

Sybase Job Interview Questions And Answers



Interview Questions Answers

<https://interviewquestionsanswers.org/>

About Interview Questions Answers

Interview Questions Answers . ORG is an interview preparation guide of thousands of Job Interview Questions And Answers, Job Interviews are always stressful even for job seekers who have gone on countless interviews. The best way to reduce the stress is to be prepared for your job interview. Take the time to review the standard interview questions you will most likely be asked. These interview questions and answers on Sybase will help you strengthen your technical skills, prepare for the interviews and quickly revise the concepts.

If you find any **question or answer** is incorrect or incomplete then you can **submit your question or answer** directly with out any registration or login at our website. You just need to visit [Sybase Interview Questions And Answers](#) to add your answer click on the *Submit Your Answer* links on the website; with each question to post your answer, if you want to ask any question then you will have a link *Submit Your Question*; that's will add your question in Sybase category. To ensure quality, each submission is checked by our team, before it becomes live. This [Sybase Interview preparation PDF](#) was generated at **Wednesday 29th November, 2023**

You can follow us on FaceBook for latest Jobs, Updates and other interviews material.
www.facebook.com/InterviewQuestionsAnswers.Org

Follow us on Twitter for latest Jobs and interview preparation guides.
<https://twitter.com/InterviewQA>

If you need any further assistance or have queries regarding this document or its material or any of other inquiry, please do not hesitate to contact us.

Best Of Luck.

Interview Questions Answers.ORG Team
<https://InterviewQuestionsAnswers.ORG/Support@InterviewQuestionsAnswers.ORG>



Sybase Interview Questions And Answers Guide.

Question - 1:

How do I tell which tables have identities in Sybase?

Ans:

You can tell if a table has identities one of two ways:

1. sp_help [tablename]: there is a field included in the sp_help output describing a table called "Identity." It is set to 1 for identity fields, 0 otherwise.
2. Within a database, execute this query:
 1. select object_name(id) "table", name "column", prec "precision"
 2. from syscolumns
 3. where convert(bit, (status & 0x80)) = 1
 4. gothis will list all the tables and the field within the table that serves as an identity, and the size of the identity field.

[View All Answers](#)

Question - 2:

How do I Configure the burn factor in Sybase?

Ans:

The number of identity values that gets "burned" upon a crash or a shutdown can be found by logging into the server and typing:

1. sp_configure "identity burning set factor"
 2. go
- the Default value set upon install is 5000. The number "5000" in this case is read as ".05% of all the potential identity values you can have in this particular case will be burned upon an unexpected shutdown." The actual number depends on the size of the identity field as you specified it when you created your table.

To set the burn factor, type:

1. sp_configure "identity burning set factor", [new value]
2. go

This is a static change; the server must be rebooted before it takes effect

[View All Answers](#)

Question - 3:

How do I configure Identities in Sybase?

Ans:

You can either create your table initially with the identity column:

1. create table ident_test
2. (text_field varchar(10),
3. ident_field numeric(5,0) identity)
4. go

Or alter an existing table and add an identity column:

1. alter table existing_table
2. add new_identity_field numeric(7,0) identity
3. go

When you alter a table and add an identity column, the System locks the table while systematically incrementing and adding unique values to each row. IF YOU DON'T SPECIFY a precision, Sybase defaults the size to 18! That's 1,000,000,000,000,000,000-1 possible values and some major problems if you ever crash your ASE and burn a default number of values... (10¹⁸ with the default burn factor will burn 5¹⁴ or 500,000,000,000,000 values...yikes).

[View All Answers](#)

Question - 4:

How can I execute dynamic SQL with ASE in Sybase?

Ans:

Adaptive Server Enterprise: System 12

ASE 12 supports dynamic SQL, allowing the following:

```
declare @sqlstring varchar(255)
select @sqlstring = "select count(*) from master..sysobjects"
exec (@sqlstring)
```



```
go
Adaptive Server Enterprise: 11.5 and 11.9
* Firstly define your local server to be a remote server using
sp_addserver LOCALSRV,sql_server[,INTERFACENAME]
go
* Enable CIS
sp_configure "enable cis",1
go
* Finally, use sp_remotesql, sending the sql to the server defined in point 1.
declare @sqlstring varchar(255)
select @sqlstring = "select count(*) from master..sysobjects"
sp_remotesql LOCALSRV,@sqlstring
go
```

[View All Answers](#)

Question - 5:

Which version of Open Client works with which ASE in Sybase?

Ans:

The TDS protocol that *is* Open Client is built so that either the client or server will fallback to a common dialect. I suppose that it is theoretically possible that both would fallback for some reason, but it seems unlikely. I was recently working with a client that was using Open/Client 4.2 to speak to a version 11.5 ASE using Powerbuilder 3 and 4! Amazing, it all worked! The main problem that you will encounter is not lack of communication but lack of features. The facility to bcp out of views was added to the 11.1.1 release. You will still be able to connect to servers with old copies of Open/Client, you just won't have all of the features.

There is also another fairly neat feature of the later releases of Open/Client, it has a very good compatibility mode for working with old applications. The client that was running Open/Client 4.2 with Powerbuilder 3 is now connecting to the database using version 11.1.1. Which is not bad when you remember that Powerbuilder 3 only talked 4.2 DBLib!

[View All Answers](#)

Question - 6:

Which should I use, RepAgent or LTM in Sybase?

Ans:

There are pros and cons to both, however, I think that it should be stated up front that RepAgents are the latest offering and I believe that Sybase would expect you to use that. Certainly the documentation for LTMs is a little buried implying that they do not consider it to be as current as LTMs.

LTM Cons:

- * Older technology. Not sure if it is being actively supported.
- * Not integrated within ASE, so there is a (small) performance penalty.
- * Separate processes, so need additional monitoring in production environments.

LTM Pros:

- * Possible to restart LTM without having to restart ASE.

RepAgent Cons

- * If it crashes it is possible that you will have to restart ASE in order to restart RepAgent.

RepAgent Pros

- * Latest, and presumably greatest, offering.
- * Tightly integrated with ASE so good performance.
- * Less to manage, no extra entries in the interfaces file.

[View All Answers](#)

Question - 7:

How can I improve throughput in Sybase?

Ans:

Check the Obvious

First, ensure that you are only replicating those parts of the system that need to be replicated. Some of this is obvious. Don't replicate any table that does not need to be replicated. Check that you are only replicating the columns you need. Replication is very sophisticated and will allow you to replicate both a subset of the columns as well as a subset of the rows.

Replicate Minimum Columns

Once the replication is set up and synchronised, it is only necessary to replicate those parts of the primary system that actually change. You are only replicating those rows and columns that need to be replicated, but you only need to replicate the actual changes. Check that each replication definition is defined using the clause:

```
create replication definition rep_def_name
with primary
```

[View All Answers](#)

Question - 8:

What is the difference between an LTM and a RepAgent in Sybase?

Ans:

Log Transfer Managers (LTMs) and RepAgents are the processes that transfer data between ASE and the Replication Server.

LTMs were delivered with the first releases of Replication Server. Each LTM is a separate process at the operating system level that runs along side ASE and Replication Server. As with ASE and Replication Server, a RUN_ and configuration file is required for each LTM. One LTM is required for each database being replicated.

Along with ASE 11.5 a new concept was introduced, that of RepAgent. I am not sure if you needed to use RepServer 11.5 as well, or whether the RepAgents could talk to earlier versions of Replication Server. Each RepAgent is, in effect, a slot-in replacement for an LTM. However, instead of running as separate operating system process, it runs as a thread within ASE. Pretty much all of the requirements for replication using an LTM apply to the RepAgents. One per database being replicated, etc. but now you do not need to have separate configuration files.

[View All Answers](#)

**Question - 9:**

What is the Difference Between Replication Server and SQL Remote in Sybase?

Ans:

Both SQL Remote and Replication Server perform replication. SQL Remote was originally part of the Adaptive Server Anywhere tool kit and is intended for intermittent replication. (The classic example is that of a salesman connecting on a daily basis to upload sales and download new prices and inventory.) Replication Server is intended for near real-time replication scenarios.

[View All Answers](#)

Question - 10:

What is Replication Server in Sybase?

Ans:

Replication Server moves transactions (insert, updates and deletes) at the table level from a source dataserer to one or more destination dataservers. The dataserer could be ASE or other major DBMS flavour (including DB2, Informix, Oracle). The source and destinations need not be of the same type.

What can it do ?

- * Move data from one source to another.
- * Move only a subset of data from source to destination. So, you can 'subscribe' to a subset of data, or a subset of the columns, in the source table, e.g. select * from clients where state = 'NY'
- * Manipulation/transformation of data when moving from source to destination. E.g. it can map data from a data-type in DB2 to an equivalent in Sybase.*
- * Provide a warm-standby system. Can be incorporated with Open Switch to provide a fairly seamless fail-over environment.
- * Merge data from several source databases into one destination database (could be for a warehouse type environment for example).
- * Move data through a complicated network down to branch offices, say, only sending the relevant data to each branch.

[View All Answers](#)

Question - 11:

What applications is ASA good for?

Ans:

ASA seems to have a number of niches. It is generally good at OLTP and can be used as a basis for a general database project. There are certainly examples of implementations supporting 100 or more users.

A major area for ASA databases is with applications that need to distribute the database with the application as a general storage area for internal components, but the database is not a major part of the deliverable. Sybase themselves have done this with the IQ meta data storage. Prior to release 11 of IQ, the meta data was stored in an ASE database. Now, with IQ 12, the meta data has moved to being stored in ASA. This makes the installation of IQ into production environments much simpler. ASA has excellent ODBC support, which makes it very attractive to tools oriented towards ODBC.

[View All Answers](#)

Question - 12:

On what platforms is ASA supported?

Ans:

- * Windows 95/98/ME, NT, 2000, CE
- * Novell NetWare
- * Solaris/SPARC
- * Solaris/Intel
- * IBM AIX
- * Linux (RedHat)
- * HP-UX

[View All Answers](#)

Question - 13:

What is ASA in Sybase?

Ans:

ASA is a fully featured DBMS with transactional integrity, automatic rollback and recovery, declarative RI, triggers and stored procedures.

