

Oracle System Architecture 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 Oracle System Architecture 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 [Oracle System Architecture 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 Oracle System Architecture category. To ensure quality, each submission is checked by our team, before it becomes live. This [Oracle System Architecture 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



Oracle System Architecture Interview Questions And Answers Guide.

Question - 1:

Tell me why oracle is called kernel package, components of oracle table, sql system variables, local and global variables?

Ans:

Please share your answers.

[View All Answers](#)

Question - 2:

Explain what is Shared SQL Area?

Ans:

Shared SQL areas are shared memory areas; any Oracle process can use a shared SQL area. The use of shared SQL areas reduces memory usage on the database server, thereby increasing system throughput. Shared SQL areas are aged out of the shared pool by way of a least recently used algorithm (similar to database buffers). To improve performance and prevent reparsing, you may want to prevent large SQL areas from aging out of the shared pool

[View All Answers](#)

Question - 3:

Explain what do Redo Log Buffers contain?

Ans:

Redo log buffers consists of change vectors where each vector contains information about which block is changed and where is the change done in the memory. This information will be useful when you are performing recovery and while writing to disk while checkpointing. There is a 1-to-1 correspondence with redo log buffers and data buffer cache buffers.

[View All Answers](#)

Question - 4:

What is Cursor in Oracle System Architecture?

Ans:

A cursor can be viewed as a "Pointer" into the result table of an SQL SELECT statement. Cursors are needed whenever the result contains more than one row. 3rd generation programs like COBOL and C cannot handle more than one row at a time. SQL produces a set-at-a-time. The cursor allows the set to be processed one row-at-a-time, just like a sequential file. COBOL and C are happy, SQL is happy. I'm happy.

[View All Answers](#)

Question - 5:

Tell me when does LGWR write to the database?

Ans:

LGWR writes when:

- 1)at commit
- 2)when 1/3rd full
- 3)when there is 1mb of redo
- 4)every 3 seconds
- 5)before DBWn writes.

[View All Answers](#)

Question - 6:

What is Shared Pool in Oracle System Architecture?

Ans:



Actually Shared pool is composed of two parts one is Library Cache and another Data Dictionary cache. The Library cache contains the shared sql areas, private sql areas, PL/SQL procedures and packages, and control structures such as locks and Library cache handles. The shared sql area contains the parse tree and execution plan where as the private sql area contains values for bind variables and runtime buffers. The data dictionary cache holds most recently used database dictionary informations.

[View All Answers](#)

Question - 7:

What is SGA in Oracle System Architecture?

Ans:

System Global Area (SGA) is component of instance which consist shared memory & control information of the instance. It consist of shared pool, database buffer cache, redo log buffer, large pool, Java pool.

[View All Answers](#)

Question - 8:

Explain when Does DBWR write to the database?

Ans:

DBWR writes when:

- 1) checkpoint occurs
- 2) dirty buffers reach threshold
- 3) there are no free buffers
- 4) timeout occurs
- 5) RAC ping request is made
- 6) when any tablespace is taken offline, read only, drop or truncated & when begin backup

[View All Answers](#)

Question - 9:

What is the function of checkpoint (CKPT)?

Ans:

The checkpoint (CKPT) process frees up limited memory space and ensures instance recovery. It is a background process that monitors the number of dirty buffers that have not been written to disk. Depending on the LOG_CHECKPOINT_INTERVAL parameter, it also gives the DBWR a wakeup call to write dirty buffer to disk. At a specified time, it updates all the data and control files with the new log file control number. This process is optionally enabled if parameter CHECKPOINT_PROCESS contains a TRUE value.

[View All Answers](#)

Question - 10:

Tell me what does LGWR do?

Ans:

LGWR is oracle background process. All changes which have been made in the database gets recorded in the redo log buffer. So when the user fires an commit statement, LGWR writes these changes into the redo log files, and also writes it into files when the redo log buffer is one third full... and also when buffer is nearly 1mb full, lgwr writes it into redo log files.

[View All Answers](#)

Question - 11:

Explain what constitute an ORACLE Instance?

