafir** 2 years ago

B C E are all correct according to Oracle "SET UNUSED Clause"

<https://docs.oracle.com/en/database/oracle/oracle-database/23/sqlrf/ALTER-TABLE.html#GUID-552E7373-BF93-477D-9DA3-B2C9386F2877>

upvoted 2 times

🗳️ 👤 **shotcom** 1 year, 9 months ago

when you query the dba\_unused\_col\_tabs view which displays a list of all tables with un-used columns, including counts of the number of columns within a table that are unused but it will not show the names of the unused columns so B is wrong

upvoted 5 times

🗳️ 👤 **Nelita** 2 years ago

**Selected Answer: CE**

CE ARE CORRECT.

upvoted 1 times

🗳️ 👤 **Sbablo** 2 years, 1 month ago

All are wrong except E

upvoted 1 times

🗳️ 👤 **shotcom** 1 year, 9 months ago

C is correct, you can specify multiple columns in a single alter statement

eg: ALTER TABLE EMPLOYEES SET UNUSED (JOB\_ID, COMMISSION\_PCT);

upvoted 1 times

🗳️ 👤 **Sbablo** 2 years, 1 month ago

Correct answer B,E

upvoted 1 times



Examine the data in the CUST\_NAME column of the CUSTOMERS table:

```
CUST_NAME
-----
Renske Ladwig
Jason Mallin
Samuel McCain
Allan MCEwen
Irene Mikkilineni
Julia Nayer
```

You want to display the CUST\_NAME values where the last name starts with Mc or MC.

Which two WHERE clauses give the required result? (Choose two.)

- A. WHERE INITCAP(SUBSTR(cust\_name, INSTR(cust\_name, ' ') + 1)) LIKE 'Mc'
- B. WHERE SUBSTR(cust\_name, INSTR(cust\_name, ' ') + 1) LIKE 'Mc'
- C. WHERE SUBSTR(cust\_name, INSTR(cust\_name, ' ') + 1) LIKE 'Mc%' OR 'MC'
- D. WHERE UPPER(SUBSTR(cust\_name, INSTR(cust\_name, ' ') + 1)) LIKE UPPER('MC')
- E. WHERE INITCAP(SUBSTR(cust\_name, INSTR(cust\_name, ' ') + 1)) IN ('MC%', 'Mc%')

**Suggested Answer:** BD

Community vote distribution

AD (100%)

🗳️ 👤 **Thameur01** 5 months, 1 week ago

**Selected Answer:** BD

A: Incorrect: INITCAP changes the first character of the extracted last name to uppercase and the rest to lowercase. However, this does not work for "MC" because INITCAP would change it to "Mc."

B: Correct: This extracts the last name and matches it directly against "Mc%" (case-sensitive). This works because "Mc" is explicitly matched.

D: Correct: UPPER ensures that both the extracted last name and the pattern are converted to uppercase, making the comparison case-insensitive. This matches both "Mc" and "MC."

upvoted 1 times

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

**Selected Answer:** AD

AD is correct LIKE is case sensitive

upvoted 1 times

🗳️ 👤 **zouve** 1 year ago

AD tested 100% guarantee

upvoted 1 times

🗳️ 👤 **Nelita** 1 year ago

**Selected Answer:** AD

AD ARE CORRECT

upvoted 1 times

🗳️ 👤 **kuff** 1 year, 1 month ago

**Selected Answer:** AD

AD is the answer

upvoted 1 times

🗳️ 👤 **pmeyer** 1 year, 2 months ago

**Selected Answer:** AD

A and D are correct, B is wrong.

upvoted 1 times

🗳️ 👤 **Poke90** 1 year, 3 months ago

Selected Answer: AD

B is wrong

upvoted 1 times

🗨️ 👤 **Rik92** 1 year, 4 months ago

Selected Answer: AD

Tried all, AD are correct.

B is incorrect as it only returns McCain and not MCEwan due to Mc and MC mismatch.

upvoted 4 times

🗨️ 👤 **greenneem** 1 year, 4 months ago

B is correct.

D is incorrect as it will also generate the last name starts with small letter m.

upvoted 1 times

🗨️ 👤 **dexdinh91** 1 year, 5 months ago

Selected Answer: AD

AD are correct

upvoted 2 times

🗨️ 👤 **Viviana3184** 1 year, 6 months ago

AD are correct

upvoted 1 times

🗨️ 👤 **Babi265** 1 year, 6 months ago

AD is the correct answer

upvoted 3 times

Which is the default column or columns for sorting output from compound queries using SET operators such as INTERSECT in a SQL statement?

- A. the first VARCHAR2 column in the first SELECT of the compound query
- B. the first column in the first SELECT of the compound query
- C. the first NUMBER column in the first SELECT of the compound query
- D. the first NUMBER or VARCHAR2 column in the last SELECT of the compound query
- E. the first column in the last SELECT of the compound query

**Suggested Answer: E**

*Community vote distribution*

B (100%)

🗳️ 👤 **yaya32** 11 months, 2 weeks ago

**Selected Answer: B**

B for me

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 1 month ago

**Selected Answer: B**

Definitely B!

upvoted 1 times

🗳️ 👤 **Mcromeo** 1 year, 2 months ago

B is right solution

upvoted 2 times

🗳️ 👤 **ESZ** 1 year, 3 months ago

**Selected Answer: B**

B IS CORRECT

upvoted 1 times

🗳️ 👤 **Nelita** 1 year, 7 months ago

**Selected Answer: B**

B IS CORRECT

upvoted 3 times

🗳️ 👤 **Nelita** 1 year, 7 months ago

B IS CORRECT

upvoted 1 times

🗳️ 👤 **kuff** 1 year, 7 months ago

**Selected Answer: B**

B is the answer

upvoted 2 times

🗳️ 👤 **pmeyer** 1 year, 8 months ago

**Selected Answer: B**

B is correct

upvoted 1 times

🗳️ 👤 **Rik92** 1 year, 10 months ago

**Selected Answer: B**

B is correct

upvoted 2 times

🗳️ 👤 **jfc1** 1 year, 11 months ago

**Selected Answer: B**

B is correct

upvoted 3 times

  **dexdinh91** 1 year, 11 months ago

**Selected Answer: B**

B is correct

upvoted 2 times

  **Viviana3184** 2 years ago

B is correct

upvoted 2 times

  **cspiro2410** 2 years ago

**Selected Answer: B**

B for sure

upvoted 2 times

  **Babi265** 2 years ago

Definitely B

upvoted 2 times

Which two statements are true about the ORDER BY clause? (Choose two.)

- A. In a character sort, the values are case-sensitive.
- B. NULLS are not included in the sort operation.
- C. Numeric values are displayed in descending order if they have decimal positions.
- D. Column aliases can be used in the ORDER BY clause.
- E. Only columns that are specified in the SELECT list can be used in the ORDER BY clause.