While it comes out of Sybase's "Mobile and Embedded" division, it is NOT limited to "small, desktop applications". There are many ASA implementations supporting over 100 concurrent users. While not as scalable as ASE, it does offer SMP support and versions for various Unix flavors as well as Netware and NT/w2k. Multi-gigabyte databases are commonly used.

ASA offers a number of features that are not to be found in ASE:

- * row level BEFORE and AFTER triggers
- * long varchar and BLOB up to 2 gigabytes
- * varchar up to 32k
- * declarative RI with cascade actions
- * all character and decimal data is stored var-len, using only the space it needs

[View All Answers](#)

Question - 14:

What is this TDS protocol in Sybase?

Ans:

Tabular Data Streams or TDS is the name given to the protocol that is used to connect Sybase clients with Sybase servers. A specification for the protocol can be obtained from Sybase, I had a copy but cannot seem to find it now.

There is a project that is reverse engineering the protocol and building a set of libraries independent of either Sybase or Microsoft, but able to connect to either of their servers. FreeTDS is a considerable way down the line, although I do not believe that it is production ready yet!



[View All Answers](#)

Question - 15:

What is the difference between DB-lib and CT-lib in Sybase?

Ans:

Both DB-lib and CT-lib are libraries that implement the TDS protocol from the client side.

DB-lib

DB-lib was Sybase's first version. It was a good first attempt, but has/had a number of inconsistencies. There are, or possibly were, a lot of applications written using DB-lib. If you are about to start a new Open Client development, consider using CT-lib, it is the preferred choice. (What version of TDS does DB-lib, is it only 4.2?) Having said that you should use CT-lib for new developments, there is one case that this may not be true for and that is 2 phase commit. 2 phase commit is supported directly by DB-lib but is not supported directly by CT-lib.

CT-lib

CT-lib is a completely re-written version of Open Client that was released in the early '90s. The API is totally different from DB-lib, and is much more consistent. Applications written using DB-lib cannot simply be compiled using CT-lib, they need a significant amount of porting effort. CT-lib is newer, more consistent and, in several people's opinions, including mine, slightly longer winded. Having said that, the future of DB-lib is uncertain and is certainly not being developed any more, as a result all new apps should be written using CT-lib.

[View All Answers](#)

Question - 16:

What is Open Client in Sybase?

Ans:

Open Client is the interface (API) between client systems and Sybase servers. Fundamentally, it comes in two forms:

Runtime

The runtime version is a set of dynamic libraries (dlls on W32 platforms) that allow client applications to connect to Sybase and Microsoft servers, or, in fact, any server that implements the Tabular Data Streams (TDS) protocol. You need some form of Open Client in order to be able to connect to ASE in any way, shape or form. Even if you are running isql on exactly the same machine as ASE itself, communication will still be via Open Client. That is not to say that client to server communication on the same machine will go via the physical network, that decision is left entirely to the protocol implementation on the machine in question.

[View All Answers](#)

Question - 17:

What is my identity burn factor vulnerability right now in Sybase?

Ans:

Identities are created type numeric, scale 0, and precision X. A precision of 9 means the largest identity value the server will be able to process is 10^9-1 , or 1,000,000,000-1, or 999,999,999. However, when it comes to burning identities, the server will burn (based on the default value of 5000) .05% of 1,000,000,000 or 500,000 values in the case of a crash. (You may think an identity precision allowing for 1 Billion rows is optimistic, but I once saw a precision set at 14...then the database crashed and their identity values jumped 5 TRILLION. Needless to say they abandoned their original design. Even worse, SQL server defaults precision to 18 if you don't specify it upon table creation...that's a MINIMUM 10,000,000,000 jump in identity values upon a crash with the absolute minimum burn factor)

[View All Answers](#)

Question - 18:

How to pad with leading zeros an int or smallint in Sybase?

Ans:

By example:

```
declare @Integer int
```

```
/* Good for positive numbers only. */
```

```
select @Integer = 1000
```

```
select "Positives Only" =
```

```
right( replicate("0", 12) + convert(varchar, @Integer), 12)
```

```
/* Good for positive and negative numbers. */
```

```
select @Integer = -1000
```

```
select "Both Signs" =
```

```
substring( "- +", (sign(@Integer) + 2), 1) +
```

```
right( replicate("0", 12) + convert(varchar, abs(@Integer)), 12)
```

```
select @Integer = 1000
```

```
select "Both Signs" =
```

```
substring( "- +", (sign(@Integer) + 2), 1) +
```

```
right( replicate("0", 12) + convert(varchar, abs(@Integer)), 12)
```

```
go
```

Produces the following results:

Positives Only

```
-----  
000000001000  
Both Signs
```

```
-----  
-000000001000  
Both Signs
```

```
-----  
+000000001000
```

```
-----
```

[View All Answers](#)

Question - 19:

Divide by zero and nulls in Sybase?



Ans:

During processing, if a divide by zero error occurs you will not get the answer you want. If you want the result set to come back and null to be displayed where divide by zero occurs do the following:

```
1. select * from total_temp
2. go
field1 field2
-----
10 10
10 0
10 NULL
(3 rows affected)
1. select field1, field1/(field2*convert(int,
substring('1',1,abs(sign(field2))))) from total_temp
2. go
field1
-----
10 1
10 NULL
10 NULL
```

[View All Answers](#)

Question - 20:

How do I bcp null dates in Sybase?

Ans:

As long as there is nothing between the field delimiters in your data, a null will be entered. If there's a space, the value will be Jan 1, 1900.

You can use sed(1) to squeeze blanks out of fields:

```
sed -e 's/[ ]*/|/g' old_file > new_file
```

[View All Answers](#)

Question - 21:

How do I pipe the output of isql to a file in Sybase?

Ans:

Fortunately this is one of the simple ones. For batch sessions, simply use the -o command line switch to specify the name of the output file. For interactive sessions, the simplest method is turn on the capture functionality of your xterm window while running the query. If this isn't possible (i.e., you are not running on UNIX), then there are a couple of options:

* Use sqsh instead

* Use the tee(1) command which takes the output from a command and writes/appends it to a file, as well as displaying it on the screen. (Again, this may be UNIX-specific; Windows users should research the CYGWIN toolset for similar functionality.)

[View All Answers](#)

Question - 22:

How do I make isql secure in Sybase?

Ans:

isql uses the open/client libraries, which have no built in means of securing the packets that I know of. However, it is possible to use ssh to do all of the work for you. It is really quite straightforward. I saw this first published on the Sybase-L list by Tim Ellis, so all of the credit goes to him.

1. You will need a server running sshd that you have access to, which also has access to the ASE server.

2. Choose a port that you are going to make your secure connection from. Just like all ASE port selections it is totally arbitrary, but you if you were setting up a number of these, then you might want to think about a strategy. Regular server + 100 or something. Just make sure that it does not, and will not, clash with any of your regular servers.

3. Edit the interfaces file on the client side and set up a new server with an IP address of localhost and the port number you chose in the previous point. You might want to call it SERVER_SSH just to make sure that you know that it is the secure one. 4. Run the following ssh command:

```
ssh -2 -N -f -L port_chosen_above:remote_server:remote_port user@ssh.server.com
```

5. Connect to the server using isql -User -SSERVER_SSH

[View All Answers](#)

Question - 23:

Are there any alternatives to isql in Sybase?

Ans:

sqsh

In my opinion, and that of quite a lot of others, this is the most useful (direct) replacement for isql that exists. It combines the usefulness of a good shell with database interaction. Looking for the ability to page the output of a long command? Look no further. Need to search a result set using a regular expression? This is the tool for you.

Like isql, sqsh is a command line tool. It supports all of the features and switches of isql with myriad of its own. There is one feature that isql has the sqsh does not, and that is the ability to read the password as the first line of an input file. If you look at a lot of the examples above, the password is piped in, sqsh does not support this with the latest release. I am not sure if this is a deliberate feature or not.

A quick summary of its features:

1. command line editing;
2. command history;
3. ability to pipe to standard filters;
4. ability to redirect output to X window;
5. shell variables
6. background execution;



[View All Answers](#)

Question - 24:

How do I pipe the output of one isql to another?

Ans:

The following example queries sysdatabases and takes each database name and creates a string of the sort sp_helpdb dbname and sends the results to another isql. This is accomplished using bourne shell sh(1) and sed(1) to strip unwanted output:

```
#!/bin/sh
PASSWD=yuk
DSQUERY=GNARLY_HAIRBALL
echo "$PASSWD print "$PASSWD"
go
select 'sp_helpdb ' + name + char(10) + 'go'
from sysdatabases
go" | isql -U sa -S $DSQUERY -w 1000 |
sed -e '/affected/d' -e '/---/d' -e '/Password:/d' |
isql -U sa -S $DSQUERY -w 1000
```

[View All Answers](#)

Question - 25:

How to remove row affected and dashes?

Ans:

If you pipe the output of isql then you can use sed(1) to remove this extraneous output:

```
echo "$PASSWD
sp_who
go" | isql -U sa -S MY_SERVER | sed -e '/affected/d'
-e '/---/d'
```

If you simply wish to eliminate the row affected line use the set nocount on switch.

[View All Answers](#)

Question - 26:

Can I use a named pipe to bcp/dump data out or in in Sybase?

Ans:

System 10 and above.

If you would like to bcp copy from one table to a named pipe and compress:

```
1. %mknod bcp.pipe p
2. %compress sysobjects.Z &
3. %bcp master..sysobjects out bcp.pipe -c -U .. > bcp.pipe
4. Use ps(1) to determine when the compress finishes.
```

To bcp from my1db..dummy_table_1 to my2db..dummy_table_2:

```
1. %mknod bcp.pipe p
2. %bcp my2db..dummy_table_2 in bcp.pipe -c -U .. &
```

To avoid confusion between the above bcp and the next, you may choose to either use a separate window or redirect the output to a file.

```
3. %bcp my1db..dummy_table_1 out bcp.pipe -c -U
```

[View All Answers](#)

Question - 27:

How do I exclude a column in Sybase?

Ans:

Open/Client 11.1.1

Create a view based on the table that you want to exclude a column from and then bcp out from the view. Open/Client Versions Older Than 11.1.1

The documentation Utility programs for Unix describes the use of format files, including the field Server Column Order. Server Column Order must equal the colid of the column, or 0 if the host file field will not be loaded into any table column.

I don't know if anyone has got this feature to work. So, here is another way of removing the column. In your example, you want to remove the last column. I am going to include another example to remove the second column and include a fourth column. Why? Because it is harder. First example will deal with removing the last column.