Ans:

It's made up of Memory Structure and Background Processes. Memory Structure (In oracle SGA (system global area)) SGA consists of several memory areas: mandatory memory area 1. shared pools = library + data dictionary cache 2. database buffer cache 3. redo log buffer. Optional 1. large pool 2. java pool. BACKGROUND Process (mandatory) 1. smon (system monitor) 2. pmon (process monitor) 3. ckpt (check point) 4. dbwr (database writer) 5. lgwr (log writer). Please correct me, if I miss any point.

[View All Answers](#)

Question - 12:

Tell me what are the functions of SMON?

Ans:

System Monitor (SMON) performs instance recovery at instance start-up. In a multiple instance system (one that uses the Parallel Server), SMON of one instance can also perform instance recovery for other instance that have failed. SMON also cleans up temporary segments that are no longer in use and recovers dead transactions skipped during crash and instance recovery because of file-read or off-line errors. These transactions are eventually recovered by SMON when the tablespace or file is brought back on-line. SMON also coalesces free extents within the database to make free space contiguous and easier to allocate.

[View All Answers](#)

Question - 13:

What is the approach taken by Oracle to pass ACID test of a database transaction?

Ans:

Please share your answers.

[View All Answers](#)



Question - 14:

What is a cluster and how it is used?

Ans:

The deployment of a single database across a cluster of servers-providing unbeatable fault tolerance, performance and scalability

[View All Answers](#)

Question - 15:

Tell me what are the components of SGA?

Ans:

Components of SGA are :

- 1> Shared pool : consists of shared sql area and shared pl/sql area
- 2> Data Base buffer cache(DB cache) : contains most recently used blocks & dirty blocks(modified but not yet written to datafile)
- 3> Redolog Buffer : contains changed/redo data
- 4> Large pool : used only in pallellel sever environment
- 5> Java pool : used when you are using java stored procedures etc.

[View All Answers](#)

Question - 16:

Why is Oracle called kernel package?

Ans:

Oracle introduced the Unbreakable Enterprise Kernel for Oracle Linux during Sep 2010 which runs directly on Oracle LINUX 5 with fast, modern, reliable kernel that is optimized for Oracle software and hardware.

[View All Answers](#)

Question - 17:

Explain what is SGA? Where it is located? Why the system performance will be slow after loading the oracle?

Ans:

sga stands fro system global area.. which is the part of the oracle instance(includes memory struture and background processes). for faster executon and better performance everthing comes in the sga.so sga is located in oracle instance,, which itself consist of two parts.

- 1>database buffercache: which stores the recently executed database blocks,, when user fires a query oracle searches that block in dictionary and stores it into buffer cache.. so if the next time when the user fires the same query ... users gets output from the buffer cache itself.. no need for parsing again.. only fetch and execute will work.
- 2> redolog buffer: it stores all the changes made in the database,, and also all DMLs,, when user fires an commit statement,, lgwr back process writes it into the redolog files.. which can be used for recovery of database during instance crash .when u install an oracle in ur machine ,, system will become slow because of the unstructured sga configuration and also ur system configuration , once the oracle instance instants starts .. oracle itself will allocate the space from system .for example : if u have system confi 512 mb and if will install oracle with specified sga cong. obviously system performance will go low.so in order to increase the system performance,,, use the system configu should be well enough to run the oracle instance.... and also allocates enough space for sga....

[View All Answers](#)

Question - 18:

Tell me what do Database Buffers contain?

Ans:

Database buffer contains the data which server process reads recently from disk. Database buffer may have three parts: 1. free buffer 2. dirty buffer 3. pinned buffer Free buffer area is free to be used to store data. Dirty buffer area contains the data which is updated since read from disk, but not written to the disk till now. Pinned buffer is a area where processing on specific data for current transaction is going on. Database buffer may be of three types: 1. keep 2. recycle 3. default Database buffer may be divided into five different block size buffers: 2k,4k,8k,16k,32k and 64k also but supported by most of the servers (HP, IBM AIX, SOLARIS 10)

[View All Answers](#)

Question - 19:

What is PGA?

Ans:

Program Global Area. The PGA is a memory region containing data and control information for a single process (server or background). One PGA is allocated for each server process; the PGA is exclusive to that server process and is read and written only by Oracle code acting on behalf of that process. A PGA is allocated by Oracle when a user connects to an Oracle database and a session is created

[View All Answers](#)

Question - 20:

What is database security? How will you use it?