**Suggested Answer:** AD

Community vote distribution

AD (100%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

Selected Answer: AD

AD is the right answer.

upvoted 1 times

🗳️ 👤 **usarovsherzod68** 9 months, 1 week ago

Selected Answer: AD

for me

upvoted 1 times

🗳️ 👤 **kddata** 9 months, 2 weeks ago

Selected Answer: AD

For me

upvoted 1 times

🗳️ 👤 **Arth1989** 11 months, 3 weeks ago

AD for sure.

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 1 month ago

Selected Answer: AD

AD is the correct answer

upvoted 2 times

🗳️ 👤 **Nelita** 1 year, 7 months ago

Selected Answer: AD

AD ARE CORRECT

upvoted 2 times

Examine the BRICKS table:

COLOUR	SHAPE	WEIGHT
red	cube	5
red	cylinder	10
blue	cube	15
blue	cylinder	20

You write this query:

```
SELECT *  
  FROM bricks b1 CROSS JOIN bricks b2  
 WHERE b1.weight < b2.weight;
```

How many rows will the query return?

- A. 4
- B. 6
- C. 16
- D. 0
- E. 1
- F. 10

**Suggested Answer: B**

Community vote distribution

B (80%)

C (20%)

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

**Selected Answer: B**

The correct answer is B which is 6 rows

Explanation:

The CROSS JOIN operation combines each row from the first table with each row from the second table. Since the BRICKS table has 4 rows, the result of the CROSS JOIN will have  $4 * 4 = 16$  rows.

The WHERE clause filters out rows where b1.weight is not less than b2.weight. This means that only rows where the weight of the first brick is less than the weight of the second brick will be included in the result. There are 6 such combinations: (5,10), (5,15), (5,20), (10,15), (10,20), and (15,20). So the final result will have 6 rows.


upvoted 8 times

 **2339ac0** Most Recent 1 month, 3 weeks ago

**Selected Answer: B**

B is the correct answer.

upvoted 1 times

 **yaya32** 11 months, 2 weeks ago

**Selected Answer: B**

Option B is correct

upvoted 1 times

 **Omar116** 1 year, 2 months ago

checked in DB B is correct

upvoted 1 times

 **Sushmasri** 1 year, 5 months ago

Option B is correct

upvoted 1 times

🗨️ 👤 **Shrimathi** 1 year, 5 months ago

return 6 rows

option B tested

upvoted 1 times

🗨️ 👤 **theleerise** 1 year, 6 months ago

**Selected Answer: C**

C is correct

upvoted 1 times

🗨️ 👤 **CyberP** 1 year, 7 months ago

**Selected Answer: C**

I think its 16

upvoted 1 times

🗨️ 👤 **CyberP** 1 year, 7 months ago

$4*4=16$

upvoted 1 times

🗨️ 👤 **CyberP** 1 year, 7 months ago

forget what i said above, correct answer is 6

upvoted 1 times

Examine this query:


```
SELECT INTERVAL '100' MONTH DURATION FROM DUAL;
```

What will be the output?

- A. an error
- B. DURATION -  
-----  
+100
- C. DURATION -  
-----  
+08
- D. DURATION -  
-----  
+08-04

**Suggested Answer:** D

Community vote distribution



**CyberP** Highly Voted 1 year, 1 month ago

**Selected Answer: D**

Explanation:

The given query is using an interval literal with a value of '100' and a unit of 'MONTH'. When using the interval literal syntax, the specified value represents the number of intervals of the given unit.

In this case, the query is requesting an interval of 100 months. Since there are 12 months in a year, the interval '100' months is equivalent to 8 years and 4 months.

The output format for intervals in Oracle is '+YY-MM', where YY represents the number of years and MM represents the number of months.

Therefore, the output will be '+08-04', indicating an interval of 8 years and 4 months.

The correct answer is D. +08-04.

upvoted 11 times

**2339ac0** Most Recent 1 month, 3 weeks ago

**Selected Answer: D**

D is the right answer. Interval Year to month.

upvoted 1 times

**lucemqy** 7 months, 2 weeks ago

**Selected Answer: D**

D is the correct answer

upvoted 1 times

**metallars** 8 months, 1 week ago

**Selected Answer: D**

D respuesta revisada y correcta

upvoted 1 times

**zouve** 1 year ago

**Selected Answer: D**



try

```
SELECT INTERVAL '100' MONTH as DURATION FROM DUAL;
```

upvoted 1 times

🗨️ 👤 **kuff** 1 year, 1 month ago

**Selected Answer: D**

D is the answer

upvoted 1 times

🗨️ 👤 **pmeyer** 1 year, 2 months ago

**Selected Answer: A**

In my test, it returned an error, even though I adjusted the quotes.

upvoted 1 times

🗨️ 👤 **pmeyer** 1 year, 1 month ago

I found my error, the correct is "D"

upvoted 2 times

🗨️ 👤 **Skye21** 1 year, 3 months ago

it's gets an error when I try it on <https://livesql.oracle.com/apex>

upvoted 2 times

🗨️ 👤 **jr\_C** 1 year, 3 months ago

Don't copy because of bad quote char, write it it's work

upvoted 1 times

🗨️ 👤 **dexdinh91** 1 year, 5 months ago

**Selected Answer: D**

beware of alias

upvoted 1 times

🗨️ 👤 **cadcadley** 1 year, 5 months ago

D is correct

upvoted 2 times

Examine this query:

```
SELECT TRUNC(ROUND(156.00,-2),-1) FROM DUAL;
```

What is the result?

- A. 150
- B. 200
- C. 160
- D. 100
- E. 16

**Suggested Answer: B**

Community vote distribution

B (75%)

A (25%)

🗳️ 👤 **yaya32** 1 year, 5 months ago

**Selected Answer: B**

Answer is B

upvoted 2 times

🗳️ 👤 **andrejshilin16** 11 months, 2 weeks ago

TRUNCATE here do nothing

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: B**

Answer is 200

upvoted 1 times

🗳️ 👤 **metallars** 1 year, 8 months ago

**Selected Answer: B**

Respuesta revisada

upvoted 1 times

🗳️ 👤 **tjsweg** 1 year, 10 months ago

can someone explain why the answer is 200?

upvoted 2 times

🗳️ 👤 **Mcromeo** 1 year, 9 months ago

because  $\text{ROUND}(156.00,-2) = 200$  and  $\text{trunc}(200,-1)=200$  (-1 => first number before coma 0 is truncated to 0)

upvoted 1 times

🗳️ 👤 **zouve** 2 years ago

B is the correct tested

upvoted 2 times

🗳️ 👤 **CMjer** 2 years, 1 month ago

**Selected Answer: B**

B is correct, run it in SQL Developer

upvoted 2 times

🗳️ 👤 **Ac37** 2 years, 3 months ago

I have run it in oracle and the answer is 200.

upvoted 1 times

🗳️ 👤 **AbdullaMA2** 2 years, 3 months ago

**Selected Answer: A**

A is the right one :

The ROUND(156.00, -2) function call rounds the number 156.00 to the nearest hundredth, resulting in the value 200.