[View All Answers](#)

Question - 28:

Would RPC or CT-Lib be better/faster?

Ans:

A netter responded stating that it was faster to ditch their algorithm and to apply a set based strategy:

The way to take your approach is to convert the row at a time processing (which is more traditional type of thinking) into a batch at a time (which is more relational type of thinking). Now I'm not trying to insult you to say that you suck or anything like that, we just need to dial you in to think in relational terms.

The idea is to do batches (or bundles) of rows rather than processing a single one at a time.

So let's take your example (since you didn't give exact values [probably out of kindness to save my eyeballs] I'll use your generic example to extend what I'm talking about):

Before:

if exists (select record) then

update record

else



insert record

New way:

1. Load all your rows into a table named new_stuff in a separate work database (call it work_db) and load it using bcp -- no third GL needed.

1. truncate new_stuff and drop all indexes

2. sort your data using UNIX sort and sort it by the clustered columns

3. load it using bcp

4. create clustered index using with sorted_data and any ancillary non-clustered index.

[View All Answers](#)

Question - 29:

Alternative to row at a time processing in Sybase?

Ans:

Someone asked how they could speed up their processing. They were batch updating/inserting gobs of information. Their algorithm was something as follows:

... In another case I do:

If exists (select record) then

update record

else

insert record

I'm not sure which way is faster or if it makes a difference. I am doing this for as many as 4000 records at a time (calling a stored procedure 4000 times!). I am interesting in knowing any way to improve this. The parameter translation alone on the procedure calls takes 40 seconds for 4000 records. I am using exec in DB-Lib.

[View All Answers](#)

Question - 30:

When should I execute an sp_recompile in Sybase?

Ans:

An sp_recompile should be issued any time a new index is added or an update statistics. Dropping an index will cause an automatic recompile of all objects that are dependent on the table.

The sp_recompile command simply increments the schemactnt counter for the given table. All dependent object counter's are checked against this counter and if they are different the SQL Server recompiles the object.

[View All Answers](#)

Question - 31:

When are stored procedures compiled in Sybase?

Ans:

Stored procedures are in a database as rows in sysprocedures, in the form of parse trees. They are later compiled into execution plans.

A stored procedures is compiled:

1. with the first EXECute, when the parse tree is read into cache

2. with every EXECute, if CREATE PROCEDURE included WITH RECOMPILE

3. with each EXECute specifying WITH RECOMPILE

4. if the plans in cache for the procedure are all in use by other processes

5. after a LOAD DATABASE, when all procedures in the database are recompiled

6. if a table referenced by the procedure can not be opened (using object id), when recompilation is done using the table's name 7. after a schema change in any referenced table, including:

1. CREATE INDEX or DROP INDEX to add/delete an index

2. ALTER TABLE to add a new column

3. sp_bindefault or sp_unbindefault to add/delete a default

4. sp_bindrule or sp_unbindrule to add/delete a rule

8. after EXECute sp_recompile on a referenced table, which increments sysobjects.schema and thus forces re-compilation

[View All Answers](#)

Question - 32:

What causes re-resolution of a stored procedure in Sybase?

Ans:

When a stored procedure references an object that is modified after the creation of the stored procedure, the stored procedure must be re-resolved. Re-resolution is the process of verifying the location of referenced objects, including the object id number. Re-resolution will occur under the following circumstances:

1. One of the tables used by the stored procedure is dropped and re-created.

2. A rule or default is bound to one of the tables (or unbound).

3. The user runs sp_recompile on one of the tables.

4. The database the stored procedure belongs to is re-loaded.

5. The database that one of the stored procedure's tables is located in is re-loaded.

6. The database that one of the stored procedure's tables is located in is dropped and re-created.

[View All Answers](#)

Question - 33:

What will cause the size of a stored procedure to grow in Sybase?

Ans:

Any of the following will result in a stored procedure to grow when it is recompiled:

1. One of the tables used in the procedure is dropped and re-created.

2. A new rule or default is bound to one of the tables or the user runs sp_recompile on one of the tables.

3. The database containing the stored procedure is re-loaded.

Other things causing a stored procedure to be re-compiled will not cause it to grow. For example, dropping an index on one of the tables used in the procedure or



doing EXEC WITH RECOMPILE.

The difference is between simple recompilation and re-resolution. Re-resolution happens when one of the tables changes in such a way that the query trees stored in sysprocedures may be invalid. The datatypes, column offsets, object ids or other parts of the tree may change. In this case, the server must re-allocate some of the query tree nodes. The old nodes are not de-allocated (there is no way to do this within a single procedure header), so the procedure grows. In time, trying to execute the stored procedure will result in a 703 error about exceeding the 64 page limit for a query.

[View All Answers](#)

Question - 34:

What are the different types of (All Page) locks in Sybase?

Ans:

First off, just to get it out of the way, Sybase does now support row level locking! (See Q6.1.11 for a description of the new features.) OK, that said and done, if you think you need row level locking, you probably aren't thinking set based -- see Q6.1.1 for set processing.

The SQL Server uses locking in order to ensure that sanity of your queries. Without locking there is no way to ensure the integrity of your operation. Imagine a transaction that debited one account and credited another. If the transaction didn't lock out readers/writers then someone can potentially see erroneous data.

Essentially, the SQL Server attempts to use the least intrusive lock possible, page lock, to satisfy a request. If it reaches around 200 page locks, then it escalates the lock to a table lock and releases all page locks thus performing the task more efficiently.

There are three types of locks:

- * page locks
- * table locks
- * demand locks

[View All Answers](#)

Question - 35:

How do I find the oldest open transaction in Sybase?

Ans:

```
select h.spid, u.name, p.cmd, h.name, h.starttime,
p.hostname, p.hostprocess, p.program_name
from master..syslogshold h,
master..sysprocesses p,
master..sysusers u
where h.spid = p.spid
and p.suid = u.suid
and h.spid != 0 /* not replication truncation point */
```

[View All Answers](#)

Question - 36:

How do I check if log truncation is blocked in Sybase?

Ans:

System 11 and beyond:

```
select h.spid, convert(varchar(20), h.name), h.starttime
from master..syslogshold h,
sysindexes i
where h.dbid = db_id()
and h.spid != 0
and i.id = 8 /* syslogs */
and h.page in (i.first, i.first+1) /* first page of log = page of oldest xact */
```

[View All Answers](#)

Question - 37:

The timestamp datatype in Sybase?

Ans:

The timestamp datatype is user-defined datatype supplied by Sybase, defined as:

varbinary(8) NULL

It has a special use when used to define a table column. A table may have at most one column of type timestamp, and whenever a row containing a timestamp column is inserted or updated the value in the timestamp column is automatically updated. This much is covered in the documentation.

What isn't covered is what the values placed in timestamp columns actually represent. It is a common misconception that timestamp values bear some relation to calendar date and/or clock time. They don't - the datatype is badly-named. SQL Server keeps a counter that is incremented for every write operation - you can see its current value via the global variable @@DBTS (though don't try and use this value to predict what will get inserted into a timestamp column as every connection shares the same counter.)

[View All Answers](#)

Question - 38:

How do I remove duplicate rows from a table in Sybase?

Ans:

There are a number of different ways to achieve this, depending on what you are trying to achieve. Usually, you are trying to remove duplication of a certain key due to changes in business rules or recognition of a business rule that was not applied when the database was originally built.

Probably the quickest method is to build a copy of the original table:

```
select *
into temp_table
from base_table
where l=0
```



Create a unique index on the columns that covers the duplicating rows with the ignore_dup_key attribute. This may be more columns that the key for the table.

```
create unique index temp_idx
on temp_table(col1, col2, ..., colN)
with ignore_dup_key
Now, insert base_table into temp_table.
insert temp_table
select * from base_table
```

[View All Answers](#)

Question - 39:

How to implement if-then-else in a select clause in Sybase?

Ans:

```
declare @val char(20)
select @val = 'grand'
select case when @val = 'small' then
'petit'
else
'grand'
end
```

However, quite a number of people are still using pre-11.5 implementations, including those people using the free 11.0.3.3 Linux release. In that case you can use the following recipe.

To implement the following condition in a select clause:

```
if @val = 'small' then
print 'petit'
else
print 'grand'
fi
```

[View All Answers](#)

Question - 40:

Fixing a Munged Log in Sybase?

Ans:

Sybase Technical Support states that this is extremely dangerous as it "jacks up the value of the timestamp" which is used for recovery purposes. This may cause potential database corruption if the system fails while the timestamp rolls over.

In 4.9.2, you could only run the dbcc rebuild_log command once and after that you would have to use bcp to rebuild the database

In System 10, you can run this command about 10 times.

In System 11 I (Pablo, previous editor) tried it about 20 times and no problem.

```
1. use master
2. go
1. select count(*) from your_database..syslogs
2. go
-----
some number
1. sp_configure "allow updates",1
2. go
1. reconfigure with override /* for system 10 and below only*/
2. go
1. begin tran
2. go
/* Save the following status to be used later... */
1. select saved_status=status from sysdatabases where name = "your_database"
2. go
1. update sysdatabases set status = -32768 where name = "your_database"
2. go
1. commit tran
2. go
1. shutdown
2. go
```

[View All Answers](#)

Question - 41:

Hiding your password to isql in Sybase?

Ans:

Single ASE on host

Script #1

Assuming that you are using bourne shell sh(1) as your scripting language you can put the password in a file and substitute the file where the password is needed.

```
#!/bin/sh
# invoke say ISQL or something...
(cat $HOME/dba/password_file
cat << EOD
dbcc ...
go
EOD ) | $SYBASE/bin/isql -Usa -w1000
Script #2
#!/bin/sh
umask 077
```



```
cat <<-endOfCat | isql -Umyuserid -Smyserver
mypassword
use mydb
go
sp_who
go
endOfCat
*
Script #3
#!/bin/sh
umask 077
cat <<-endOfCat | isql -Umyuserid -Smyserver
`myScriptForGeneratingPasswords myServer`
use mydb
go
sp_who
go
endOfCat
*
Script #3
#!/bin/sh
umask 077
isql -Umyuserid -Smyserver <<-endOfIsql
mypassword
use mydb
go
sp_who
go
endOfIsql
*
```

[View All Answers](#)

Question - 42:

Optimistic versus Pessimistic locking in Sybase?

