

510-015

Sybase

ASE 12.0 Server Administration Professional

OfficialCerts.com is a reputable IT certification examination guide, study guides and audio exam provider. We ensure that you pass your 510-015 exam in first attempt and also get high scores to acquire Sybase certification.

If you use OfficialCerts 510-015 Certification questions and answers, you will experience actual 510-015 exam questions/answers. We know exactly what is needed and have all the exam preparation material required to pass the exam. Our Sybase exam prep covers over 95% of the questions and answers that may be appeared in your 510-015 exam. Every point from pass4sure 510-015 PDF, 510-015 review will help you take Sybase 510-015 exam much easier and become Sybase certified.

Here's what you can expect from the OfficialCerts Sybase 510-015 course:

- * Up-to-Date Sybase 510-015 questions as experienced in the real exam.*
- * 100% correct Sybase 510-015 answers you simply can't find in other 510-015 courses.*
- * All of our tests are easy to download. Your file will be saved as a 510-015 PDF.*
- * Sybase 510-015 brain dump free content featuring the real 510-015 test questions.*

Sybase 510-015 certification exam is of core importance both in your Professional life and Sybase certification path. With Sybase certification you can get a good job easily in the market and get on your path for success. Professionals who passed Sybase 510-015 exam training are an absolute favorite in the industry. You will pass Sybase 510-015 certification test and career opportunities will be open for you.

<http://janetdailey.com/?certs=exams.asp?examcode=510-015>



QUESTION 1

Given

The employees table has an index created as follows:

```
create index emp_index1  
on employees(last_name,first_name desc)
```

Which of the following queries would not require a sort? (Choose 2)

- A. Select * from employees order by last_name, first_name
- B. Select * from employees order by last_name ASC, first_name DESC
- C. Select * from employees order by first_name DESC, last_name ASC
- D. Select * from employees order by last_name DESC, first_name ASC

Answer: B , D

QUESTION 2

Given

number of worker processes = 10

parallel degree =5

max scan parallel degree = 4.

Table A has the following characteristics:

Datarows locking, No indexes, a column named price,5000 rows, 4 partitions.

A user executes the following query. Select* from A where price <\$10.

Which of the following conditions would cause the query to run in serial instead of parallel? (choose 2)

- A. there are not enough worker processes available at run time
- B. the database option select into/bulkcopy/pllsort is not set to true
- C. the partition skew is 2.3
- D. the database is set for single user mode
- E. max parallel degrees does not match the number of partitions

Answer: A , C

QUESTION 3

Which of the following statements describes readpast locking?

(Choose 2)

- A. It can be specified at the session, transaction, and table level.
- B. It allows insert, update, and delete commands to read past any incompatible lock.
- C. it allows readers to not block writers.
- D. It allows select and readtext queries to silently skip all rows or pageslocked with incompatible locks.

Answer: B , D

QUESTION 4

The ACME company has an 8 CPU SMP system running Adaptive Enterprise (ASE). The ASE server has been allocated 1.5 GB of 2.0 GB of memory available. The ASE server has been allocated 2 engines. There are 1,000 concurrent users of the ASE server.

The sp_sysmon report shows the spinlock contention rate averages 70-80% for the default data cache.

Which steps can be taken to correct this performance problem?

- A. Add a 16KB buffer pool to the server.
- B. Increase the size of the tempdb database.
- C. Add two engines to the ASE server.
- D. Add named caches to and bind the appropriate objects to the named caches.
- E. Add more locks to the server.

Answer: D

QUESTION 5

What is the purpose of dsync flag?

- A. to synchronize primary and mirrored devices.
- B. to synchronize primary and secondary databases.
- C. to control whether writes to operating system files are buffered or not.
- D. to control Backup Server's access to a dump device.

Answer: c

QUESTION 6

All of the following execution class associations exist in the ASE when user FRED, running application ctisql, logs in. Which one will be applied to FRED's session?

- A. sp_bindexclass ctisql,AP,NULL,EC1
- B. sp_bindexclass ctisql,AP,FRED,EC2
- C. sp_bindexclass FRED,LG,NULL,EC1
- D. sp_bindexclass FRED,LG,ctisql,EC2
- E. It depends on the order of the creations of the bindings

Answer: B

QUESTION 7

The reorg command may be used to:(Choose 2)

- A. Compact a table that uses datarows locking
- B. Reclaim unused space in a table that uses allpages locking
- C. Reclaim unused space in an index that uses datapages locking

D. Rebuild a table that uses allpages locking

Answer: A , C

QUESTION 8

What are some ways to reduce last data page lock contention for inserts?(Choose 3)

- A. alter the table to use the Datarows locking scheme
- B. alter the table to use the Datapages locking scheme
- C. partition the table
- D. create a clustered index on a random key
- E. increase the server configuration for the number of locks
- F.increase the server configuration for the deadlock checking peroid

Answer: A, C , D

QUESTION 9

The "abstract plan language" is:

- A. A language used to tell the optimizer how to process a query
- B. The internal command language used by sp_sysmon
- C. A new type of command language for creating very complicated queries
- D. The language used internally by the server for planning the most efficient use of memory

Answer: A

QUESTION 10

Which two actions can cause locking on system tables in tempdb?(Choose 2)

- A. create table in tempdb
- B. select into a #tempdb
- C. worktables created for reformatting
- D. selecting from a large table in tempdb

Answer: A , B

QUESTION 11

Review the sp_sysmon section below.

Procedure Cache Management per sec per xact count % of total

Procedure Request 242.6 5.4 150267 n/a

Procedure Reads from Disk 1.8 0.0 1104 0.7%

Procedure Writes to Disk 0.0 0.0 0 0.0%

Procedure Removals 3.5 0.1 2180 n/a

Which statement is correct about this server's stored procedure activity?