Ans:

That is Acl . Those levels are

- Manager,
- designer,
- editor,
- author,
- reader,
- depositor,
- no access



[View All Answers](#)

Question - 21:

Tell me how many Dispatcher Processes are created?

Ans:

Atleast one Dispatcher process is created for every communication protocol in use.

[View All Answers](#)

Question - 22:

What is the function of Lock (LCKn) Process in System Architecture?

Ans:

Lock (LCKn) are used for inter-instance locking when the ORACLE Parallel Server option is used.

[View All Answers](#)

Question - 23:

Do you know what is function of RECO?

Ans:

RECOVer (RECO) is used to resolve distributed transactions that are pending due to a network or system failure in a distributed database. At timed intervals, the local RECO attempts to connect to remote databases and automatically complete the commit or rollback of the local portion of any pending distributed transactions.

[View All Answers](#)

Question - 24:

Explain what is the function of Dispatcher (Dnnn)?

Ans:

Dispatcher (Dnnn) process is responsible for routing requests from connected user processes to available shared server processes and returning the responses back to the appropriate user processes.

[View All Answers](#)

Question - 25:

Explain what are functions of PMON?

Ans:

Process Monitor (PMON) performs process recovery when a user process fails. PMON is responsible for cleaning up the cache and freeing resources that the process was using. PMON also checks on dispatcher and server processes and restarts them if they have failed.

[View All Answers](#)

Question - 26:

Tell me what is the function of ARCH?

Ans:

Archiver (ARCH) copies the on-line redo log files to archival storage when they are full. ARCH is active only when a database's redo log is used in ARCHIVELOG mode.

[View All Answers](#)

Question - 27:

Tell me what is the maximum number of Lock Processes used?

Ans:

Though a single LCK process is sufficient for most Parallel Server systems, up to ten locks (LCK0, ..., LCK9) are used for inter-instance locking.

[View All Answers](#)

Question - 28:

Tell us what does DBWR do?

Ans:

Database writer writes modified blocks from the database buffer cache to the data files.

[View All Answers](#)

Question - 29:

Tell me what is a Multi-threaded Server Configuration?

Ans:

In a Multi-threaded Server Configuration many user processes share a group of server processes.

[View All Answers](#)

**Question - 30:**

Do you know what is a Parallel Server option in ORACLE?

Ans:

A configuration for loosely coupled systems where multiple instance share a single physical database is called Parallel Server.

[View All Answers](#)

Question - 31:

Tell me what is Dedicated Server Configuration?

Ans:

In a Dedicated Server Configuration a Server Process handles requests for a Single User Process.

[View All Answers](#)

Question - 32:

What is User Process in Oracle System Architecture?

Ans:

A user process is created and maintained to execute the software code of an application program. It is a shadow process created automatically to facilitate communication between the user and the server process.

[View All Answers](#)

Question - 33:

What is Server Process in Oracle System Architecture?

Ans:

Server Process handle requests from connected user process. A server process is in charge of communicating with the user process and interacting with ORACLE carry out requests of the associated user process.

[View All Answers](#)

Question - 34:

Tell me the names of ORACLE Background Process?

Ans:

Mandatory Processes DBWR -
Database Writer.LGWR -
Log WriterCKPT -
Check PointSMON -
System MonitorPMON -
Process Monitor Optional Process ARCH -
ArchiverRECO -
RecoverDnnn -
Dispatcher, andLCKn -
LockSnnn -
Server.

[View All Answers](#)

Question - 35:

Explain what are the different Oracle Kernel Layers?