Ans:

Solution #1

Use a timestamp on a header table that would be shared by the common data. This timestamp field is a Sybase datatype and has nothing to do with the current time. Do not attempt to do any operations on this column other than comparisons. What you do is when you grab data to present to the end-user, have the client software also grab the timestamp column value. After some time, if the end-user wishes to update the database, compare the client timestamp with what's in the database and if it's changed, then you can take appropriate action: again this is dictated by the business.

Problem #1

If users are sharing tables but columns are not shared, there's no way to detect this using timestamps because it's not sufficiently granular.

Solution #2

... Also are you coding to ct-lib directly? If so there's something that you could have done, or may still be able to do if you are using cursors.

[View All Answers](#)

Question - 43:

How do I force an index to be used in Sybase?

Ans:

System 11

In System 11, the binding of the internal ordinal value is alleviated so that instead of using the ordinal index value, the index name can be used instead:

select ... from my_table (index my_first_index)

Sybase 4.x and Sybase System 10

All indexes have an ordinal value assigned to them. For example, the following query will return the ordinal value of all the indexes on my_table:

```
select name, indid
from sysindexes
where id = object_id("my_table")
```

Assuming that we wanted to force the usage of index numbered three:

```
select ... from my_table(3)
```

Note: using a value of zero is equivalent to forcing a table scan. Whilst this sounds like a daft thing to do, sometimes a table scan is a better solution than heavy index scanning.

It is essential that all index hints be well documented. This is good DBA practice. It is especially true for Sybase System 10 and below.

One scheme that I have used that works quite well is to implement a table similar to sysdepends in the database that contains the index hints.

```
create table idxdepends
(
tblname varchar(32) not null -- Table being hinted
,depname varchar(50) not null -- Proc, trigger or app that
-- contains hint.
,idxname varchar(32) not null -- Index being hinted at
--,hintcount int null -- You may want to count the
-- number of hints per proc.
)
```

[View All Answers](#)

Question - 44:

Why place tempdb and log on low numbered devices?



Ans:

System 10 and below.

In System 10 and Sybase 4.X, the I/O scheduler starts at logical device (ldev) zero and works up the ldev list looking for outstanding I/O's to process. Taking this into consideration, the following device fragments (disk init) should be added before any others:

1. tempdb
2. log

[View All Answers](#)

Question - 45:

Should I run 32 or 64 bit ASE with Solaris?

Ans:

The following was written by the prior maintainer:

Sybase' first foray into 64-bit was with release 11.9.3. I do not know much about that release, but I seem to remember that it was always lagging behind its sister release of 11.9.2.

With ASE 12, Sybase have both 32-bit and 64-bit versions at the same release level. This is a big improvement, since it cuts out some concern that was prevalent with 11.9.3 as to why they were on different numbers. The releases are supposed to be identical in terms of functionality, save the fact that the 64-bit version can address more memory.

So, why not just be done with it and have just the one version? Firstly, I suppose that not everyone who can run Solaris has the capability to run the 64-bit version. There are still a lot of 32-bit Sparc chips around and a lot of people use them. It is also possible to run 32-bit Solaris on a 64-bit machine. In order to be able to run 64-bit Sybase you will have to be running 64-bit Solaris.

If you have a 64-bit environment, you still need to choose between which Sybase version to run. If you have more than 4G bytes of memory on your machine and you would like Sybase to take advantage of it, then the 64-bit version is for you. If not, then the word on the street, and from Sybase themselves, is that in identical environments, the 32-bit version runs slightly faster. I have heard a couple of explanations as to why this is so, but nothing that I find 100% convincing.

[View All Answers](#)

Question - 46:

What is Intimate Shared Memory or ISM?

Ans:

Intimate Shared Memory or ISM is a specific feature of Sun Solaris. The feature was developed so that when multiple processes (at OS level) try to access a shared memory region, they do not use multiple TLBs (Translation Lookaside Buffers) at OS kernel level. This saves lot of kernel memory space.

I don't think that does a whole lot for Sybase, more for Oracle I suppose. However, there is a side effect that is useful. If there is enough memory available on the machine, typically Solaris will not swap out process memory marked as ISM if it can possibly help it.

Swapping in Solaris is done in three phases, reserved, allocated and used. Locking the shared memory has the advantage of increasing performance. Of course, if there are lot's of processes on the machine and if new processes starve for memory, there is a potential that ISM will get swapped.

For performance reasons, it is worth ensuring that Sybase can allocated its shared memory segment using ISM. ASE tries by default to use ISM and will display an error message during start up if this is not possible. It is probably worth starting Sybase soon after a machine is rebooted to give it the best possible chance of using ISM.

[View All Answers](#)

Question - 47:

How do I set TS Role in order to run certain DBCCs?

Ans:

Some DBCC commands require that you set TS Role in order to run them. Here's how to set it:

Login to Server as sa and perform the following:

```
sp_role "grant", sybase_ts_role, sa
go
set role "sybase_ts_role" on
go
```

[View All Answers](#)

Question - 48:

Differences between clustered and non-clustered in Sybase?

Ans:

Clustered Index

A phone book lists everyone by last name. We have an A section, we have a B section and so forth. Within each section my phone book is clever enough to list the starting and ending names for the given page.

The phone book is clustered by last name.

create clustered index on phone_book (last_name)

It's fast to perform the following queries on the phone book:

- * Find the address of those whose last name is Cisar.
- * Find the address of those whose last name is between Even and Fa

Searches that don't work well:

- * Find the address of those whose phone number is 440-1300.
- * Find the address of those whose prefix is 440 In order to determine the answer to the two above we'd have to search the entire phone book. We can call that a table scan.

Non-Clustered Index

To help us solve the problem above we can build a non-clustered index.

create nonclustered index on phone_book (phone_number)

Our non-clustered index will be built and maintained by our Mythical ASE as follows:

1. Create a data structure that will house a phone_number and information where the phone_number exists in the phone book: page number and the row within the page. The phone numbers will be kept in ascending order.
2. Scan the entire phone book and add an entry to our data structure above for each phone number found.
3. For each phone number found, note along side it the page number that it was located and which row it was in.



[View All Answers](#)

Question - 49:

How do I capture a process's SQL in Sybase?

Ans:

This is a bit of a wide question, and there are many answers to it. Primarily, it depends on why you are trying to capture it. If you are trying to debug a troublesome stored procedure that is behaving differently in production to how it did in testing, then you might look at the DBCC method. Alternatively, if you wanted to do some longer term profiling, then auditing or one of the third party tools might be the way forward. If you know of methods that are not included here, please let me know.

DBCCs

If you want to look at the SQL a particular process is running at the moment, one of the following should work. Not sure which versions of ASE these work with. Remember to issue dbcc traceon(3604) before running any of the dbcc's so that you can see the output at your terminal.

* dbcc sqltext(spид)

* dbcc pss(0, spид, 0)

The first of the commands issues the SQL of the spид only a bit like this:

[27] BISCAY.master.1> dbcc sqltext(9)

[27] BISCAY.master.2> go

SQL Text: select spид, status, suser_name(suid), hostname,

db_name(dbid), cmd, cpu, physical_io, memusage,

convert(char(5),blocked) from master..sysprocesses

DBCC execution completed. If DBCC printed error messages, contact a user with

System Administrator (SA) role.

[28] BISCAY.master.1>

[View All Answers](#)

Question - 50:

What are the catches?

Ans:

Well, nothing truly comes for free. CIS is not a means of providing true load sharing, although you will find nothing explicitly in the documentation to tell you this. Obviously there is a performance hit which seems to affect cursors worst of all. CIS itself is implemented using cursors and this may be part of the explanation. OK, so how do I use it?

Easy! Add the remote server using sp_addserver. Make sure that you define it as type sql_server or ASEnterprise. Create an "existing" table using the definition of the remote table. Update statistics on this new "existing" table. Then simply use it in joins exactly as if it were a local table.

[View All Answers](#)

Question - 51:

What is CIS and how can I use it?

Ans:

CIS is the new name for Omni ASE. The biggest difference is that CIS is included with Adaptive Server Enterprise as standard. Actually, this is not completely accurate; the ability to connect to other ASEs and ASEs, including Microsoft's, is included as standard. If you need to connect to DB2 or Oracle you have to obtain an additional licence. So, what is it?

CIS is a means of connecting two servers together so that seamless cross-server joins can be executed. It is not just restricted to selects, pretty much any operation that can be performed on a local table can also be performed on a remote table. This includes dropping it, so be careful!

What servers can I connect to?

* Sybase ASE

* Microsoft SQL Server

* IBM DB2

* Oracle

[View All Answers](#)

Question - 52:

If the master device is full, how do I make the master database bigger?

Ans:

It is not possible to extend the master database across another device, so the following from Eric McGrane (recently of Sybase Product Support Engineering) should help.

* dump the current master database

* Pre-12.5 users use buildmaster to create a new master device with a larger size. ASE 12.5 users use dataserver to build the new, larger, master database.

* start the server in single user mode using the new master device

* login to the server and execute the following tsq1:

select * from sysdevices

* take note of the high value

* load the dump of the master you had just taken

* restart the server (as it will be shut down when master is done loading), again

in single user mode so that you can update system tables

* login to the server and update sysdevices setting high for master to the value that you noted previously

* shut the server down and start it back up, but this time not in single user mode.

[View All Answers](#)

Question - 53:

What is cmaxpktsz good for in Sybase?

Ans:



cmxpksz corresponds to the parameter "maximum network packet size" which you can see through sp_configure. I recommend only updating this value through sp_configure. If some of your applications send or receive large amounts of data across the network, these applications can achieve significant performance improvement by using larger packet sizes. Two examples are large bulk copy operations and applications reading or writing large text or image values. Generally, you want to keep the value of default network packet size small for users performing short queries, and allow users who send or receive large volumes of data to request larger packet sizes by setting the maximum network packet size configuration variable.

caddnetmem corresponds to the parameter "additional netmem" which you can see through sp_configure. Again, I recommend only updating this value through sp_configure. "additional netmem" sets the maximum size of additional memory that can be used for network packets that are larger than ASE's default packet size. The default value for additional netmem is 0, which means that no extra space has been allocated for large packets. See the discussion below, under maximum network packet size, for information on setting this configuration variable. Memory allocated with additional netmem is added to the memory allocated by memory. It does not affect other ASE memory uses.