Answer: D

QUESTION 18

Give the following cursor code fragment,

```
declare title_cursor cursor
```

```
for select price from title
```

```
for update
```

```
go
```

```
declare @price money
```

```
open title_cursor
```

```
fetch title_cursor into @price
```

What locks are held on the row or data page with each fetch?

- A. Update locks
- B. Exclusive locks
- C. Shared locks
- D. No locks are held

Answer: A

QUESTION 19

Increasing the number of pre-allocated extents to 16 is likely to help the performance of which of the following operations?

- A. index creation
- B. bcp in
- C. select into
- D. alter table lock datarows
- E. dbcc checkstorage

Answer: B

QUESTION 20

If a DBA detects many page splits after data has been sequentially insert by an application, what is the best option to minimize the number of future page splits?

- A. deactivate the ascinserts option on the table
- B. set the max_row_per_page to a low value
- C. set the fillfactor to a low value at server level
- D. activate the ascinserts option on the table
- E. set the fillfactor to a low value at table level

Answer: D

QUESTION 21

Which of the following operations require "quiesce database"?

- A. altering the database to a new size
- B. performing non-Sybase unmirroring and backup operations
- C. truncating the transaction log
- D. dumping the database using Sybase Backup Server

Answer: B

QUESTION 22

All of the following tactics may reduce contention for cache spinlocks in an ASE configured to support multiple online engines,except

- A. Increasing the spinlock ratio
- B. Creating additional named caches and binding frequently used tables to them
- C. Creating partitions in heavily used caches
- D. Returning to a single engine configuration

Answer: A

QUESTION 23

Which conditions must be true to create a nonclustered index in parallel?(choose 3)

- A. the server must be configured for parallel access
- B. the table must be partitioned
- C. the database option select into/bulkcopy/pllsort must set to true
- D. the data must reside on a user-defined segment
- E. the number of pages in the table must be at least eight times the configured of sort buffers

Answer: A, C, E.

QUESTION 24

An application has been implemented in the production environment, and there has been a slowdown. To investigate that problem, you are looking at the procedure cache management section of the sp_sysmon output. The sp_sysmon output has reported a high 'percentage reads from disk'. What will be your action based on the above fact?

- A. No action is necessary based on this fact
- B. The application is dynamically generating stored procedures and the servers tuning will not help
- C. This indicated that the procedure cache is too small and you will increase the procedure cache size
- D. You will bind syscomments and sysobjects table to its own cache which will reduce the high percentage of 'procedures reads from disk'.

Answer: C

QUESTION 25

Given the following statement:

```
create clustered index idx1 on titles(title_id) on new_segment
```

Where:

The titles table has 3 partitions

The target segment new_segment spans 10 devices

The max parallel degree parameter is 25

The session-level parallel_degree parameter is set to 2

How will the index be created?

- A. With a parallel sort with 2 procedures and 1 consumer
- B. The statement will fail
- C. With a serial sort
- D. With a parallel sort with 2 procedures and 2 consumers

Answer: B

QUESTION 26

Which of the following cache or IO tuning methods can reduce the spinlock contention on a single "hot" table that is already bound to a named cache?

- A. Create 16K buffer pool in the named cache
- B. Change the cache replacement strategy of the named cache to relax LRU
- C. Partition the named cache
- D. Increase the named cache size

Answer: C

QUESTION 27

What is the maximum row size for DOL tables?

- A. 1024 bytes of data
- B. 1958 bytes of data
- C. 1962 bytes of data
- D. 2000 bytes of data

Answer: B

QUESTION 28

What is the degree of parallelism used in a nested-loop join?

- A. the product of the worker processes that access individual tables in the join
- B. the sum of the worker processes that access individual tables in the join

- C. a nested-loop join is always performed using a serial scan
- D. two(1 coordinating process and 1 worker process)

Answer: A

QUESTION 29

Given: Adaptive Server physical storage space is divided into a hierarchy. What is the hierarchy sequence, from largest to smallest units of storage?

- A. allocation unit, extent, page, database fragment, device
- B. page, extent, database fragment, device, allocation unit
- C. device, database fragment, allocation unit, extent, page
- D. extent, allocation unit, page, device, database fragment

Answer: C

QUESTION 30

Certkiller .com needs to connect 5000 concurrent sessions to an Adaptive Server Enterprise server. The errorlog at startup of the ASE server shows the following:
00:00000:00000:1998/04/01 08:37:00.01 Kernel Network and device connection limit is 1014.

The DBA uses sp_configure to set "number of user connections" to 5000. What is the minimum number of engines, configured with sp_configure "max online engines", that must be used for this server to allow the 5000 connections to simultaneously occur?

- A. 1 engine
- B. 5 engines
- C. 6 engines
- D. 4 engines
- E. 10 engines

Answer: B

QUESTION 31

Which of the following can NOT have an assigned execution class

- A. database
- B. A stored procedure
- C. An application
- D. A login

Answer: A

QUESTION 32

When a row is deleted on a table with datarows locking:

- A. Rows on the page are moved up so the empty space is at the end of the page
- B. The space occupied by the deleted row is not reclaimed immediately
- C. The space Occupied by the deleted row is filled with zeros
- D. The last row on the page is moved into the space created by the deleted row

Answer: B

QUESTION 33

Which of the following is true about the configuration option "max scan parallel degree"?