The TRUNC(200, -1) function call truncates the number 200 to the nearest ten, resulting in the value 150.



Therefore, the final result of the SQL query is 150.

upvoted 2 times

  **Fargo01** 1 year, 3 months ago

Just run the query and you'll get the result : 200.

upvoted 1 times

  **Rik92** 2 years, 4 months ago

TRUNC does nothing here. The rounding of -2 gives closest order of 100, 156->200. Then 200 is truncated to 200

upvoted 2 times

You want to write a query that prompts for two column names and the where condition each time it is executed in a session but only prompts for the table name the first time it is executed.

The variables used in your query are never undefined in your session.

Which query can be used?

- A. 

```
SELECT &&col1, &&col2
FROM &table
WHERE &&condition = &&cond;
```
- B. 

```
SELECT '&&col1', '&&col2'
FROM &table
WHERE '&&condition' = '&cond';
```
- C. 

```
SELECT &&col1, &&col2
FROM &table
WHERE &&condition;
```
- D. 



```
SELECT &col1, &col2
FROM &table
WHERE &condition;
```
- E. 

```
SELECT &col1, &col2
FROM "&table"
WHERE &condition;
```

**Suggested Answer: C**


Community vote distribution

D (100%)

  **lucemqy** 7 months, 2 weeks ago



**Selected Answer: D**

D is the correct answer  
upvoted 1 times

  **kuff** 1 year, 1 month ago



**Selected Answer: D**

D is the answer  
upvoted 3 times

  **pmeyer** 1 year, 2 months ago


**Selected Answer: D**

D is the correct answer  
upvoted 2 times

  **Rik92** 1 year, 4 months ago

**Selected Answer: D**

D  
&prompts always, && prompts once as it created a session variable  
upvoted 4 times

  **jfc1** 1 year, 5 months ago

**Selected Answer: D**


D is the correct answer  
upvoted 2 times

  **dexdinh91** 1 year, 5 months ago

Selected Answer: D

D is correct

upvoted 2 times

  **cadcadley** 1 year, 5 months ago

C is correct look the question the where condition each time it is executed in a session but only prompts for the table name the first time it is executed.

upvoted 1 times

  **abdullah\_barham** 1 year, 6 months ago

Selected Answer: D

D is the answer

upvoted 2 times

  **Babi265** 1 year, 6 months ago

D is the answer

upvoted 2 times

Which three statements are true about indexes and their administration in an Oracle database? (Choose three.)

- A. The same table column can be part of a unique and non-unique index.
- B. A descending index is a type of function-based index.
- C. An INVISIBLE index is not maintained when DML is performed on its underlying table.
- D. If a query filters on an indexed column then it will always be used during execution of the query.
- E. An index can be created as part of a CREATE TABLE statement.
- F. An UNUSABLE index is maintained when DML is performed on its underlying table.

**Suggested Answer:** BEF

Community vote distribution

ABE (100%)

🗳️ **TATSSSS** 9 months, 2 weeks ago

**Selected Answer:** ABE

C and F are wrong, oracle documentation says:

Unlike unusable indexes, an invisible index is maintained during DML statements.

<https://docs.oracle.com/en/database/oracle/oracle-database/19/admin/managing-indexes.html#GUID-3A66938F-73C6-4173-844E-3938A0DBBB54>

upvoted 2 times

🗳️ **lucemqy** 1 year, 1 month ago

**Selected Answer:** ABE

ABE correct

upvoted 1 times

🗳️ **ESZ** 1 year, 3 months ago

**Selected Answer:** ABE

ABE correct

upvoted 1 times

🗳️ **zouve** 1 year, 6 months ago

BEF for me

upvoted 1 times

🗳️ **dilshod** 1 year, 6 months ago

**Selected Answer:** ABE

These are correct

upvoted 2 times

🗳️ **Poke90** 1 year, 9 months ago

**Selected Answer:** ABE

invisible indexes are maintained, unusable indexes are not. C and F are wrong

upvoted 1 times

🗳️ **Poke90** 1 year, 9 months ago

F is wrong but C is correct:

[https://docs.oracle.com/cd/E18283\\_01/server.112/e17120/indexes004.htm](https://docs.oracle.com/cd/E18283_01/server.112/e17120/indexes004.htm)

so either ACE or BCE

upvoted 1 times

🗳️ **jr\_C** 1 year, 9 months ago

C is false, it is explained on the link you provided on the section: Understand When to Use Unusable or Invisible Indexes

upvoted 1 times

🗳️ **greenneem** 1 year, 10 months ago

F is wrong. it is not maintained by DML.

upvoted 2 times

  **Viviana3184** 2 years ago

ABE are correct

upvoted 2 times

  **Babi265** 2 years ago

Answer is ABE.

F is wrong, unused index is not mentained

upvoted 3 times

Examine this description of the EMP table:

Name	Null?	Type
-----	-----	-----
EMPNO	NOT NULL	NUMBER(4)
ENAME		VARCHAR2(10)
SAL		NUMBER(7,2)
DEPTNO		NUMBER(2)

You execute this query:

```
SELECT deptno AS "departments", SUM(sal) AS "salary"
  FROM emp
 GROUP BY 1
HAVING SUM(sal) > 3000;
```

What is the result?

- A. only departments where the total salary is greater than 3000, returned in no particular order
- B. only departments where the total salary is greater than 3000, ordered by department
- C. all departments and a sum of the salaries of employees with a salary greater than 3000
- D. an error

**Suggested Answer:** D

 **andrejshilin16** 11 months, 2 weeks ago


Answer D, since in GROUP BY clause must used column names instead positioned numbers like here, like GROUP BY deptno - works fine  
upvoted 1 times

 **zouve** 2 years ago

D  
ORA-00979: not a GROUP BY expression  
upvoted 1 times

 **katsoufis** 2 years, 3 months ago

why D? A is correct  
upvoted 1 times

 **Rik92** 2 years, 3 months ago

Because the Group by requires all items in the select clause  
upvoted 4 times

 **Rik92** 2 years, 3 months ago

Correction: One can order by 1 but you have to name the group by columns  
upvoted 2 times



Which two are true about virtual columns? (Choose two.)

- A. They can be indexed.
- B. They can be referenced in the column expression of another virtual column.
- C. They cannot have a data type explicitly specified.
- D. They can be referenced in the set clause of an update statement as the name of the column to be updated.
- E. They can be referenced in the where clause of an update or delete statement.

**Suggested Answer: AE**

Community vote distribution

AE (100%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer: AE**

AE is the right answer.

upvoted 1 times

🗳️ 👤 **salander71** 9 months ago

Correct answer is: AB

A. Virtual columns can be indexed just like regular columns. This is useful when queries frequently filter or sort data based on the virtual column, as indexing improves performance.