[View All Answers](#)

Question - 54:

What is Trace Flags -- 5101 and 5102 in Sybase?

Ans:

5101

Normally, each engine issues and checks for its own Disk I/O on behalf of the tasks it runs. In completely symmetric operating systems, this behavior provides maximum I/O throughput for ASE. Some operating systems are not completely symmetric in their Disk I/O routines. For these environments, the server can be booted with the 5101 trace flag. While tasks still request disk I/O from any engine, the actual request to/from the OS is performed by engine 0. The performance benefit comes from the reduced or eliminated contention on the locking mechanism inside the OS kernel. To enable I/O affinity to engine 0, start ASE with the 5101 Trace Flag.

Your errorlog will indicate the use of this option with the message:

Disk I/O affinity to engine: 0

This trace flag only provides performance gains for servers with 3 or more datasever engines configured and being significantly utilized.

Use of this trace flag with fully symmetric operating systems will degrade performance!

[View All Answers](#)

Question - 55:

How to clear a log suspension?

Ans:

A connection that is in a log suspend state is there because the transaction that it was performing couldn't be logged. The reason it couldn't be logged is because the database transaction log is full. Typically, the connection that caused the log to fill is the one suspended. We'll get to that later.

In order to clear the problem you must dump the transaction log. This can be done as follows:

```
dump tran db_name to data_device
```

```
go
```

At this point, any completed transactions will be flushed out to disk. If you don't care about the recoverability of the database, you can issue the following command:

```
dump tran db_name with truncate_only
```

If that doesn't work, you can use the with no_log option instead of the with truncate_only.

After successfully clearing the log the suspended connection(s) will resume.

Unfortunately, as mentioned above, there is the situation where the connection that is suspended is the culprit that filled the log. Remember that dumping the log only clears out completed transaction. If the connection filled the log with one large transaction, then dumping the log isn't going to clear the suspension.

v System 10

What you need to do is issue an ASE kill command on the connection and then un-suspend it:

[View All Answers](#)

Question - 56:

How to Shrink a Database?

Ans:

It has historically been difficult to shrink any database except tempdb (because it is created fresh every boot time). The two methods commonly used have been:

1. Ensure that you have scripts for all your objects (some tools like SA Companion, DB Artisan or dbschema.pl from Sybperl can create scripts from an existing database), then bcp out your data, drop the database, recreate it smaller, run your scripts, and bcp in your data.

2. Use a third-party tool such as DataTool's SQL Backtrack, which in essence automates the first process.

This technote outlines a third possibility that can work in most cases.

An Unsupported Method to Shrink a Database

This process is fairly trivial in some cases, such as removing a recently added fragment or trimming a database that has a log fragment as its final allocation, but can also be much more complicated or time consuming than the script and bcp method.

General Outline

The general outline of how to do it is:

1. Make a backup of the current database
2. Migrate data from sysusages fragments with high lstart values to fragments with low lstart values.
3. Edit sysusages to remove high lstart fragments that no longer have data allocations.
4. Reboot ASE.

[View All Answers](#)

Question - 57:

How do I audit the SQL sent to the server?

Ans:

This does not seem to be well documented, so here is a quick means of auditing the SQL text that is sent to the server. Note that this simply audits the SQL sent to the server. So, if your user process executes a big stored procedure, all you will see here is a call to the stored procedure. None of the SQL that is executed as part of the stored procedure will be listed.

Firstly, you need to have installed Sybase security (which involves installing the sybsecurity database and loading it using the script \$\$SYBASE/scripts/installsecurity). Read the Sybase Security Administration Manual, you may want to enable a threshold procedure to toggle between a couple of audit tables. Be warned, that the default configuration option "suspend auditing when device full" is set to 1. This means that the server will suspend all normal SQL



operations if the audit database becomes full and the sso logs in and gets rid of some data. You might want to consider changing this to 0 unless yours is a particularly sensitive installation.

[View All Answers](#)

Question - 58:

What is the best value for cschedspins?

Ans:

It is crucial to understand that cschedspins is a tunable parameter (recommended values being between 1-2000) and the optimum value is completely dependent on the customer's environment. cschedspins is used by the scheduler only when it finds that there are no runnable tasks. If there are no runnable tasks, the scheduler has two options:

1. Let the engine go to sleep (which is done by an OS call) for a specified interval or until an event happens. This option assumes that tasks won't become runnable because of tasks executing on other engines. This would happen when the tasks are waiting for I/O more than any other resource such as locks. Which means that we could free up the CPU resource (by going to sleep) and let the system use it to expedite completion of system tasks including I/O.
2. Go and look for a ready task again. This option assumes that a task would become runnable in the near term and so incurring the extra cost of an OS context switch through the OS sleep/wakeup mechanism is unacceptable. This scenario assumes that tasks are waiting on resources such as locks, which could free up because of tasks executing on other engines, more than they wait for I/O.

[View All Answers](#)

Question - 59:

What is Trace Flag Definitions in Sybase?

Ans:

To activate trace flags, add them to the RUN_* script. The following example is using the 1611 and 260 trace flags. Note that there is no space between the '-T' and the traceflag, despite what is written in some documentation.

Use of these traceflags is not recommended by Sybase. Please use at your own risk.

```
% cd ~sybase/install
% cat RUN_BLAND
#!/bin/sh
#
# SQL Server Information:
# name: BLAND
# master device: /usr/sybase/dbf/BLAND/master.dat
# master device size: 25600
# errorlog: /usr/sybase/install/errorlog_BLAND
# interfaces: /usr/sybase
#
/usr/sybase/dataserver -d/usr/sybase/dbf/BLAND/master.dat
-sBLAND -e/usr/sybase/install/errorlog_BLAND -i/usr/sybase
-T1611 -T260
```

[View All Answers](#)

Question - 60:

How do I run multiple versions of Sybase on the same server?

Ans:

The answer to this relies somewhat on the platform that you are using.

Unix

ASE Versions Before 12.0

This applies to Unix and variants, Linux included. Install the various releases of software into logical places within your filesystem. I like to store all application software below a single directory for ease of maintenance, choose something like /sw. I know that some are keen on /opt and others /usr/local. It is all down to preference and server usage. If you have both Oracle and Sybase on the same server you might want /sw/sybase or /opt/sybase. Be a little careful here if your platform is Linux or FreeBSD. The standard installation directories for Sybase on those platforms is /opt/sybase. Finally, have a directory for the release, say ASE11_9_2 or simply 11.9.2 if you only ever have Sybase ASE running on this server. A little imagination is called for!

So, now you have a directory such as /sw/sybase/ASE/11.9.2 (my preferred choice :-), and some software installed under the directories, what now? In the most minimal form, that is all you need. Non of the environment variables are essential. You could quite successfully run

[View All Answers](#)

Question - 61:

How do I turn off marked suspect on my database in Sybase?

Ans:

One of your database is marked suspect as the SQL Server is coming up. Here are the steps to take to unset the flag.

Remember to fix the problem that caused the database to be marked suspect after switching the flag.

System 11

1. sp_configure "allow updates", 1
2. select status - 320 from sysdatabases where dbid = db_id("my_hosed_db") -- save this value.
3. begin transaction
4. update sysdatabases set status = -32768 where dbid = db_id("my_hosed_db")
5. commit transaction
6. shutdown
7. startserver -f RUN_*
8. fix the problem that caused the database to be marked suspect
9. begin transaction
10. update sysdatabases set status = saved_value where dbid = db_id("my_hosed_db")
11. commit transaction
12. sp_configure "allow updates", 0



13. reconfigure
14. shutdown
15. startserver -f RUN_*

[View All Answers](#)

Question - 62:

On startup, the transaction log of a database has filled and recovery has suspended, what can I do?

Ans:

You might find the following in the error log:

00:00000:00001:2000/01/04 07:43:42.68 server Can't allocate space for object 'syslogs' in database 'DBbad' because 'logsegment' segment is full/has no free extents.
If you ran out of space in syslogs, dump the transaction log. Otherwise, use ALTER DATABASE or sp_extendsegment to increase size of the segment.
00:00000:00001:2000/01/04 07:43:42.68 server Error: 3475, Severity: 21, State: 7
00:00000:00001:2000/01/04 07:43:42.68 server There is no space available in SYSLOGS for process 1 to log a record for which space has been reserved. This process will retry at intervals of one minute. The internal error number is -4.

[View All Answers](#)

Question - 63:

Why do my page locks not get escalated to a table lock after 200 locks?

Ans:

Several reasons why this may be happening.

* Are you doing the updates from within a cursor?

The lock promotion only happens if you are attempting to take out 200 locks in a single operation ie a single insert, update or delete. If you continually loop over a table using a cursor, locking one row at time, the lock promotion never fires. Either use an explicit mechanism to lock the whole table, if that is required, or remove the cursor replacing it with an appropriate join.

* A single operation is failing to escalate?

Even if you are performing a single insert, update or delete, Sybase only attempts to lock the whole table when the lock escalation point is reached. If this attempt fails because there is another lock which prevents the escalation, the attempt is aborted and individual page locking continues.

[View All Answers](#)

Question - 64:

How Do I Take Advantage of Table Partitioning with bcp in Sybase?

Ans:

You can take advantage of table partitioning with bcp in by following these guidelines:

1. Break up the data file into multiple files and simultaneously run each of these files as a separate bcp job against one table.

Running simultaneous jobs increases throughput.

2. Choose a number of partitions greater than the number of bcp jobs.

Having more partitions than processes (jobs) decreases the probability of page lock contention.

3. Use the batch option of bcp in. For example, after every 100 rows, force a commit. Here is the syntax of this command:

bcp table_name in filename -b100

Each time a transaction commits, ASE randomly assigns a new partition for the next insert. This, in turn, reduces the probability of page lock contention

[View All Answers](#)

Question - 65:

How to manually drop a table in Sybase?

Ans:

Occasionally you may find that after issuing a drop table command that the ASE crashed and consequently the table didn't drop entirely. Sure you can't see it but that sucker is still floating around somewhere.