- A. By default hash based scans will not occur
- B. It must be greater than or equal to "max parallel degree"
- C. It is a static parameter
- D. It sets the maximum number of worker processes for any query
- E. "set scan parallel degree n" may be used to temporarily use more processes than "max scan parallel degree"

Answer: A

QUESTION 34

Given that all other user tasks are running with the default execution class, assigning execution class EC1 to a long-running I/O intensive application will have which of the following effects:

- A. improve the responsiveness of the long-running task, possibly shortening its elapsed time
- B. cause the long running task to consume all available CPU time, starving other tasks
- C. increase the rate at which deadlocks occur
- D. substantially decrease the frequency of scheduling the long running task, slowing it down
- E. permit other tasks to saturate the available CPU resources, causing a bottleneck

Answer: A

QUESTION 35

Scheduling priority is a direct attribute of which of the following:

- A. Engine Group
- B. Execution class
- C. Application
- D. Stored procedure
- E. Database

Answer: B

QUESTION 36

Given the following sp_sysmon output,
Procedure Cache Management per sec per xact count% of total
Procedure Requests 242.6 5.4 14556 n/a
Procedure Reads from Disk 167.6 0.0 10000 68.7 %
Procedure Writes to Disk 0.1 0.0 40 0.1 %
Procedure Removals 0.0 0.0 0 n/a

If the above indicates that a performance problem exists, indicate what steps may be taken to alleviate the problem?

- A. The number of times a procedure was read from disk is high. Therefore, increase procedure cache size
- B. The number of times a procedure was written to disk is low. Therefore, decrease procedure cache size
- C. The number of times a procedure was read from disk is high. Therefore, decrease procedure cache size
- D. The above output does not show a performance problem

Answer: D

QUESTION 37

Where are table/index level statistics stored?

- A. sysstatistics
- B. systabstats
- C. On the leaf level of the index
- D. The GAM page
- E. The distribution page

Answer: B

QUESTION 38

Which one of the following queries would not benefit from parallel storing?

- A. Queries that use distinct
- B. Queries that perform merge-joins requiring sorts
- C. Queries that use the reformatting strategy
- D. Queries that use union all

Answer: D

QUESTION 39

Which of the following stored procedures is used to display the current sessions scheduling and engine affinity properties

- A. sp_showcontrolinfo

- B. sp_displaylogin
- C. sp_who
- D. sp_showpsex
- E. sp_showexeclass

Answer: D

QUESTION 40

Which of the following applications can benefit from data-only locking?

- A. The table is a partitioned heap table that has a high rate of inserts
- B. Applications need to maintain an extremely high transaction rate even though contention is low
- C. Applications require clustered access to the data rows due to range queries or order by clauses
- D. A nonclustered index is required and all the inserts occur at the end of the index

Answer: D

QUESTION 41

Given: number of worker processes= 10, max parallel degree= 5, max scan parallel degree=4. A table has 3 partitions. The first partition has 70 pages, the second partition has 70 pages the third partition has 10 pages. Which of the following statements are true about the table?

(choose 2)

- A. the partition skew is less than 2.0
- B. the partition skew is greater than 2.0
- C. queries against the table cannot be optimized for partition-based access because max parallel degree is greater than the number of partitions
- D. queries against the table can be optimized for partition-based parallel access because the average number of pages per partition is greater than 20
- E. queries against the table cannot be optimized for partition-based parallel access because each partition does not have at least 20 pages

Answer: A,E

QUESTION 42

There are no named caches in the server. Only a 4K pool has been added to the default data cache. The following query is executed:

```
select count(*) from authors
```

```
Where state= 'MA'
```

```
/* 'authors' table does not have any index on state and there are around 10000 rows in the table */
```

Which statement is correct?

- A. The optimizer cannot consider a large I/O as the table is not bound to a named cache
- B. It is a point query and hence large I/O will not be performed
- C. As there is no 16K pool to perform a table scan, the optimizer will choose 2K pool.
- D. There will be a table scan and the optimizer will select 4K pool to perform I/O

Answer: D

QUESTION 43

Which ASE commands or utilities can be used to examine partition skew? (choose 2)

- A. sp_helpdb
- B. sp_helppartition
- C. sp_help
- D. optdiag
- E. reorg

Answer: B, C.

QUESTION 44

When an engine looks for a tasks to run, what is the order in which it searches the priority run queues?

- A. It is user defined
- B. high, medium and slow
- C. low, medium and high
- D. longest queue to shortest queue

Answer: B

QUESTION 45

Which statement is not true about cache partitioning?

- A. Default data cache can be partitioned
- B. A named cache with a 16K buffer pool can be partitioned
- C. Cache partitioning requires a server reboot
- D. A 2MB named cache can be divided into 8 cache partitions
- E. Named caches using the relaxed LRU cache replacement strategy can be partitioned

Answer: D

QUESTION 46

If the command "lock table titles in share mode wait2" is issued and the lock can not be acquired within the time period.

- A. the batch is aborted
- B. the transaction is rolled back

- C. an information message is generated
- D. the command waits indefinitely

Answer: C

QUESTION 47

What effect does the use of the command "set sort_merge on" have on the server?

- A. It changes the server level setting to allow the use of merge joins
- B. It overrides the server level setting to allow the use of merge joins in the current session
- C. It overrides the server level setting to allow the use of merge joins by all currently connected users
- D. It allows the use of merge joins in the current session only if the server level is set to "enabled".

Answer: B

QUESTION 48

Transaction Isolation Level 2 is directly supported in the following lock scheme(s):

- A. Allpages
- B. Allpages, Datapages
- C. Allpages, Datapages, Datarows
- D. Datapages, Datarows

Answer: D

QUESTION 49

Which of the following are true of the Datarows Locking Scheme? (choose 2)

- A. Index pages are locked
- B. No transaction locks are held on index pages
- C. Server uses page locks and row locks, but no table locks
- D. Server uses table locks and row locks, but not page locks

Answer: B, D.

QUESTION 50

The following query is executed frequently against the database

```
select ta.title_id,a.au_lname  
from titleauthor ta,authors a  
where ta.au_id=a.au_id
```

Which database denormalization technique could be used to improve the performance of this query?