B. Virtual columns can be used in the expression of another virtual column. This allows for more complex derived data to be generated from other columns, including other virtual columns.

upvoted 1 times

🗳️ 👤 **salander71** 9 months ago

E is false. Although virtual columns can be used in a WHERE clause of queries to filter data, they cannot be used in the WHERE clause of an UPDATE or DELETE statement because these statements modify data, and virtual columns are not physically stored.

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: AE**

AE They can be indexed and used in where clause of UPDATE and DELETE statements

upvoted 1 times

🗳️ 👤 **Poke90** 2 years, 3 months ago

<https://oracle-base.com/articles/11g/virtual-columns-11gr1>

upvoted 4 times

🗳️ 👤 **alic\_alex** 2 years, 3 months ago

**Selected Answer: AE**

more details about virtual columns: <https://www.gpsos.es/2021/03/virtual-columns-in-oracle-use-and-limitations/?lang=en>

upvoted 2 times

A session's NLS\_DATE\_FORMAT is set to DD Mon YYYY.

Which two queries return the value 1 Jan 2019? (Choose two.)

- A. SELECT TO\_DATE('2019-01-01') FROM DUAL;
- B. SELECT DATE '2019-01-01' FROM DUAL;
- C. SELECT '2019-01-01' FROM DUAL;
- D. SELECT TO\_DATE('2019-01-01', 'YYYY-MM-DD') FROM DUAL;
- E. SELECT TO\_CHAR('2019-01-01') FROM DUAL;

Suggested Answer: *BD*

🗲 👤 **2339ac0** 1 month, 3 weeks ago

Selected Answer: **BD**

BD is right answer due to internal nls date format and explicit format parameter.  
upvoted 1 times

🗲 👤 **kay000001** 8 months, 1 week ago

Selected Answer: **BD**

Answers are: B & D.  
upvoted 1 times

🗲 👤 **lucemqy** 1 year, 7 months ago

BD Works  
upvoted 1 times

🗲 👤 **Festo** 1 year, 8 months ago

BD IS CORRECT  
upvoted 1 times

🗲 👤 **TheOracleWasTaken** 1 year, 11 months ago

B wouldnt work right bc it doesnt match the nls\_parameter thing?.  
D works tho  
upvoted 2 times

🗲 👤 **lucemqy** 1 year, 7 months ago

B would work the date displayed will be the default NLS\_DATE\_FOMAT set  
upvoted 1 times

Examine this SQL statement:

```
DELETE FROM employees e
WHERE EXISTS
  (SELECT 'dummy'
   FROM emp_history
   WHERE employee_id = e.employee_id);
```

Which two are true? (Choose two.)

- A. The DELETE statement executes successfully even if the subquery selects multiple rows.
- B. The subquery is executed before the DELETE statement is executed.
- C. The subquery is not a correlated subquery.
- D. All existing rows in the EMPLOYEES table are deleted.
- E. The subquery is executed for every row in the EMPLOYEES table.

**Suggested Answer: AE**

Community vote distribution

AE (100%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer: AE**

AE is the correct answer

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 1 month ago

**Selected Answer: AE**

AE is correct

upvoted 1 times

🗳️ 👤 **zouve** 1 year, 5 months ago

**Selected Answer: AE**

A. The DELETE statement executes successfully even if the subquery selects multiple rows. This is because the EXISTS condition only checks for the existence of at least one row in the subquery, regardless of how many rows are returned.

E. The subquery is executed for every row in the EMPLOYEES table. This is because the subquery is a correlated subquery, which means that it references a column from the outer query (e.employee\_id). As a result, the subquery must be executed once for each row in the EMPLOYEES table to determine whether the EXISTS condition is true or false for that row.

upvoted 2 times

🗳️ 👤 **zouve** 1 year, 5 months ago

in other words the subquery is evaluated once for each row selected by the outer query

upvoted 1 times

🗳️ 👤 **karols32** 1 year, 7 months ago

A,D.

Subquery isn't query before outer query.

upvoted 1 times

🗳️ 👤 **auwia** 1 year, 8 months ago

AB are correct.

upvoted 1 times



🗳️ 👤 **Poke90** 1 year, 9 months ago

Why is B wrong?

upvoted 2 times

🗳️ 👤 **Tof0ra** 10 months, 4 weeks ago

Understand that one thing is the execution of the DELETE command other is the deletion of the row.  
upvoted 1 times

  **Tof0ra** 10 months, 4 weeks ago

The DELETE starts running and then it calls the subquery.  
upvoted 1 times

Examine this constraint information:

TABLE NAME	REFERENCE CONSTRAINT NAME	COLUMN NAME	CONSTRAINT TYPE	CONSTRAINT NAME	SEARCH CONDITION
DEPT	CC_DEPT	DEPTNO	C		deptno > 9
DEPT	SYS_C0012476	DNAME	C		"DNAME" IS NOT NULL
DEPT	SYS_C0012478	DEPTNO	P		
EMP	CC_COMM	COMMISSION	C		commission < salary
EMP	CC_COMM	SALARY	C		commission < salary
EMP	CC_EMPNO	EMPNO	C		empno > 10
EMP	CC_SAL	SALARY	C		salary > 1000
EMP	CR_DEPT	DEPTNO	R	SYS_C0012478	
EMP	CR_MGR	MANAGER	R	SYS_C0012484	
EMP	SYS_C0012479	ENAME	C		"ENAME" IS NOT NULL
EMP	SYS_C0012480	JOB	C		"JOB" IS NOT NULL
EMP	SYS_C0012484	EMPNO	P		

Which three statements are true? (Choose three.)

- A. The SALARY column must have a value.
- B. The DEPTNO column in the EMP table can contain NULLS.
- C. The COMMISSION column can contain negative values.
- D. The DEPTNO column in the EMP table can contain the value 1.
- E. The MANAGER column is a foreign key referencing the EMPNO column.
- F. The DNAME column has a unique constraint.
- G. An index is created automatically in the MANAGER column.