Here's a list of instructions to follow when trying to drop a corrupt table:

1. sp_configure allow, 1

go

reconfigure with override

go

2. Write db_id down.

use db_name

go

select db_id()

go

3. Write down the id of the bad_table:

select id

from sysobjects

where name = bad_table_name

go

4. You will need these index IDs to run dbcc extntzap. Also, remember that if the table has a clustered index you will need to run extntzap on index "0", even though there is no sysindexes entry for that indid.

select indid

from sysindexes

where id = table_id

go

[View All Answers](#)

Question - 66:

Why not max out all my columns Sybase?

**Ans:**

People occasionally ask the following valid question:

v Suppose I have varying lengths of character strings none of which should exceed 50 characters.

Is there any advantage of last_name varchar(50) over this last_name varchar(255)?

That is, for simplicity, can I just define all my varying strings to be varchar(255) without even thinking about how long they may actually be? Is there any storage or performance penalty for this.

There is no performance penalty by doing this but as another netter pointed out:

If you want to define indexes on these fields, then you should specify the smallest size because the sum of the maximal lengths of the fields in the index can't be greater than 256 bytes.

and someone else wrote in saying:

Your data structures should match the business requirements. This way the data structure themselves becomes a data dictionary for others to model their applications (report generation and the like).

[View All Answers](#)

Question - 67:

What is a good example of a transaction in Sybase?

Ans:

This answer is geared for Online Transaction Processing (OTLP) applications.

To gain maximum throughput all your transactions should be in stored procedures - see Q1.5.8. The transactions within each stored procedure should be short and simple. All validation should be done outside of the transaction and only the modification to the database should be done within the transaction. Also, don't forget to name the transaction for sp_who - see Q9.2.

The following is an example of a good transaction:

```
/* perform validation */
select ...
if ... /* error */
/* give error message */
else /* proceed */
begin
begin transaction acct_addition
update ...
insert ...
commit transaction acct_addition
end
```

[View All Answers](#)

Question - 68:

What is a natural key in Sybase?

Ans:

Let me think back to my database class... okay, I can't think that far so I'll paraphrase... essentially, a natural key is a key for a given table that uniquely identifies the row. It's natural in the sense that it follows the business or real world need.

For example, assume that social security numbers are unique (I believe it is strived to be unique but it's not always the case), then if you had the following employee table:

```
employee:
ssn char(09)
f_name char(20)
l_name char(20)
title char(03)
```

Then a natural key would be ssn. If the combination of _name and l_name were unique at this company, then another natural key would be f_name, l_name. As a matter of fact, you can have many natural keys in a given table but in practice what one does is build a surrogate (or artificial) key.

The surrogate key is guaranteed to be unique because (wait, get back, here it goes again) it's typically a monotonically increasing value. Okay, my mathematician wife would be proud of me... really all it means is that the key is increasing linearly: i+1

[View All Answers](#)

Question - 69:

How to Save space when inserting rows monotonically?

Ans:

If the columns that comprise the clustered index are monotonically increasing (that is, new row key values are greater than those previously inserted) the following System 11 dbcc tune will not split the page when it's half way full. Rather it'll let the page fill and then allocate another page:

```
dbcc tune(ascinserts, 1, "my_table")
```

By the way, SyBooks is wrong when it states that the above needs to be reset when ASE is rebooted. This is a permanent setting.

To undo it:

```
dbcc tune(ascinserts, 0, "my_table")
```

[View All Answers](#)

Question - 70:

How to compute database fragmentation in Sybase?

Ans:

Command

```
dbcc traceon(3604)
```

go

```
dbcc tab(production, my_table, 0)
```

go

Interpretation



A delta of one means the next page is on the same track, two is a short seek, three is a long seek. You can play with these constants but they aren't that important.

A table I thought was unfragmented had L1 = 1.2 L2 = 1.8

A table I thought was fragmented had L1 = 2.4 L2 = 6.6

How to Fix

You fix a fragmented table with clustered index by dropping and creating the index. This measurement isn't the correct one for tables without clustered indexes. If your table doesn't have a clustered index, create a dummy one and drop it.

[View All Answers](#)

Question - 71:

How to implement database security?

Ans:

This is a brief run-down of the features and ideas you can use to implement database security: Logins, Roles, Users, Aliases and Groups

* sp_addlogin - Creating a login adds a basic authorisation for an account - a username and password - to connect to the server. By default, no access is granted to any individual databases.

* sp_adduser - A user is the addition of an account to a specific database.

* sp_addalias - An alias is a method of allowing an account to use a specific database by impersonating an existing database user or owner.

* sp_addgroup - Groups are collections of users at the database level. Users can be added to groups via the sp_adduser command.

A user can belong to only one group - a serious limitation that Sybase might be addressing soon according to the ISUG enhancements requests. Permissions on objects can be granted or revoked to or from users or groups.

* sp_role - A role is a high-level Sybase authorisation to act in a specific capacity for administration purposes. Refer to the Sybase documentation for details.

[View All Answers](#)

Question - 72:

How to view the log files in sybase?

Ans:

cd \$SYBASE/\$SYBASE_ASE/install

[View All Answers](#)

Question - 73:

Can I Run Any Transact-SQL Command on a Partitioned Table in Sybase?

Ans:

No. Once you have partitioned a table, you cannot use any of the following Transact-SQL commands on the table until you unpartition it:

1. drop table
2. sp_placeobject
3. truncate table
4. alter table table_name partition n

On releases of ASE prior to 11.5 it was not possible to create a clustered index on a partitioned table either.

[View All Answers](#)

Question - 74:

How Does Partition Assignment Relate to Transactions in Sybase?

Ans:

A user is assigned to a partition for the duration of a transaction. Assignment of partitions resumes with the first insert in a new transaction. The user holds the lock, and therefore partition, until the transaction ends.

For this reason, if you are inserting a great deal of data, you should batch it into separate jobs, each within its own transaction. See "How Do I Take Advantage of Table Partitioning with bcp in?", for details.

[View All Answers](#)

Question - 75:

Can Two Tasks Be Assigned to the Same Partition in Sybase?

Ans:

Yes. ASE randomly assigns partitions. This means there is always a chance that two users will vie for the same partition when attempting to insert and one would lock the other out.

The more partitions a table has, the lower the probability of users trying to write to the same partition at the same time.

[View All Answers](#)

Question - 76:

Must I Use Multiple Devices to Take Advantage of Partitions in Sybase?

Ans:

It depends on which type of performance improvement you want.

Table partitioning improves performance in two ways: primarily, by decreasing page contention for inserts and, secondarily, by decreasing i/o contention. "What Is Table Partitioning?" explains each in detail.

If you want to decrease page contention you do not need multiple devices. If you want to decrease i/o contention, you must use multiple devices.

[View All Answers](#)

Question - 77:

How Do I Create A Partitioned Table That Spans Multiple Devices in Sybase?



Ans:

Creating a partitioned table that spans multiple devices is a multi-step procedure. In this example, we assume the following:

* We want to create a new segment rather than using the default segment.

* We want to spread the partitioned table across three devices, data_dev1, data_dev2, and data_dev3.

Here are the steps:

1. Define a segment:

```
sp_addsegment newsegment, my_database, data_dev1
```

2. Extend the segment across all three devices:

```
sp_extendsegment newsegment, my_database, data_dev2
```

```
sp_extendsegment newsegment, my_database, data_dev3
```

3. Create the table on the segment:

```
create table my_table
```

```
(names, varchar(80) not null)
```

```
on newsegment
```

4. Partition the table:

```
alter table my_table partition 30
```

[View All Answers](#)

Question - 78:

Given a table which contains some rows with duplicate keys, how would you remove the duplicates?

Ans:

create temporary table with same table structure and then create unique index on this temporary table with option ignore_dup_row.

now insert data into temp table from the table in which duplicate records exists.

while inserting record into temp table, duplicate rows will get ignored.

Finally temp table will have unique records

Ex:

```
select distinct * from table_1 into temp_table_1
```

```
truncate table_1
```

```
insert into table_1
```

```
select * from temp_table_1
```

This is one of the ways to eliminate duplicates.

[View All Answers](#)

Question - 79:

What Is Table Partitioning in Sybase?

Ans:

Table partitioning is a procedure that creates multiple page chains for a single table.

The primary purpose of table partitioning is to improve the performance of concurrent inserts to a table by reducing contention for the last page of a page chain.

Partitioning can also potentially improve performance by making it possible to distribute a table's I/O over multiple database devices. Page Contention for Inserts

By default, ASE stores a table's data in one double-linked set of pages called a page chain. If the table does not have a clustered index, ASE makes all inserts to the table in the last page of the page chain.

When a transaction inserts a row into a table, ASE holds an exclusive page lock on the last page while it inserts the row. If the current last page becomes full, ASE allocates and links a new last page.

As multiple transactions attempt to insert data into the table at the same time, performance problems can occur. Only one transaction at a time can obtain an exclusive lock on the last page, so other concurrent insert transactions block each other.

[View All Answers](#)

Question - 80:

Can I Partition Any Table in Sybase ?

Ans:

No. You cannot partition the following kinds of tables:

1. Tables with clustered indexes (as of release 11.5 it is possible to have a clustered index on a partitioned table)

2. ASE system tables

3. Work tables

4. Temporary tables

5. Tables that are already partitioned. However, you can unpartition and then re-partition tables to change the number of partitions.

[View All Answers](#)

Question - 81:

How Do I Choose Which Tables To Partition in Sybase?

Ans:

You should partition heap tables that have large amounts of concurrent insert activity. (A heap table is a table with no clustered index.) Here are some examples:

1. An "append-only" table to which every transaction must write

2. Tables that provide a history or audit list of activities

3. A new table into which you load data with bcp in. Once the data is loaded in, you can unpartition the table. This enables you to create a clustered index on the table, or issue other commands not permitted on a partition table.

[View All Answers](#)

Question - 82:

Does Table Partitioning Require User-Defined Segments in Sybase?

Ans:



No. By design, each table is intrinsically assigned to one segment, called the default segment. When a table is partitioned, any partitions on that table are distributed among the devices assigned to the default segment.

In the example under "How Do I Create A Partitioned Table That Spans Multiple Devices?", the table sits on a user-defined segment that spans three devices.

[View All Answers](#)

Question - 83:

We have lost the sa password, what can we do?

Ans:

Remember Douglas Adams famous quote "Don't panic" is the first thing!

I know that most people use the 'sa' account all of the time, which is fine if there is only ever one dba administering the system. If you have more than one person accessing the server using the 'sa' account, consider using sa_role enabled accounts and disabling the 'sa' account. Funnily enough, this is obviously what Sybase think because it is one of the questions in the certification exams.