510-015

- A. Add the redundant column au_id to the titleauthors table
- B. Add the redundant column au_lname to the titleauthors table
- C. Move the title id column to the authors table
- D. Move the au_lname column to the titleauthors table
- E. Change the associative table titleauthor to use the natural key for both the title and the author

Answer: B

QUESTION 51

A query and any views in a query can reference how many user tables?

- A. 12
- B. 16
- C. 32
- D. 50
- E. 192

Answer: D

QUESTION 52

A decision support (DSS) application is communicating with the ASE server. The server configuration has been tuned to support the needs of the DSS application and the following configuration parameters have been set:

default network packet size = 8192

maximum network packet size = 8192

additional network memory = 75776

The DSS application is connecting to the server using the value of "default network packet size" parameter as its requested packet size

A new online transaction processing (OLTP) application is added to the server. A typical insert

or update statement from the OLTP application is less than 200 bytes.

Without impacting the performance of the existing DSS application, which two changes should be made to improve the network performance of the OLTP application? (choose 2)

- A. Decrease the size of the "default network packet size" parameter to 512
- B. Decrease the size of the "max network packet size" parameter to 4096
- C. Increase the value of the "additional network memory" parameter to 151552 (or 2*75776)
- D. Modify the DSS client to connect to the server using the value of the "max network packet size" parameter as its requested packet size
- E. Modify the DSS client to connect to the server using the value of the "additional network memory" parameter as its requested packet size

Answer: A, D.

QUESTION 53

Consider the following scenario: An update has to be performed to the price column in the titles table. Two are considered, a SQL update statement versus a cursor-based approach updating the price row-by-row. What might be an advantage of cursors in this scenario?

- A. Faster than set-oriented processing
- B. May increase concurrence by committing each row as a transaction
- C. Reduces the number of log records
- D. Enables users to access the table using the cursor's security privileges

Answer: B

QUESTION 54

Which of the following steps is NOT always necessary for performing a tuning benchmark?

- A. record a baseline measurement
- B. change one variable, holding all others constant
- C. make a new measurement
- D. restart Adaptive Server

Answer: D

QUESTION 55

systabstats does not contain information on:

- A. Number of data pages for a table
- B. Average length of data rows and leaf rows
- C. Stopping points for the reorg command
- D. Column histograms

Answer: D

QUESTION 56

When does the optimizer consider a hash based table scan on an APL table?

- A. if it is a heap table
- B. if it has a clustered index
- C. if it is a table with a clustered index and a non clustered index
- D. if it is a partitioned table with a clustered index

Answer: A

QUESTION 57

What type of access methods is used for a covered query?

- A. parallel index scan
- B. serial index scan
- C. hash based index scan
- D. hash based table scan

Answer: B

QUESTION 58

When you run optdiag command on a allpages-locked table with 1000 pages, what will be value of "Forwarded row count"?

- A. 1000
- B. 500
- C. 0
- D. 1

Answer: C

QUESTION 59

Examine this sp_sysmon output. Device Activity Details:

Device Activity Details:

/dev/rdisk/tl3d0s6				
data_dev1	per sec	per xact	count	% of total
-----	-----	-----	-----	-----
Reads				
	APF	0.7	1395	75.7 %
	non-APF	0.1	286	15.5 %
Writes		0.1	162	8.8 %
-----	-----	-----	-----	-----
Total I/O s		0.9	1843	6.1 %
Device Semaphore Granted		0.9	1839	60.0 %
Device Semaphore Waited		0.4	643	40.0 %

/dev/rdisk/tl3d0s7				
data_dev2	per sec	per xact	count	% of total
-----	-----	-----	-----	-----

Which action might be taken to eliminate this bottleneck?

510-015

- A. Change buffer replacement strategy from MRU to LRU
- B. Create multiple segments on data_dev1 and place tables on separate segments
- C. Remove clustered index on tables, which reside entirely on data_dev1
- D. Create segments spanning multiple devices and place objects on those segments

Answer: D

QUESTION 60

What is the maximum number of key columns allowed in an index?

- A. 16
- B. 31
- C. 32
- D. 63

Answer: B

QUESTION 61

The sp_sysmon output 'Large I/Os denied' reports the number of times large I/O could not be performed. The ASE cannot perform large I/O because of which three reasons? (choose 3)

- A. The disk I/O controller does not allow large I/O
- B. The page in a buffer already resides in another pool
- C. There are no buffers available in the requested pool to perform large I/O
- D. The table is not placed in its own segment and hence, large I/O can not be performed
- E. The first extent of an allocation unit is always read into a 2K pool. Hence, for a table scan, at least one large I/O will be denied

Answer: B, C, , E.

QUESTION 62

What effect does the use of the command "set sort_merge on" have on the server?

- A. It changes the server level setting to allow the use of merge joins
- B. It overrides the server level setting to allow the use of merge joins in the current session
- C. It overrides the server level setting to allow the use of merge joins by all currently connected users
- D. It allows the use of merge joins in the current session only if the server level is set to "enabled"

Answer: B

QUESTION 63

Consider the following scenario: An update has to be performed to the price column in

510-015

the titles tables. two are considered,a SQL update statement versus a cursor-based approach updating the price row-by-row. What might be an advantage of cursors in this scenario?

- A. Faster than set-oriented processing
- B. May increase concurrency by committing each row as a transaction
- C. Reduces the number of log records
- D. Enables users to access the table using the cursor's security privileges

Answer: B

QUESTION 64

The following query is executed frequently against the database

```
select a.au_fname,a.au_lname,  
ad.street, ad.city,  
ad.state, ad.zip  
from authors a, author_address ad  
where a.au_id=ad.au_id
```

What database denormalization technique could be used to improve the performance of this query?