**Suggested Answer: BDE**

Community vote distribution



**Viviana3184** Highly Voted 2 years, 6 months ago

i think BCE are correct  
upvoted 17 times

**bfb7c7d** Most Recent 8 months ago

As per the rule Primary and Composite Key can not contain Null value Then why B option is right  
I think A is right  
upvoted 1 times

**rani\_12** 10 months, 3 weeks ago

but how did we know that the manager is referencing the empno column ?!  
upvoted 3 times

**ArslanAltaf** 1 year, 2 months ago

BCE is Correct

- A. The SALARY column must have a value. ....FALSE
- B. The DEPTNO column in the EMP table can contain NULLS.....TRUE
- C. The COMMISSION column can contain negative values. ....TRUE
- D. The DEPTNO column in the EMP table can contain the value 1..... TRUE WITH CONDITION. (If dept table primary key has 1 values)
- E. The MANAGER column is a foreign key referencing the EMPNO column..... TRUE
- F. The DNAME column has a unique constraint. FALSE
- G. An index is created automatically in the MANAGER column....FALSE

upvoted 2 times

**RIDA\_007** 1 year, 1 month ago

A is correct!! The SALARY column is part of a Check constraint CC\_COMM that requires the SALARY to be greater than 100, which means it must have a value and cannot be NULL

upvoted 1 times

🗨️ 👤 **ArslanAltaf** 1 year ago

Its correct there is a check constraint which mean it has to be  $> 100$ , yet it will not prevent a null value to be inserted in the column. In order to prevent not null, this specific constraint has to be added. Test it. A is wrong

upvoted 1 times

🗨️ 👤 **yaya32** 1 year, 5 months ago

**Selected Answer: CDE**

I think B is wrong : The DEPTNO column in the EMP table can contain NULLS no because it refers to a PK so not possible.

upvoted 2 times

🗨️ 👤 **Nico\_Nikko** 1 year, 3 months ago

NULL is actually the only exception. AN FK Column can have NULL Values even though the PK Column it is referencing does not.

upvoted 2 times

🗨️ 👤 **yanoolthecool** 1 year, 6 months ago

**Selected Answer: BCE**

BCE, D is wrong because DEP\_NO has a constraint in the department table that prevents it from having a value less than 10, very tricky, those who answered D might get surprises after the exam.

upvoted 1 times

🗨️ 👤 **winston** 1 year, 6 months ago

BDE Is correct but why is C wrong?

upvoted 1 times

🗨️ 👤 **Mcromeo** 1 year, 6 months ago

**Selected Answer: BCE**

BCE is correct

upvoted 1 times

🗨️ 👤 **zangado** 1 year, 7 months ago

**Selected Answer: BCE**

BCE are correct.

D is wrong because DEPTNO in EMP table is referencign DEPTNO in DEPT and there is check constraint which demands DEPTNO in DEPT to be greater than 9

upvoted 1 times

🗨️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: BDE**

BDE is the correct answer

upvoted 1 times

🗨️ 👤 **Mcromeo** 1 year, 6 months ago

D is not correct bcz deptno  $> 9$

upvoted 1 times

🗨️ 👤 **ramonciks** 1 year, 9 months ago

why E ?

upvoted 1 times

🗨️ 👤 **Sangiii** 1 year, 10 months ago

BCE - is the answer

upvoted 1 times

🗨️ 👤 **TheOracleWasTaken** 1 year, 11 months ago

**Selected Answer: CE**

deptno is a pk so cannot be null

upvoted 1 times

🗨️ 👤 **Tof0ra** 1 year, 4 months ago

It's a PK in dept not in emp

upvoted 1 times

🗨️ 👤 **CMjer** 2 years, 1 month ago

**Selected Answer: BCE**

BCE seems correct


upvoted 1 times

  **AV80** 2 years, 6 months ago

**Selected Answer: CDE**

CDE seems to be right option

upvoted 1 times

  **anzac** 2 years, 6 months ago

D is not correct because DEPTNO is foreign key referencing DEPTNO in DEPT table and there DEPTNO must > 9

upvoted 2 times

  **lucemqy** 1 year, 7 months ago

DEPTNO is a foreign key referencing a primary key so it should be able to be 1

upvoted 1 times

Which two are true about creating tables in an Oracle database? (Choose two.)

- A. Creating an external table will automatically create a file using the specified directory and file name.
- B. A system privilege is required.
- C. The same table name can be used for tables in different schemas.
- D. A primary key constraint is mandatory.
- E. A CREATE TABLE statement can specify the maximum number of rows the table will contain.

**Suggested Answer:** BC

Community vote distribution

BC (83%)

AC (17%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer:** BC

BC is right answer.

upvoted 1 times

🗳️ 👤 **Thameur01** 5 months ago

**Selected Answer:** BC

BC is correct

upvoted 1 times

🗳️ 👤 **hmatinnn** 10 months, 2 weeks ago

**Selected Answer:** BC

Should be BC

upvoted 1 times

🗳️ 👤 **yanoolthecool** 1 year ago

**Selected Answer:** BC

BC cuz google says:

'System privileges allow a user to perform a particular database operation or class of database operations. For example, to create a table, the user needs the create table privilege.'

'Objects have privileges associated with them, such as insert, update and delete a table.'

upvoted 1 times

🗳️ 👤 **zangado** 1 year ago

**Selected Answer:** BC

To create a relational table in your own schema, you must have the CREATE TABLE system privilege. To create a table in another user's schema, you must have the CREATE ANY TABLE system privilege.

A) is wrong because external table creates a file only if it's created with 'CREATE AS SELECT' statement

upvoted 3 times

🗳️ 👤 **lucemqy** 1 year, 1 month ago

**Selected Answer:** AC

Should be AC since CREATE TABLE does not require system privilege just object privilege. CREATE ANY TABLE require system privilege

upvoted 1 times

🗳️ 👤 **yanoolthecool** 1 year ago

'System privileges allow a user to perform a particular database operation or class of database operations. For example, to create a table, the user needs the create table privilege.'

'Objects have privileges associated with them, such as insert, update and delete a table.'

upvoted 1 times



Examine this partial statement:

```
SELECT ename, sal, comm FROM emp
```

Now examine this output:

ENAME	SAL	COMM
MARTIN	1250	1400
WARD	1250	500
ALLEN	1600	300
TURNER	1500	0
ADAMS	1100	
BLAKE	2850	
CLARK	2450	
FORD	3000	
JAMES	950	
JONES	2975	
KING	5000	
MILLER	1300	
SCOTT	3000	
SMITH	800	

Which ORDER BY clause will generate the displayed output?

- A. ORDER BY comm DESC NULLS LAST, ename
- B. ORDER BY NVL(comm, 0) ASC NULLS FIRST, ename
- C. ORDER BY NVL(comm, 0) ASC NULLS LAST, ename
- D. ORDER BY NVL(comm, 0) DESC, ename

**Suggested Answer: A**

Community vote distribution

A (100%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer: A**

A is the right answer.

upvoted 1 times

🗳️ 👤 **zzz02** 9 months, 1 week ago

**Selected Answer: D**

I'd give it D

upvoted 1 times

🗳️ 👤 **9ramone7** 1 year, 2 months ago

Tested on Oracle live sql.

A is correct query for the output

B comm nulls are at the top - wrong

C Turner, Allen, Ward, Martin at the bottom - wrong

D Only Turner at the bottom - wrong

upvoted 1 times

🗳️ 👤 **yaya32** 1 year, 5 months ago

**Selected Answer: A**

A for me

upvoted 1 times

🗳️ 👤 **Ramunele** 1 year, 9 months ago

D is not correct, because Turner is last in the table. A is correct answer

upvoted 3 times

🗳️ 👤 **sunilrajusb** 1 year, 10 months ago

D also correct answer,, Tested and its giving expected output

upvoted 3 times

  **Misi\_Oracle** 1 year, 1 month ago

NVL lets you replace null (returned as a blank) with a string in the results of a query. If expr1 is null, then NVL returns expr2. If expr1 is not null, then NVL returns expr1. Meaning that the output would be replaced by 0. The output still remains as null.