If you see that someone is logged using the 'sa' account or is using an account with 'sa_role' enabled, then you can do the following:

```
sp_configure "allow updates to system tables",1
go
update syslogins set password=null where name = 'sa'
go
sp_password null,newPassword
go
```

[View All Answers](#)

Question - 84:

What is Open Server in Sybase?

Ans:

Sybase open server is a vital component of Sybase and Microsoft Open Data Services . The Sybase open server is a major component for handling distributed database systems and thus forms a vital and powerful server side component of Sybase systems . One of the major feature of this is Sybase open server is interoperability.

[View All Answers](#)

Question - 85:

How do I set a password to be null?

Ans:

Since ASE 11 (I cannot remember if it was with the very first release of 11, but certainly not before) the password column in syslogins has been encrypted. Setting this column to NULL does not equate to that login having a NULL password. A NULL password still requires the correct binary string to be in place.

In release 12 and above, set the minimum password length to be 0 using sp_configure and give that account a null password, and all should be fine.

Before 12, it is not possible to set the minimum password length, so the direct approach is not possible. So, update the relevant record in syslogins setting the password column to be the same as that of an account with a NULL password already.

How does one get the correct binary value? When a new ASE is built, the 'sa' account has a NULL password to start with. Setting an account to have the same binary value as such an 'sa' account should work. Remember that the binary string is going to be specific to the operating system and the exact release of ASE etc. Obviously, if you have set the password of your 'sa' accounts to be something other than NULL (sensible move), then you are going to have to build yourself a dummy server just to get the correct string. If this is important to you, then you may wish to store the value somewhere safe once you have generated it.

[View All Answers](#)

Question - 86:

Does Sybase support Row Level Locking?

Ans:

With Adaptive Server Enterprise 11.9 Sybase introduced row level locking into its product. In fact it went further than that, it introduced 3 different locking levels:

- * All Pages Locking

This is the scheme that is implemented in all servers prior to 11.9. Here locks are taken out at the page level, which may included many rows. The name refers to the fact that all of the pages in any data manipulation statement are locked, both data and index.

- * Data Page Locking

The other two locking schemes are bundled together under the title Data Page Locking, referring to the fact that only data pages are ever locked in the conventional sense. Data Page Locking is divided into two categories

- * Data Only Locking

This locking scheme still locks a page at a time, including all of the rows contained within that page, but uses a new mechanism, called latches, to lock index pages for the shortest amount of time. One of the consequences of this scheme is that it does not update index pages. In order to support this Sybase has introduced a new concept, forwarded rows. These are rows that have had to move because they have grown beyond space allowed for them on the page they were created. 2002 bytes per page.

- * Row Level Locking

[View All Answers](#)

Question - 87:

How To Change varchar(m) to varchar(n) in Sybase?

Ans:

Before you start:

```
select max(datalength(column_name))
from affected_table
```

In other words, please be sure you're going into this with your head on straight.

How To Change System Catalogs

This information is Critical To The Defense Of The Free World, and you would be Well Advised To Do It Exactly As Specified:

use master



```
go
sp_configure "allow updates", 1
go
reconfigure with override /* System 10 and below */
go
use victim_database
go
select name, colid
from syscolumns
where id = object_id("affected_table")
go
begin tran
go
update syscolumns
set length = new_value
where id = object_id("affected_table")
and colid = value_from_above
go
update sysindexes
set maxlen = maxlen + increase/decrease?
where id=object_id("affected_table")
and indid = 0
go
```

[View All Answers](#)

Question - 88:

How can you change to get the decimal separator as comma instead of the decimal point (.)?

Ans:

Use stuff(convert(varchar,qty),charindex(".",convert(varchar,qty)),1,",")

[View All Answers](#)

Question - 89:

How Can I Run the ASE Upgrade Manually?

Ans:

How to Run the ASE Upgrade Manually

This document describes the steps required to perform a manual upgrade for ASE from release 4.x or 10.0x to release 11.02. In most cases, however, you should use sybinit to perform the upgrade.

BE SURE TO HAVE GOOD BACKUPS BEFORE STARTING THIS PROCEDURE.

1. Use release 11.0x sybinit to run the pre-eligibility test and Check Reserved words. Make any necessary changes that are mentioned in the sybinit log. The sybinit log is located in \$\$SYBASE/init/logs/logxxxx.yyy.

2. Use isql to connect to the 4.x or 10.0x ASE and do the following tasks:

a. Turn on option to allow updates to system tables:

1. sp_configure "allow updates", 1

2. go

b. Checkpoint all databases:

1. use "dbname"

2. go

1. checkpoint

2. go

c. Shutdown the 4.x or 10.0x ASE.

1. shutdown

2. go

3. Copy the interfaces file to the release 11.0x directory.

4. Set the environment variable SYBASE to the release 11.0x directory.

5. Copy the runserver file to the release 11.0x \$\$SYBASE/install directory.

6. Edit the \$\$SYBASE/install/RUN_SYBASE (runserver file) to change the path from the 4.x or 10.x dataserver directory to the new release 11.0x directory.

7. Start ASE using the new runserver file.

[View All Answers](#)

Question - 90:

Raw partitions or regular files?

Ans:

As always, this answer depends on the vendor's implementation on a cooked file system for the ASE... Performance Hit (synchronous vs asynchronous)

If on this platform, the ASE performs file system I/O synchronously then the ASE is blocked on the read/write and throughput is decreased tremendously.

The way the ASE typically works is that it will issue an I/O (read/write) and save the I/O control block and continue to do other work (on behalf of other connections). It'll periodically poll the workq's (network, I/O) and resume connections when their work has completed (I/O completed, network data xmit'd...).

Performance Hit (bcopy issue)

Assuming that the file system I/O is asynchronous (this can be done on SGI), a performance hit may be realized when bcopy'ing the data from kernel space to user space.

[View All Answers](#)

Question - 91:

How can I tell the datetime my Server started?

Ans:



Method #1

The normal way would be to look at the errorlog, but this is not always convenient or even possible. From a SQL session you find out the server startup time to within a few seconds using:

```
select "Server Start Time" = crdate
from master..sysdatabases
where name = "tempdb"
```

Method #2

Another useful query is:

```
select * from sysengines
```

which gives the address and port number at which the server is listening.

[View All Answers](#)

Question - 92:

How do I interpret the tli strings in the interface file?

Ans:

The tli string contained with Solaris interface files is a hex string containing port and IP address. If you have an entry SYBSRV

```
master tli tcp /dev/tcp x000204018196c4510000000000000000
```

Then it can be interpreted as follows:

x0002 no user interpretation (header info?)

0401 port number (1025 decimal)

81 first part of IP address (129 decimal)

96 second part of IP address (150 decimal)

c4 third part of IP address (196 decimal)

51 fourth part of IP address (81 decimal)

So, the above tli address is equivalent to

SYBSRV

```
master tcp ether sybhost 1025
```

[View All Answers](#)

Question - 93:

How to clear tempdb when the tempdb gets filled in Sybase?

Ans:

Dump tran with truncate_only if it is not cleared the go for select lct_admin("abort", {process_id [, database_id]})

[View All Answers](#)

Question - 94:

Explain about ETL data integration software?

Ans:

ETL supports data consolidation; it brings data from multiple sources and places it in a single repository. It brings scalable grid architecture for efficient management of data.

[View All Answers](#)

Question - 95:

Explain about Mobile enterprise application enablement?

Ans:

Custom based applications for mobiles, existing applications, etc can be brought onto a mobile phone with this suite. It provides core integration, messaging, synchronization, data management of mobile phone, etc. It enables information exchange and transportability between developers and enterprise systems.

[View All Answers](#)

Question - 96:

What is SQL Server and ASE?

Ans:

A ASE (SQL Server) is simply a Unix process. It is also known as the database engine. It has multiple threads to handle asynchronous I/O and other tasks. The number of threads spawned is the number of engines (more on this in a second) times five. This is the current implementation of Sybase System 10, 10.0.1 and 10.0.2 on IRIX 5.3.

Adaptive Server Enterprise (ASE) has long been noted for its reliability, low total cost of ownership and superior performance. With its latest version, ASE 15, it has been dramatically enhanced to deliver capabilities urgently needed by enterprises today. It lays the long-term foundation for strategic agility and continuing innovation in mission-critical environments. ASE 15 provides unique security options and a host of other new features that boost performance while reducing operational costs and risk. Find out how you can exploit new technologies such as grids and clusters, service-oriented architectures and real-time messaging.

[View All Answers](#)

Question - 97:

How to start/stop ASE when CPU reboots?

Ans:

Below is an example of the various files (on Irix) that are needed to start/stop an ASE. The information can easily be extended to any UNIX platform.

The idea is to allow as much flexibility to the two classes of administrators who manage the machine:

- * The System Administrator

- * The Database Administrator



Any errors introduced by the DBA will not interfere with the System Administrator's job.
With that in mind we have the system startup/shutdown file /etc/init.d/sybase invoking a script defined by the DBA:
/usr/sybase/sys.config/{start,stop}.sybase
/etc/init.d/sybase

[View All Answers](#)

Question - 98:

What is the command for quitting from Sybase?

Ans:

syb_quit(0) command can be used to exit from Sybase database .

[View All Answers](#)

Question - 99:

How do I move tempdb off of the Master Device?

Ans:

There used to be a section in the FAQ describing how to drop all of tempdb's devices physically from the master device. This can make recovery of the server impossible in case of a serious error and so it strongly recommended that you do not do this but simply drop the segments as outlined below.

Sybase TS Preferred Method of Moving tempdb off the Master Device.

This is the Sybase TS method of removing most activity from the master device:

1. Alter tempdb on another device:

1. alter database tempdb on ...

2. go

2. Use the tempdb:

1. use tempdb

2. go

3. Drop the segments:

1.sp_dropsegment "default", tempdb, master

2. go

1. sp_dropsegment "logsegment", tempdb, master

2. go

1. sp_dropsegment "system", tempdb, master

2. go

Note that there is still some activity on the master device. On a three connection test that I ran:

```
while ( 1 = 1 )
```

```
begin
```

```
create table #x (col_a int)
```

```
drop table #x
```

```
end
```

[View All Answers](#)

Question - 100:

How do I correct timeslice -201?