- A. Move the au_lname and au_fname columns to the author_address table
- B. Collapse the author_address and authors table into a single table
- C. Move the street, city, state, and zip columns to the authors table
- D. Create a view called mailing_address to eliminate the two table join from the query
- E. Add the redundant column au_id to the author_address table

Answer: B

QUESTION 65

Examine this sp_sysmon output. (choose 2)

Lock Summary per sec per xact count % of total

.....

Last page locks on heaps

Granted 120.1 1.0 7258 80.0%

Waited 24.0 0.2 150 20.0%

Which two actions might be taken to eliminate this bottleneck? (choose 2)

- A. Increase ASE memory
- B. Increase data cache size
- C. Create partitions for tables, which are candidates for heavy inserts
- D. Add clustered index to distribute inserts randomly across the data pages
- E. Bind tables to named caches

Answer: C, D.

QUESTION 66

Worker Process Management

```
-----  
per sec per xact count % of total  
.....  
Worker Process Requests  
Requests Granted 0.1 8.0 16 50.0%  
Requests Denied 0.1 8.0 16 50.0%  
-----  
Total Requests 0.2 16.0 32  
Worker Process Usage  
Total Used 0.4 39.0 78 n/a  
Max Ever Used During Sample 0.1 12.0 24 n/a  
Memory Requests for Worker Processes  
Succeeded 4.5 401 902 100.0%  
Failed 0 0 0 0.0%  
-----  
Total Requests 4.5 401.0 902  
Avg Mem Ever Used by a WP (in bytes) n/a n/a 311.7 n/a  
What should you do if you saw the following sp_sysmon output
```

- A. Increase "max parallel degree.
- B. Increase "max scan parallel degree".
- C. Increase "number of worker processes".
- D. Increase "memory per worker processes"

Answer: C

QUESTION 67

Which of the following cache or IO tuning methods can reduce the spinlock contention on a single "hot" table that is already bound to a named cache?

- A. Create 16K buffer pool in the named cache
- B. Change the cache replacement strategy of the named cache to relaxed LRU
- C. Partition the named Cache
- D. Increase the named cache size

Answer: C

QUESTION 68

You are tuning an ASE which is at "steady state". If you suspect, based on current usage, that too much memory is allocated to open object structures in your server, which of the following procedures should you use to most directly confirm or refute your hypothesis?

- A. sp_helpconfig
- B. sp_countmetadata

- C. sp_objects_stats
- D. sp_monitorconfig
- E. sp_configure

Answer: D

QUESTION 69

What are the three attributes of an execution class?

- A. login name, host, and client
- B. client, stored procedure, and time slice
- C. base priority, time slice and engine group
- D. session id, worker process, and engine group

Answer: C

QUESTION 70

How many predefined execution classes are there in ASE?

- A. 4
- B. 3
- C. 2
- D. 1

Answer: B

QUESTION 71

Which of the following statements best describes a DOL table with a cluster ratio of 0.97:

- A. The table is poorly structured and should be fixed by a reorg rebuild
- B. Dense table packing will increase the probability of data page splits
- C. Large I/O strategies will be particularly effective on this table
- D. The clustered index is almost as big as the table itself
- E. Row formatting is extremely unlikely due to large allocation gaps

Answer: C

QUESTION 72

Which of the following applications can benefit from data-only locking?

- A. The table is a partitioned heap table that has a high rate of inserts
- B. Applications need to maintain an extremely high transaction rate even though contention is low
- C. Applications require clustered access to the data rows due to range queries or order by clauses
- D. A nonclustered index is required and all the inserts occur at the end of the index

Answer: D

QUESTION 73

Which stored procedure can the DBA use to change a login's priority on-the-fly?

- A. sp_setpsex
- B. sp_bindexclass
- C. sp_unbindexclass
- D. sp_addengine
- E. sp_modifylogin

Answer: A

QUESTION 74

All of the following tactics may reduce contention for cache spinlocks in an ASE configured to support multiple online engines, except

- A. Increasing the spinlock ratio
- B. Creating additional named caches and binding frequently used tables to them
- C. Creating partitions in heavily used caches
- D. Creating a second buffer pool in a heavily used named cache
- E. Returning to a single engine configuration

Answer: A

QUESTION 75

An application uses stored procedures to perform updates, inserts and deletes for an OLTP application. Some update stored procedures optimize with a join order of table_A->table_B. Other stored procedures optimize with a join order of table_B->table_

A. Which solutions needs to be added within the stored procedures to guarantee that deadlocks are reduced or eliminated by always making the join order table_A->table_B?

- A. Non-correlated subqueries to force outside-in processing
- B. SET FORCEPLAN ON
- C. (index table order 1)- an index hint applied to table_A
- D. SET SHOWPLAN ON
- E>Define join-ordered Views on the affected tables

Answer: B

QUESTION 76

Increasing "runnable process search count" will have the effect of:

510-015

- A. enabling more online engines to work together in a cluster
- B. increasing the amount of CPU time ASE consumes when all users are sleeping
- C. making ASE less responsive when an I/O intensive workload is present
- D. increasing the number of user connections that can be scheduled at once
- E. increasing the aggressiveness of APF

Answer: B

QUESTION 77

Which of the following is NOT true of the DOL Data Page Format?

- A. The page header for DOL pages larger than for APL pages
- B. There are no data page chains (nextpg/prevpg pointers are not maintained)
- C. Free space is dynamically maintained as one contiguous chunk
- D. The same page format is used for both Datarows and Datapages locked tables

Answer: C

QUESTION 78

Given the following statement

```
create index idx1 on titles(title_id) on new_segment
```

Where:

The titles table has 3 partitions

The target segment new_segment spans 10 devices

The max parallel degree parameter is 25

If available worker processes at runtime is 5, what is the total number of procedures used for the parallel sort?