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/NVL.html>

upvoted 1 times

Examine the description of the CUSTOMERS table:

CUSTOMER_ID	CUSTOMER_NAME
10	MARK
20	Mandy
30	Mary
40	MARVIN
50	MARTIN

Which two SELECT statements will return these results: (Choose two.)

CUSTOMER_NAME
Mandy
Mary

- A. SELECT customer\_name FROM customers WHERE UPPER(customer\_name) LIKE 'MA\*';
- B. SELECT customer\_name FROM customers WHERE customer\_name = '\*Ma\*';
- C. SELECT customer\_name FROM customers WHERE customer\_name LIKE 'Ma\*';
- D. SELECT customer\_name FROM customers WHERE UPPER(customer\_name) LIKE 'MA%';
- E. SELECT customer\_name FROM customers WHERE customer\_name LIKE '%a%';
- F. SELECT customer\_name FROM customers WHERE customer\_name LIKE 'Ma%';
- G. SELECT customer\_name FROM customers WHERE customer\_name LIKE '\*Ma\*';

**Suggested Answer:** EF

Community vote distribution

EF (100%)

826cb95 1 month, 4 weeks ago

**Selected Answer:** F

F -for sure

C - is second closest if it was customer\_name LIKE 'Ma%' but not when it is 'Ma\*' as shown  
upvoted 1 times

yaya32 11 months, 2 weeks ago

**Selected Answer:** EF

EF is correct

upvoted 1 times

lucemqy 1 year, 1 month ago

**Selected Answer:** EF

EF since like is case sensitive  
upvoted 1 times

Sangiii 1 year, 4 months ago

EF is correct

upvoted 2 times

Orxan\_H 1 year, 6 months ago

EF are correct

upvoted 1 times



CMjer 1 year, 7 months ago

**Selected Answer:** EF

D is false - it will return all names (checked with SQL Developer)  
upvoted 2 times

Cecilia\_Soso 1 year, 8 months ago

D and F, %a% will return all names with letter a  
upvoted 1 times

  **cadcadley** 1 year, 11 months ago  
EF is correct  
upvoted 4 times

The PRODUCT\_INFORMATION table has a UNIT\_PRICE column of data type NUMBER(8,2).

Evaluate this SQL statement:

```
SELECT TO_CHAR(unit_price, '$9,999') FROM product_information;
```

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

- A. A row whose UNIT\_PRICE column contains the value 10235.95 will be displayed as \$1,0236.
- B. A row whose UNIT\_PRICE column contains the value 1023.95 will be displayed as \$1,024.
- C. A row whose UNIT\_PRICE column contains the value 10235.95 will be displayed as \$1,023.
- D. A row whose UNIT\_PRICE column contains the value 10235.95 will be displayed as #####.
- E. A row whose UNIT\_PRICE column contains the value 1023.99 will be displayed as \$1,023.

**Suggested Answer:** BD

Community vote distribution

BD (100%)

🗨️ 👤 **yaya32** 11 months, 2 weeks ago

**Selected Answer: BD**

BD is correct

upvoted 1 times

🗨️ 👤 **lucemqy** 1 year, 1 month ago

**Selected Answer: BD**

BD is the correct answer

upvoted 1 times

🗨️ 👤 **ESZ** 1 year, 9 months ago

Not sure why D is correct. Could someone explain please. Thanks

upvoted 1 times

🗨️ 👤 **MoonLight** 1 year, 8 months ago

You have to change the format to '\$99,999'


upvoted 2 times

Which two statements are true about Oracle databases and SQL? (Choose two.)

- A. Updates performed by a database user can be rolled back by another user by using the ROLLBACK command.
- B. A query can access only tables within the same schema.
- C. The database guarantees read consistency at select level on user-created tables.
- D. A user can be the owner of multiple schemas in the same database.
- E. When you execute an update statement, the database instance locks each updated row.

**Suggested Answer:** CE

Community vote distribution



**2339ac0** 1 month, 3 weeks ago

**Selected Answer:** CE

CE is right answer.

upvoted 1 times

**archit4321** 8 months ago

**Selected Answer:** CE

C and E are the most accurate

upvoted 2 times

**billysunday1** 9 months, 1 week ago

**Selected Answer:** CE

C and E.

[https://docs.oracle.com/cd/B28359\\_01/server.111/b28318/schema.htm#CNCPT111](https://docs.oracle.com/cd/B28359_01/server.111/b28318/schema.htm#CNCPT111)

A schema is a collection of logical structures of data, or schema objects. A schema is owned by a database user and has the same name as that user. Each user owns a single schema. Schema objects can be created and manipulated with SQL and include the following types of objects:

upvoted 1 times

**billysunday1** 9 months, 1 week ago

**Selected Answer:** CD

Answer should be C and D.

C is ACID which Oracle SQL always do

[https://docs.oracle.com/cd/B13789\\_01/server.101/b10759/statements\\_8003.htm](https://docs.oracle.com/cd/B13789_01/server.101/b10759/statements_8003.htm)

```
CREATE USER my_user IDENTIFIED BY my_password DEFAULT TABLESPACE tbspace1 QUOTA UNLIMITED ON tbspace1;
```

```
GRANT schema1, schema2 TO my_user;
```

upvoted 1 times

**ogi33** 9 months ago

GRANT schema1, schema2 TO my\_user make error , can grant only role or privileges

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/GRANT.html#GUID-20B4E2C0-A7F8-4BC8-A5E8-BE61BDC41AC3>

upvoted 1 times

Which statement is true about TRUNCATE and DELETE?

- A. For tables with multiple indexes and triggers, DELETE is faster than TRUNCATE.
- B. You can never TRUNCATE a table if foreign key constraints would be violated.
- C. You can DELETE rows from a table with referential integrity constraints.
- D. For large tables, DELETE is faster than TRUNCATE.