Ans:

The Oracle call interface (OCI) The Oracle call interface is the lowest level at which client programs are intended to interact with Oracle. This interface is well documented and provides access to most of the functionality of Oracle, including advanced features such as object navigation, and sophisticated transaction and session control. Applications with advanced requirements have to use OCI directly, in order to access the features that are not available in Oracle's other development tools. The user program interface (UPI) OCI is based on the user program interface. There are some UPI facilities that are not yet available via OCI, and so some of the Oracle tools actually call this interface directly. Precompiler programs also call the user program interface, but indirectly via the SQLLIB library, which is an undocumented alternative to OCI. The Oracle program interface (OPI) The user program interface is the lowest layer of the client-side call stack, and the Oracle program interface is the highest layer of the server-side call stack. In most configurations, Net8 bridges the gap between UPI and OPI. However, in single-task executables there is no gap, and the UPI calls correspond directly to OPI calls. The compilation layer (KK) This is the top layer of the Oracle kernel proper. This layer is responsible for the parsing and optimization of SQL statements and for the compilation of PL/SQL program units. The execution layer (KX) This layer handles the binding and execution of SQL statements and PL/SQL program units. It is also responsible for the execution of recursive calls for trigger execution, and for the execution of SQL statements within PL/SQL program units. The distributed execution layer (K2) The distributed execution layer establishes the transaction branches for distributed transactions, and handles the management of the two-phase commit protocol. The network program interface (NPI) When remote objects are referenced in a SQL statement, the network program interface sends the decomposed statement components to the remote database instances and receives the data in return. The security layer (KZ) This layer is called by the compilation and execution layers to validate the required object and system privileges. The query layer (KQ) This layer provides rows to the higher layers. In particular, the query layer is responsible for caching rows from the data dictionary, for use by the security and compilation layers. The recursive program interface (RPI) The recursive program interface is used to populate the dictionary cache from the data dictionary. Row cache recursive SQL statements are executed in a separate call context, but are not parsed and optimized in the compilation layer. The access layer (KA) The access layer is responsible for access to database segments. This is the first layer of the lower half of the kernel. The data layer (KD) This layer is responsible for the management and interpretation of data within the blocks of database segments such as tables, clusters, and indexes. The transaction layer (KT) This layer is responsible for the allocation of transactions to rollback segments, interested transaction list changes within data blocks, changes to rollback segment blocks for undo generation, transaction control facilities such as savepoints, and read consistency. The transaction layer is also responsible for space management, both at the level of segment free lists and at the level of tablespace extent allocation. The cache layer (KC) The cache layer manages the database buffer cache. It uses operating system dependent facilities for data file I/O, provides concurrency control facilities for local access to the cache buffers, and provides parallel cache management (PCM) instance locking facilities for Oracle parallel server. The other main responsibility of the cache layer is the control of redo generation into the log buffer, and the writing of redo to the log files. The cache layer also caches control file information. The services layer (KS) The services layer provides low-level services that are



used by all the higher layers, such as error handling, debugging, and tracing facilities, as well as parameter control and memory services. In particular, the service layer is responsible for generic concurrency control facilities such as latches, event waits, enqueue locks, and instance locks. This layer is also responsible for the management of the data structures for background and user processes and sessions, as well as state objects, inter-process messages, and system statistics. The lock management layer (KJ) This layer is responsible for the locking used for synchronization and communication between the instances of a parallel server database. The generic layer (KG) The generic layer provides management for the generic data structures that are used by the higher layers, such as linked lists. Of particular interest are the library cache and the memory allocation heaps used for the shared pool and session memory. The operating system dependencies (S) Oracle uses operating system facilities for I/O, process scheduling, memory management, and other operations. The implementation details are operating system dependent, and so these details are isolated into a separate layer.

[View All Answers](#)

Question - 36:

Tell me what are the two types of Server Configurations?

Ans:

Dedicated Server Configuration and Multi-threaded Server Configuration.

[View All Answers](#)

Oracle Most Popular & Related Interview Guides

- 1 : [Oracle PL-SQL Interview Questions and Answers.](#)
- 2 : [Oracle DBA Interview Questions and Answers.](#)
- 3 : [Oracle D2K Interview Questions and Answers.](#)
- 4 : [OCI Interview Questions and Answers.](#)
- 5 : [Forms Reports Interview Questions and Answers.](#)
- 6 : [Oracle RMAN Interview Questions and Answers.](#)
- 7 : [Oracle Scenarios Interview Questions and Answers.](#)
- 8 : [Database Management Interview Questions and Answers.](#)
- 9 : [Oracle AOL Interview Questions and Answers.](#)
- 10 : [Oracle Developer 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)