- A. 10
- B. 3
- C. 25
- D. 1

Answer: D

QUESTION 79

An 8-engine Adaptive Server is undergoing a high rate of Logical I/O and spinlock contention in the default cache. The DBA has used Sybase ASE Monitor and has found that two 4MB tables are table-scanning within the application. This is causing 50% of the server's Logical I/Os. The code within the application cannot be changed.

The task for the DBA is to increase throughput of the Adaptive Server and lower the amount of spinlock contention without changing indexes, tables or the application.

Which two are valid operations which should be done to meet this objective? (choose 2)

- A. Create two named caches of 2 MB each and then recycle the Adaptive Server to activate them.

510-015

- B. Create two named caches of 5 MB each and then recycle the Adaptive Server tables and activate them.
- C. Create two 5MB 16KB Page Pools within the default data cache. Bind the two small tables to these page pools.
- D. Bind the two tables to their own newly-created and active named cache
- E. Compress the tables which join to the two small tables so that they are chosen as outer tables. This will eliminate the need to table scan the small tables.

Answer: B, D

QUESTION 80

What is an abstract query plan?

- A. It is a query plan that gets generated at run time by the optimizer
- B. It describes the execution plan for a query using a language created for that purpose
- C. It is a system table
- D. It is a stored procedure

Answer: B

QUESTION 81

With sp_sysmon one can determine (choose 3)

- A. Tables that have had the most logical i/os
- B. data caches that have been used the most
- C. buffer pools that have been used the most
- D. user log caches that have been flushed the most
- E. stored procedures that have been executed the most
- F. engines that have been used the most

Answer: B, C, , F.

QUESTION 82

Which two are advantages of use max_rows_per_page to reduce the row density?(choose 2)

- A. It may decrease the disk space used
- B. It may decrease the contention
- C. It may decrease the number of page splits
- D. It may decrease the number of page I/Os
- E. It may help to decrease row movement of frequently updated varchar columns

Answer: B, E.

QUESTION 83

The reorg command may be used to: (choose 2)

- A. Compact a table that uses datarows locking
- B. Reclaim unused space in a table that uses allpages locking
- C. Reclaim unused space in an index that uses datapages locking
- D. Rebuild a table that uses allpages locking

Answer: A, C.

QUESTION 84

Which of the following are valid batches?(choose 2)

- A. create table tab 1 (a int)
execute ("select*from tab 1")
- B. execute ("create table tab 1(a int)")
select* from tab 1
- C. execute ("create table tab 1 (a int)")
execute ("select* from tab 1")
go

Answer: A, C.

QUESTION 85

If an index with a matching SARG is available on a small DOL table the optimizer will do which of the following?

- A. Uses a table scan
- B. Uses the index if the number of pages in the table is less than or equal to the "concurrency_opt_threshold" configuration.
- C. Uses the index if the number of pages in the table is more than the "concurrency_opt_threshold" configuration.
- D. Uses the index if the table has " concurrency_opt_threshold" set to 0

Answers : B

QUESTION 86

Parallel sorting can be used for: (choose 3)

- A. Sort Merge joins
- B. slow bcp
- C. Reformatting
- D. Creating indexes
- E. reorg compact

Answer: A, C, , D.

QUESTION 87

Which of the following are true of the Datarows Locking Scheme?(choose 2)

- A. Index pages are locked
- B. No transaction locks are held on index pages
- C. Server uses pages locks and row locks, but not table locks
- D. Server uses table locks and row locks, but not page locks

Answer: B, D.

QUESTION 88

How do you reclaim space from a fragmented allpages heap table?

- A. run reorg command
- B. dump and load the database
- C. create and drop a clustered index
- D. truncate the least used pages

Answer: C

QUESTION 89

Engine affinity a task to run on:

- A. only one engine
- B. an engine when it is made online using dbcc online (engine)
- C. an engine or a group of engines
- D. none of the above

Answer: C

QUESTION 90

Given:

A table named books:

title char(100),

title_id char(15),

type char(15),

price money not null

The books table has an index created as follows:

```
create index price index 1
```

```
on books(price desc)
```

Which of the following queries would be covered and not require access to the data pages? (choose 5)

- A. select count(*) from books
- B. select sum(price) from books
- C. select type, sum(price) from books group by type
- D. select price from books
- E. select price from books order by price ASC

F.select price from books order by price DESC

Answer: A,B,D,E,, F.

QUESTION 91

Increasing "housekeeper free write percent" is likely to improve which of the following aspects of server performance (choose 3)

- A. increase the frequency of "free checkpoints"
- B. increase the responsiveness of deadlock checking
- C. frees space consumed by logically deleted rows
- D. decrease the number of buffers "found in wash"
- E. compensate for I/O overloads induced by overly aggressive APF

Answer A,C,D

QUESTION 92

How does the optimizer choose the optimal degree of parallelism for partition based scans?

- A. It will use the maximum number of available worker processes
- B. It will use the maximum allowed worker processes based on "max parallel degree"
- C. It will use the maximum allowed worker processes based on "max scan parallel degree"
- D. It will use at most one worker process for every partition in the table
- E. It will use no more than 3 worker process.

Answer: D

QUESTION 93

Which of the following statements describes readpast locking? (choose 2)

- A. It can be specified at the session, transaction, and table level
- B. It allows insert, update, and delete commands to read past any incompatible lock
- C. It allows reader to not block writers
- D. It allows select and readtext queries to silently skip all rows or pages locked with incompatible locks

Answer: B,D

QUESTION 94

Which of the following statements are true of table statistics: (choose 2)

- A. Creating statistics on unindexed columns can improve the performance of some queries
- B. Truncating a table deletes the column-level statistics in sysstatistics
- C. Dropping an index will drop statistics on its column(s)

D. Histograms are kept on a per-column basis

Answer: A, D

QUESTION 95

Where are table/index level statistics stored?