**Suggested Answer: D**

Community vote distribution

B (71%)

C (29%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer: B**

B is right answer.

upvoted 1 times

🗳️ 👤 **Bissto** 10 months ago

**Selected Answer: B**

B is correct

upvoted 1 times

🗳️ 👤 **billysunday1** 1 year, 3 months ago

**Selected Answer: B**

Answer is B

ORA-02266: unique/primary keys in table referenced by enabled foreign keys

upvoted 1 times

🗳️ 👤 **hmatinnn** 1 year, 4 months ago

**Selected Answer: B**

So answer should be B

upvoted 1 times

🗳️ 👤 **hmatinnn** 1 year, 4 months ago

The answer is B

Here is explanation : In this option it does not mean you can never truncate a table

it says if foreign key activated and would be violated you can NEVER. The keyword is IF <<<<

upvoted 1 times

🗳️ 👤 **yanoolthecool** 1 year, 6 months ago

**Selected Answer: C**

C CAN work if we have CASCADE, B is wrong cuz it says NEVER, the documentation states:

'You cannot truncate the parent table of an enabled foreign key constraint. You must disable the constraint before truncating the table. An exception is that you can truncate the table if the integrity constraint is self-referential.'

upvoted 1 times

🗳️ 👤 **yanoolthecool** 1 year, 6 months ago

there's also this:

CASCADE

If you specify CASCADE, then Oracle Database truncates all child tables that reference table with an enabled ON DELETE CASCADE referential constraint. This is a recursive operation that will truncate all child tables, grandchild tables, and so on, using the specified options.

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: C**

C is correct since you can delete using cascade option

B is incorrect since it says never but you can by disable the constraint

upvoted 3 times

🗳️ 👤 **Mcromeo** 1 year, 8 months ago

**Selected Answer: B**

B is correct. Truncate is faster than delete

upvoted 1 times

🗳️ 👤 **jm9999** 1 year, 9 months ago

Very confusing. It seems truncate is always faster than delete. And for B and C, B is false if you use cascade in the command and constraint was created with on delete cascade. Similarly with C, if constraint was created with on delete cascade then you can delete those rows.

upvoted 1 times

🗳️ 👤 **jm9999** 1 year, 8 months ago

The problem with B is that there is an exception for self referential constraint and B uses the word 'never'. I think it's C because of the 'on delete cascade option'.

upvoted 1 times

🗳️ 👤 **Roneros** 1 year, 9 months ago

**Selected Answer: B**

You cannot truncate the parent table of an enabled foreign key constraint. You must disable the constraint before truncating the table.

<https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/TRUNCATE-TABLE.html>

upvoted 1 times

🗳️ 👤 **TofOra** 1 year, 4 months ago

From the link: You cannot truncate the parent table of an enabled foreign key constraint. You must disable the constraint before truncating the table. An exception is that you can truncate the table if the integrity constraint is self-referential.

upvoted 1 times

🗳️ 👤 **Orxan\_H** 1 year, 10 months ago

**Selected Answer: B**

B is correct

upvoted 2 times

🗳️ 👤 **Shrimathi** 1 year, 11 months ago

C is correct,

We can truncate table contains referential integrity, one by drop the constraint/ while adding constraint we can include on delete cascade . Truncate table colors cascade; this statement will run successfully.

Reference <https://www.examttopics.com/discussions/oracle/view/15064-exam-1z0-071-topic-1-question-268-discussion/>

upvoted 2 times

🗳️ 👤 **ESZ** 1 year, 11 months ago

This needs to be updated!

upvoted 1 times

🗳️ 👤 **zouve** 2 years ago

Option B is true. When using the TRUNCATE statement, if there are foreign key constraints referencing the table to be truncated, you will not be able to truncate the table. Truncating a table removes all the data in the table, and if there are dependent foreign key constraints, truncating the table would violate those constraints. In such cases, you would need to either disable or drop the foreign key constraints before truncating the table.

upvoted 1 times

🗳️ 👤 **TofOra** 1 year, 4 months ago

It says "NEVER", but is possible for self referenced tables.

upvoted 1 times

🗳️ 👤 **Cecilia\_Soso** 2 years, 2 months ago

B -You cannot truncate the parent table of an enabled foreign key constraint. You must disable the constraint before truncating the table. An exception is that you can truncate the table if the integrity constraint is self-referential.

upvoted 1 times

🗳️ 👤 **aamm74200** 2 years, 5 months ago





C.

A and D: False --> Truncate is better and has better performance than delete and Drop except for very small tables.

B False --> You can truncate a table with a FK on Cascade. You can disable FK and then Truncate (except if FK is self-referential).

upvoted 4 times

  **jfc1** 2 years, 5 months ago

**Selected Answer: B**

D is wrong . B is correct

[https://docs.oracle.com/database/121/SQLRF/statements\\_10007.htm#SQLRF01707](https://docs.oracle.com/database/121/SQLRF/statements_10007.htm#SQLRF01707)

upvoted 2 times

Which two statements are true? (Choose two.)

- A. CASE is a function and DECODE is not.
- B. Neither CASE nor DECODE is a function.
- C. All conditions evaluated using CASE can also be evaluated using DECODE.
- D. All conditions evaluated using DECODE can also be evaluated using CASE.
- E. DECODE is a function and CASE is not.
- F. Both CASE and DECODE are functions.

**Suggested Answer:** DF

Community vote distribution

DE (100%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer: DE**

DE is right answer.

upvoted 1 times

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

**Selected Answer: DE**

DE

Case is a statement DECODE is a function

upvoted 2 times

🗳️ 👤 **WingL** 11 months, 3 weeks ago

DE are correct.

CASE is a expression not a function.

[https://docs.oracle.com/cd/B13789\\_01/appdev.101/b10807/13\\_elems004.htm#:~:text=The%20CASE%20statement%20evaluates%20a,first%20one%20that%](https://docs.oracle.com/cd/B13789_01/appdev.101/b10807/13_elems004.htm#:~:text=The%20CASE%20statement%20evaluates%20a,first%20one%20that%20)

upvoted 1 times

🗳️ 👤 **zouve** 1 year ago

case is statement while decode is a function

upvoted 2 times

🗳️ 👤 **dexdinh91** 1 year, 5 months ago

**Selected Answer: DE**

DE are correct

upvoted 3 times

🗳️ 👤 **Viviana3184** 1 year, 6 months ago

DE are correct

upvoted 1 times

🗳️ 👤 **Babi265** 1 year, 6 months ago

Should be DE ?

upvoted 2 times

Examine these statements executed in a single Oracle session:

```
CREATE TABLE product (pcode NUMBER(2), pname VARCHAR2(20));
INSERT INTO product VALUES (1, 'pen');
INSERT INTO product VALUES (2, 'pencil');
INSERT INTO product VALUES (3, 'fountain pen');
SAVEPOINT a;
UPDATE product SET pcode = 10 WHERE pcode = 1;
COMMIT;
DELETE FROM product WHERE pcode = 2;
SAVEPOINT b;
UPDATE product SET pcode = 30 WHERE pcode = 3;
SAVEPOINT c;
DELETE FROM product WHERE pcode = 10;
ROLLBACK TO SAVEPOINT b;
COMMIT;
```

Which three statements are true? (Choose three.)

- A. The code for pen is 10.
- B. There is no row containing fountain pen.
- C. There is no row containing pen.
- D. There is no row containing pencil.
- E. The code for fountain pen is 3.
- F. The code for pen is 1.