Ans:

Why Increase It?

Basically, it will allow a task to be scheduled onto the CPU for a longer time. Each task on the system is scheduled onto the CPU for a fixed period of time, called the timeslice, during which it does some work, which is resumed when its next turn comes around.

The process has up until the value of ctimemax (a config block variable) to finish its task. As the task is working away, the scheduler counts down ctimemax units. When it gets to the value of ctimemax - 1, if it gets stuck and for some reason cannot be taken off the CPU, then a timeslice error gets generated and the process gets infected.

On the other hand, ASE will allow a server process to run as long as it needs to. It will not swap the process out for another process to run. The process will decide when it is "done" with the server CPU. If, however, a process goes on and on and never relinquishes the server CPU, then Server will timeslice the process. Potential Fix

1. Shutdown the ASE

2. %buildmaster -dyour_device -yctimemax=2000

3. Restart your ASE. If the problem persists contact Sybase Technical Support notifying them what you have done already.

[View All Answers](#)

Question - 101:

What is RAID and Sybase?

Ans:

RAID

RAID means several things at once. It provides increased performance through disk striping, and/or resistance to hardware failure through either mirroring (fast) or parity (slower but cheaper).

RAID 0

RAID 0 is just striping. It allows you to read and write quickly, but provides no protection against failure.

RAID 1

RAID 1 is just mirroring. It protects you against failure, and generally reads and writes as fast as a normal disk. It uses twice as many disks as normal (and sends twice as much data across your SCSI bus, but most machines have plenty of extra capacity on their SCSI busses.)

Sybase mirroring always reads from the primary copy, so it does not increase read performance.

RAID 0+1

RAID 0+1 (also called RAID 10) is striping and mirroring together. This gives you the highest read and write performance of any of the raid options, but uses twice as many disks as normal.

[View All Answers](#)

**Question - 102:**

What is the relationship between JDBC and Sybase?

Ans:

JDBC is a Application programming Interface which helps Java Developers to interact with Database like Sybase via DML Statements (Insert, Update, Delete and so on).

[View All Answers](#)

Question - 103:

How to swap a db device with another?

Ans:

Here are four approaches. Before attempting any of the following: Backup, Backup, Backup.

Dump and Restore

1. Backup the databases on the device, drop the databases, drop the devices.

2. Rebuild the new devices.

3. Rebuild the databases (Make sure you recreate the fragments correctly - See Ed Barlow's scripts (<http://www.tiac.net/users/sqltech/>) for an sp that helps you do this if you've lost your notes. Failure to do this will possibly lead to data on log segments and log on data segments).

4. Reload the database dumps!

Twiddle the Data Dictionary - for brave experts only.

1. Shut down the server.

2. Do a physical dump (using dd(1), or such utility) of the device to be moved.

3. Load the dump to the new device

4. Edit the data dictionary (sysdevices.physname) to point to the new device.

[View All Answers](#)

Question - 104:

What is a Coalesce? What is the equivalent of Oracle/Db2 Coalesce Function in Sybase?

Ans:

Coalesce is also available in Sybase

This functions evaluates the listed expressions and returns the first non-null value. If all the expressions are null, coalesce returns null

USAGE: coalesce(expression, expression [, expression]...)

EXAMPLE: coalesce(initialqty,finalqty,middleqty)

Returns the first occurrence of a non-NULL value in either the initialqty or finalqty or middleqty

[View All Answers](#)

Question - 105:

Explanation about replication server software?

Ans:

Replication server software makes you to manage multiple data platforms very easily. Some of the features which it provides are heterogeneous replication and synchronization across various applications bi directionally. It provides services from the client to server and server to client.

[View All Answers](#)

Question - 106:

Explain about Open switch business continuity software?

Ans:

Open switch business software helps you to work on the back up system without any hassles making your primary database server secured. It gives you connectivity to the primary database from the secondary server even in the middle of the application.

[View All Answers](#)

Question - 107:

Explain about Mirror Activator Disk mirroring?

Ans:

This software works with storage replication systems and replicates the database to an available ASE or any database back up devices such as oracle. It increases the security of applications and also reduces the fail over time for data applications.

[View All Answers](#)

Question - 108:

What is procedure cache and data cache in Sybase?

Ans:

Procedure cache is the memory area where compiled query tree run e.g, procedure, batch query. Data Cache is the memory area where data that is required for the current queries running is bought from Disk on to memory for building result sets.

[View All Answers](#)

Question - 109:

Give the benefits of Data integration suite?

Ans:



Some of the benefits which you can get from a Data integration suite are

1) Flexibility: -This suite gives a greater flexibility to incorporate data from various sources such as web, file systems, etc and it can transfer this data to third party applications such as SQL, production systems, etc.

[View All Answers](#)

Question - 110:

Explain about Enterprise scalability and ROI of Data integration suite?

Ans:

Enterprise scalability provides increased capacity without any new addition of hardware. It increases the capacity by increasing the availability in the enterprise data layer.

ROI: - Streamlining of application projects and reducing data management costs, data federation increases ROI. Complexity and superior decision making choices make ROI best in its class.

[View All Answers](#)

Question - 111:

What are the vital features of Sybase?

Ans:

Absolutely there are numerous features but most important are as follows

1. Manages User Connection
 2. Enforces Business Rules
 3. Provide Recovery and BackUps
 4. Understand T-SQL
 5. Optimize Query Execution
 6. Maintain Databases with NO Redundancy
 7. Supports Data and Referential Integrity
- many more....

[View All Answers](#)

Question - 112:

What is the difference in storage type of numeric and decimal data types in sybase ?

Ans:

The numeric and decimal types are identical in all respects but one:

Only numeric types with a scale of 0 can be used for the IDENTITY column

[View All Answers](#)

Question - 113:

Explain about the Data integration suite features?

Ans:

Data integration suite offers advanced development and management tools. Integration of data is very easy and efficient as you can navigate your data to a specified address without much difficulty. The five important features are ETL, Data federation, Replications, Real time events and search events.

[View All Answers](#)

Question - 114:

Explain about Power Designer Data modeling software?

Ans:

Power designer has two unique tools known as Link and Sync technologies. This technology offers Data base designers to design efficient data models through which one can offer efficient design. Also Data base designers can share resources among them which will improve the productivity and performance.

[View All Answers](#)

Question - 115:

Explain about the mainframe connect tools from Sybase?

Ans:

This connectivity tool gives access to high performance tools which connects you very fast to the main frame data sources. It gives you read and write access to databases. It also has a facility by which you can connect the main frame data to secondary data storage facilities.

[View All Answers](#)

Question - 116:

Explain the benefits you can get from Mainframe connect?

Ans:

These are the following benefits you can get from mainframe connect: -

- 1) Client applications can be enabled. External software's can be connected to the mainframe connect for higher performance.
- 2) Mainframe applications connect to external data through LAN.
- 3) New applications can be created which enables you to connect to external data or source of information.

[View All Answers](#)

**Question - 117:**

Explain about the features of Sybase IQ?

Ans:

Sybase IQ stores data base in the form of columns of data rather in rows. This storage model improves the accessibility and search criteria for the data. Performance of the search criteria improves greatly because it searches only columns rather than the rows and columns of the database. This is also environment friendly.

[View All Answers](#)

Question - 118:

Explain the structure of Adaptive server enterprise?

Ans:

Structure of ASE installation hosts several databases and it comprises of one data server. Also system files occupy only one meta data. User databases are stored in the form of tables. Security and information privacy can be protected by setting permissions.

[View All Answers](#)

Question - 119:

Explain What is Update Statistics in Sybase?

Ans:

The update statistics command helps the server make the best decisions about which indexes to use when it processes a query, by providing information about the distribution of the key values in the indexes. The update statistics commands create statistics, if there are no statistics for a particular column, or replaces existing statistics if they already exist. The statistics are stored in the system tables systabstats and sysstatistics.

[View All Answers](#)

Question - 120:

Explain about Sybase 365?

Ans:

This product from Sybase is the world leader in messaging services through mobile. It distributes mobile content via MMS, SMS, UMTS and WAP. It process more than 8 billion messages monthly and with this product you have access to enterprise mobility anywhere.

[View All Answers](#)

Question - 121:

State some of the criticism about Sybase IQ?

Ans:

Sybase storage systems are optimized for Dataware housing applications where data cannot be modified can be modified by searching and accessing the data. Its column based search criteria gives high performance for search but it negates the writing capability to the data. Performance for row based relational based databases is worse because it is very difficult for updating and modifying them.

[View All Answers](#)

Question - 122:

State some of the products of Sybase?

Ans:

Sybase has worlds leading products some of them are

- * ASE
- * Avantgo
- * Power builder and power designer
- * SQL ianywhere
- * Sybase IQ
- * Sybase replication server
- * Sybase RTE
- * Sybase data federation
- * Main frame connect

[View All Answers](#)

Question - 123:

Explain about the Sybase group?

Ans:

Sybase exclusively focuses on mobilizing and managing information. It is known to be the third largest database management company after Oracle and IBM. Sybase products on Investment banking are well known to the world over and it offers many customized solutions for business intelligence.

[View All Answers](#)

Databases Programming Most Popular & Related Interview Guides

- 1 : [RDBMS Interview Questions and Answers.](#)
- 2 : [SQL Interview Questions and Answers.](#)
- 3 : [SSRS Interview Questions and Answers.](#)
- 4 : [Database Administrator \(DBA\) Interview Questions and Answers.](#)
- 5 : [Teradata Interview Questions and Answers.](#)
- 6 : [PostgreSQL Interview Questions and Answers.](#)
- 7 : [MongoDB Interview Questions and Answers.](#)
- 8 : [MS SQL Server Interview Questions and Answers.](#)
- 9 : [IBM DB2 Interview Questions and Answers.](#)
- 10 : [Data Modeling Interview Questions and Answers.](#)

Follow us on FaceBook

www.facebook.com/InterviewQuestionsAnswers.Org

Follow us on Twitter

<https://twitter.com/InterviewQA>

For any inquiry please do not hesitate to contact us.

Interview Questions Answers.ORG Team

[https://InterviewQuestionsAnswers.ORG/
support@InterviewQuestionsAnswers.ORG](https://InterviewQuestionsAnswers.ORG/support@InterviewQuestionsAnswers.ORG)