- A. sysstatistics
- B. systabstats
- C. On the leaf level of the index
- D. The GAM page
- E. The distribution page

Answer: B

QUESTION 96

Assuming that all other user tasks are running with the default execution class, assigning execution class EC3 to a long-running CPU intensive application will have which of the following effects:

- A. improve the responsiveness of the long-running task, possibly shortening its elapsed time
- B. cause the long running task to consume all available CPU time, starving other tasks
- C. increase the rate at which deadlocks occur
- D. substantially decrease the frequency of scheduling the long running task, slowing it down.
- E. increase the amount of unspent idle time available for future demand

Answer: D

QUESTION 97

What are some ways to reduce last data page lock contention for inserts? (choose 3)

- A. alter the table to use the Datarows locking scheme
- B. alter the table to use the Datapages locking scheme
- C. partition the table
- D. create a clustered index on a random key
- E. increase the server configuration for the number of locks
- F. increase the server configuration for the deadlock checking period

Answer: A, C, , D.

QUESTION 98

When configuring the "additional network memory" parameter, the number of buffers required per connection must be taken into consideration. How many buffers does each connection require?

- A. 1
- B. 2
- C. 3
- D. 4
- E. 5

Answer: C

QUESTION 99

What type of access method can be used for an unpartitioned APL table with a useful clustered index?

- A. hash based table scan
- B. serial table scan
- C. serial index scan
- D. clustered index partition scan

Answer: C

QUESTION 100

Under what conditions can a clustered index be created using the with sorted_data option? (choose 3)

- A. when dropping and recreating the clustered index on a table that uses the Allpages locking scheme
- B. when dropping and recreating the clustered index on a table with a clustered index cluster ratio below 1
- C. after using bcp to load data into an unpartitioned empty Allpages table that has no indexes, if the data is already sorted in order of the key
- D. after using bcp to load data into an unpartitioned empty Data-only locking scheme table that has no indexes, if the data is already sorted in order of the key
- E. after creating a nonclustered index on the same key that the clustered index will be created on

Answer: A, C, , D.

QUESTION 101

Which are the two ways for boosting the performance of parallel sort during index creation? (choose 2)

- A. "trunc log on chkpt" database option set to true in the target database
- B. configuring a pool of 16K buffers in the cache used by the sort
- C. create index...with consumers=N
- D. set scan_parallel_degree N
(where N=number of target partitions)

E. Keep tempdb on one contiguous device

Answer: B,C.

QUESTION 102

What type of access method is used when you delete from a partitioned table with a useful clustered index?

- A. hash based table scan
- B. partitioned index scan
- C. serial index scan
- D. serial table scan

Answer: C

QUESTION 103

What is the most important criterion when ASE chooses a parallel query plan?

- A. total amount of work done
- B. number of locks
- C. response time
- D. number of logical I/Os performed

Answer: C

QUESTION 104

Which of the following operations require "quiesce database"?

- A. altering the database to a new size
- B. performing non-Sybase unmirroring and backup operations
- C. truncating the transaction log
- D. dumping the database using Sybase Backup Server

Answer: B

QUESTION 105

Which statement is not true about cache partitioning?

- A. Default data cache can be partitioned
- B. A named cache with a 16K buffer pool can be partitioned
- C. Cache partitioning requires a server reboot
- D. A 2MB named cache can be divided into 8 cache partitions
- E. Named caches using the relaxed LRU cache replacement strategy can be partitioned

Answer: D

QUESTION 106

What is the maximum number of tables allowed in a join query?

- A. 50 User tables and 14 work tables
- B. 32 User tables and 14 work tables
- C. 30 User tables and 12 work tables
- D. 16 User tables and 12 work tables

Answer: A

QUESTION 107

Where can the last page of a partition be found for an APL table?

- A. sysindexes
- B. root page
- C. OAM page
- D. allocation page
- E. control page

Answer: E

QUESTION 108

When a row is deleted on a table with datarows locking:

- A. Rows on the page are moved up so the empty space is at the end of the page.
- B. The space occupied by the deleted row is not reclaimed immediately
- C. The space occupied by the deleted row is filled with zeros
- D. The last row on the page is moved into the space created by the deleted row

Answer: B

QUESTION 109

Choose the transaction isolation level and table lock scheme that is described by the following sp_lock output.

```
fid spid loid locktype table_id page row dbname class context
```

```
-----
```

```
0 11 22 Ex_intent 16003088 0 0 pubs2 Non Cursor Lock
```

```
0 11 22 Ex_row 16003088 810 3 pubs2 Non Cursor Lock Range 0 11 22 sh_row
```

```
16003088 824 0 pubs2 Non
```

```
Cursor Lock Inf Key
```

- A. transaction isolation level 3 and a datarows table
- B. transaction isolation level 3 and a datapages table
- C. transaction isolation level 2 and a datarows table
- D. transaction isolation level 3 and an allpages tables
- E. transaction isolation level 2 and an allpages table

Answer: A

QUESTION 110

Given FREDs login is assigned EC2. When he uses a stored procedure bound to EC3 his scheduling priority is:

- A. high
- B. medium
- C. low
- D. unable to be determined

Answer: B

QUESTION 111

Spinlock contention is high in a named cache. A small database is bound to the named cache. This database fits into the named cache entirely. Cache misses do not occur within this named cache. This named cache is 20 MB in size and uses the strict-policy of buffer replacement. Which two will lower spinlock contention for the named cache?(choose 2)

- A. Set the buffer replacement policy for the specific named cache to "relaxed"
- B. Increase the number of online engines for this Adaptive Server Enterprise. Larger numbers of engines can spread spinlock contention out across more engines.
- C. Use the Adaptive Server Enterprise parallel features to process the queries in parallel for this database and lower spinlock contention on the serial process.
- D. Create a second named cache and bind some objects from the database to a second cache

Answer: A, D.