**Suggested Answer:** ADE

*Community vote distribution*

ADE (100%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer:** ADE

ADE is the right answer.

upvoted 1 times

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

**Selected Answer:** ADE

ADE is the correct answer

upvoted 3 times

Which is true about the & and && prefixes with substitution variables? (Choose all that apply.)

- A. Both & and && can prefix a substitution variable name in queries and DML statements.
- B. An & prefix to an undefined substitution variable, which is referenced twice in the same query, will prompt for a value twice.
- C. & can prefix a substitution variable name only in queries.
- D. An && prefix to an undefined substitution variable, which is referenced multiple times in multiple queries, will prompt for a value once per query.
- E. The && prefix will not prompt for a value even if the substitution variable is not previously defined in the session.

**Suggested Answer: AB**

Community vote distribution

AB (100%)

🗳️ 👤 **yanoolthecool** 1 year ago

**Selected Answer: AB**

D is not correct because it will prompt the user once, and only once for all queries, not once for every query.

upvoted 1 times

🗳️ 👤 **mrcsp** 1 year, 4 months ago

D too

select &&a, &&a from dual;

select &&a from dual;

upvoted 2 times

🗳️ 👤 **Oraclestd24** 1 year, 3 months ago

select &&a col1, &&a col2 from dual;

select &&a col3 from dual;

D can't be answer , because it says " will prompt for a value once per query." but that is not the case.

upvoted 1 times

🗳️ 👤 **lucemqy** 1 year, 7 months ago

**Selected Answer: AB**

AB is correct

<https://docs.oracle.com/en/database/oracle/oracle-database/21/sqpug/using-substitution-variables-sqlplus.html#GUID-0BEEC1D7-876B-495C-9327-17037652D3D2>

upvoted 1 times

🗳️ 👤 **WingL** 1 year, 11 months ago

A.B correct.

reference for B

If a single ampersand prefix is used with an undefined variable, the value you enter at the prompt is not stored. Immediately after the value is substituted in the statement the variable is discarded and remains undefined. If the variable is referenced twice, even in the same statement, then you are prompted twice.

upvoted 1 times

Which statement will return a comma-separated list of employee names in alphabetical order for each department in the EMP table?

- A. 

```
SELECT deptno, LISTAGG(ename, ',') WITHIN GROUP (ORDER BY ename) AS employee_list
FROM emp
GROUP BY deptno;
```
- B. 

```
SELECT deptno, LISTAGG(ename, ',') WITHIN GROUP (GROUP BY deptno) AS employee_list
FROM emp
ORDER BY ename;
```
- C. 

```
SELECT deptno, LISTAGG(ename, ',') WITHIN GROUP AS employee_list
FROM emp
GROUP BY deptno;
```
- D. 

```
SELECT deptno, LISTAGG(ename, ',') WITHIN GROUP AS employee_list
FROM emp
GROUP BY deptno
ORDER BY ename;
```

**Suggested Answer: A**

Community vote distribution

A (100%)

🗳️ 👤 **2339ac0** 1 month, 3 weeks ago

**Selected Answer: A**

A is right answer.

upvoted 1 times

🗳️ 👤 **ArslanAltaf** 6 months, 3 weeks ago

A Correct

select department\_id, listagg(last\_name, ',') within group (order by last\_name) from employees group by department\_id;

upvoted 1 times

🗳️ 👤 **ArslanAltaf** 8 months, 1 week ago

B, C, D will throw error.

A is the only correct option

upvoted 1 times

🗳️ 👤 **amizh** 10 months, 2 weeks ago

**Selected Answer: A**

tried it in sql developer

upvoted 1 times

Examine the data in the COLORS table:

RGB_HEX_VALUE	COLOR_NAME
FF0000	red
00FF00	green
0000FF	blue

Examine the data in the BRICKS table:


BRICK_ID	COLOR_RGB_HEX_VALUE
1	FF0000
2	00FF00
3	FFFFFF

Which two queries return all the rows from COLORS? (Choose two.)

- A. `SELECT *  
FROM bricks b  
RIGHT JOIN colors c  
ON b.color_rgb_hex_value = c.rgb_hex_value;`
- B. `SELECT *  
FROM bricks b  
FULL JOIN colors c  
ON b.color_rgb_hex_value = c.rgb_hex_value;`
- C. `SELECT *  
FROM colors c  
LEFT JOIN bricks b  
USING (rgb_hex_value);`
- D. `SELECT *  
FROM colors c  
LEFT JOIN bricks b  
ON b.color_rgb_hex_value = c.rgb_hex_value  
WHERE b.brick_id > 0;`
- E. `SELECT *  
FROM bricks b  
JOIN colors c  
ON b.color_rgb_hex_value = c.rgb_hex_value;`

**Suggested Answer:** DE

Community vote distribution



**anzac** Highly Voted 2 years, 6 months ago

Checked: A, B

Check:

```
CREATE TABLE COLORS_120
(
  RGB_HEX_VALUE VARCHAR2(100)
  , COLOR_NAME VARCHAR2(100)
);
```

```
CREATE TABLE BRIKS_120
```

```
(  
BRICK_ID NUMBER  
, COLOR_RGB_HEX_VALUE VARCHAR2(100)  
);
```

```
INSERT INTO COLORS_120  
SELECT 'FF0000','red' FROM DUAL  
UNION ALL  
SELECT '00FF00','green' FROM DUAL  
UNION ALL  
SELECT '0000FF','blue' FROM DUAL;
```

```
INSERT INTO BRIKS_120  
SELECT 1,'FF0000' FROM DUAL  
UNION ALL  
SELECT 2,'00FF00' FROM DUAL  
UNION ALL  
SELECT 3,'FFFFFF' FROM DUAL;
```

```
A;  
SELECT *  
FROM BRIKS_120 b  
RIGHT JOIN COLORS_120 c  
ON b.COLOR_RGB_HEX_VALUE = c.RGB_HEX_VALUE;
```



```
B;  
SELECT *  
FROM BRIKS_120 b  
FULL JOIN COLORS_120 c  
ON b.COLOR_RGB_HEX_VALUE = c.RGB_HEX_VALUE;
```

```
C;  
SELECT *  
FROM c  
FULL JOIN BRIKS_120 b  
USING(RGB_HEX_VALUE);
```

```
D;  
SELECT *  
FROM COLORS_120 c  
LEFT JOIN BRIKS_120 b  
ON b.COLOR_RGB_HEX_VALUE = c.RGB_HEX_VALUE  
WHERE b.brick_id > 0
```

;

```
E;  
SELECT *  
FROM BRIKS_120 b  
LEFT JOIN COLORS_120 c  
ON b.COLOR_RGB_HEX_VALUE = c.RGB_HEX_VALUE;  
upvoted 13 times
```

  **WingL** 1 year, 11 months ago  
thanks for coding.  
upvoted 3 times

  **2339ac0** Most Recent 1 month, 3 weeks ago