QUESTION 112

The best way to keep a clustered APL table compact when inserting data that is naturally ordered in the same sequence as the index key is to:

- A. enable MRU buffer replacement on the table
- B. reduce the identity_gap of the table
- C. force use of a "safe" nonclustered index
- D. increase the reservedpagegap for the table
- E. enable asinserts on the table

Answer: E

QUESTION 113

The following query and showplan output is an example of what type of merge join operation?

```
select t1.title_id, t2.title
```

from titles t1, titles t2

where t1.title_id = t2.title_id

QUERY PLAN FOR STATEMENT 1 (at line 1).

STEP 1

The type of query is INSERT.

The update mode is direct.

FROM TABLE

t2

Nested iteration.

Table Scan.

Forward scan.

Positioning at start of table.

Using I/O Size 2 Kbytes for data pages.

With LRU Buffer Replacement Strategy for data pages.

TO TABLE

Worktable1.

Worktable created for sort merge join

STEP 2

The type of query is SELECT.

b

FROM TABLE

t1

Merge join (outer table).

Serial data merge.

Index : 1

Forward scan.

Positioning by key.

Index contains all needed columns, Base table will not be read.

Keys are:

Using I/O Size 2 Kbytes for index leaf pages.

With LRU Buffer Replacement Strategy for index leaf pages.

FROM TABLE

Worktable.

Merge join (inner table).

Table Scan

Forward scan

Positioning at start of table

Using I/O Size 2 Kbytes for data pages.

With LRU Buffer Replacement Strategy for data pages.

- A. Full-Merge Join
- B. Left-Merge Join
- C. Right-Merge Join
- D. Sort-Merge Join

Answer: B

QUESTION 114

The "abstract plan language "is:

- A. A language used to tell the optimizer how to process a query
- B. The internal command language used by sp_sysmon
- C. A new type of command language for creating very complicated queries
- D. The language used internally by the server for planning the most efficient use of memory

Answer: A

QUESTION 115

Which of the following create table can prevent or minimize row forwarding in DOL tables?

- A. fillfactor
- B. max_rows_per_page
- C. exp_row_size
- D. reservepagegap

Answer: C

QUESTION 116

Given: Adaptive Server physical storage space is divided into a hierarchy. What is the hierarchy sequence, from largest to smallest units of storage?

- A. allocation unit, extent, page, database fragment, device
- B. page, extent, database fragment, device, allocation unit
- C. device, database fragment, allocation unit, extent, page
- D. extent, allocation unit, page, device, database fregment

Answer: C

QUESTION 117

Adaptive Server sets a demand lock to indicate that a transaction is next in queue to lock a table or page.Which statement best describes demand locking behaviour in serial or parallel execution?

- A. When queries are running in serial,demand locking treats all the shared locks from a family of worker processes as if they were a single task
- B. When queries are running in parallel, demand locking treats all the shared locks from a family of worker processes as if they were a single task
- C. When queries are running in parallel, demand locking treats all the shared locks from a family of worker processes as if they were a different task
- D. When queries are running in serial there is no need for demand locking

Answer: B

QUESTION 118

The titles table has 50,000 rows and a nonclustered index on the price column. Which of the following queries use less logical I/Os?

- A. `select "Min" =min(price), "Max" =max(price) from titles`
- B. `select "Min" = (select min(price)from titles), "Max" =(select max(price)from titles)`
- C. They will both use the same number of I/Os

Answer: B

QUESTION 119

Given a user of application ctisql is assigned EC2. When she uses a stored procedure bound to EC1, her scheduling priority is:

- A. high
- B. medium
- C. low
- D. unable to be determined

Answer: A

QUESTION 120

If `sp_sysmon` output consistently shows a high number of Voluntary Yields, which configuration parameter could be modified to improve performance?

- A. user log cache size
- B. time slice
- C. identity grab size
- D. stack size
- E. cpu accounting flush interval

Answer: B

OfficialCerts.com Certification Exam Full Version Features;

- Verified answers researched by industry experts.
- Exams **updated** on regular basis.
- Questions, Answers are downloadable in **PDF** format.
- **No authorization** code required to open exam.
- **Portable** anywhere.
- 100% success **Guarantee**.
- **Fast**, helpful support 24x7.

View list of All exams we offer;

<http://www.officialcerts.com/allexams.asp>

To contact our Support;

<http://www.officialcerts.com/support.asp>

View FAQs

<http://www.officialcerts.com/faq.asp>

Download All Exams Samples

<http://www.officialcerts.com/samples.asp>

To purchase Full Version and updated exam;

<http://www.officialcerts.com/allexams.asp>



Shop now using **PayPal**



3COM	CompTIA	Filemaker	IBM	LPI	OMG	Sun
ADOBE	ComputerAssociates	Fortinet	IISFA	McAfee	Oracle	Sybase
APC	CWNP	Foundry	Intel	McData	PMI	Symantec
Apple	DELL	Fujitsu	ISACA	Microsoft	Polycom	TeraData
BEA	ECCouncil	GuidanceSoftware	ISC2	Mile2	RedHat	TIA
BICSI	EMC	HDI	ISEB	NetworkAppliance	Sair	Tibco
CheckPoint	Enterasys	Hitachi	ISM	Network-General	SASInstitute	TruSecure
Cisco	ExamExpress	HP	Juniper	Nokia	SCP	Veritas
Citrix	Exin	Huawei	Legato	Nortel	See-Beyond	Vmware
CIW	ExtremeNetworks	Hyperion	Lotus	Novell	Google	

You have made the
Right Choice

You are becoming member of most comprehensive, accurate, highest quality and lowest cost certification resource in the world.

