

Oracle DBA Course

Mark InfoTech, DBA Administration course will help you acquire advanced skills in oracle to work as Oracle DBA. Students will develop, enhance and build confidence to manage any kind of large databases in real time world. Training will include basic Prerequisites and Complete Life cycle of Oracle database administration.

You Will Learn

High level, following were the areas we will cover.

- 1) Unix Overview and Basic Administration
- 2) Shell Scripting
- 3) Architecture & Configuration
- 4) Oracle Database Management
- 5) Oracle Object Management
- 6) Monitoring Oracle
- 7) Performance Management
- 8) Oracle Database Backup, Recovery Strategies, High Availability, RAC etc.
- 9) Oracle Internals, Oracle Wait Interface, SQL Tuning
- 10) Oracle Upgrade approaches and Methods.

Who Should Attend

Contact our office with your resume, for evaluation.

Course Content

1. Architecture & Configuration

1-1: Overview of Oracle DBA tasks

- Oracle as a flexible, complex & robust RDBMS
- The evolution of hardware and the relation to Oracle
- Different DBA job roles (VP of DBA, developer DBA, production DBA, database babysitter)
- The changing job role of the Oracle DBA
- Environment management (network, CPU, disk and RAM)
- Instance management (managing SGA regions)
- Oracle table and index management

1-2: Instance Architecture

- Instance vs. database
- Components of an instance
- Creating the OFA file structure (\$DBA, bdump, udump, pfile)

1-3: Oracle Instance Internal

- SGA vs. PGA
- Background processes
- Interfaces with server and disk I/O subsystem

1-4: Using SQL*Plus for DBA management

- Connecting and executing SQL
- Using the | as sysdba | syntax
- Overview of SQL*Plus DBA commands (startup, etc.)

1-5: Control files, UNDO and REDO management

- Explaining the use of control files
- Listing the Contents of the control File
- File locations for control Files
- Obtaining Control File Information

- Listing control file contents
- Displaying and Creating Undo segments
- Altering Undo Segments
- Determining the Number and Size of Undo segments
- Understanding flashback technology
- Troubleshooting Undo snapshot too old
- Redo log concepts for recovery
- Online redo log (log_buffer) online redo logs and archived redo logs
- Oracle ARCH and LGWR background processes
- Redo log dictionary queries
- Redo log switch frequency and performance
- Multiplexing the Online Redo Log Files
- Archiving the Oracle Redo Logs
- Recovery using the redo log files
- 1-6: User and privilege management
 - The three security methods (VPD, Grant security/role-based security, grant execute)
 - Creating New Database Users
 - Using pre-spawned Oracle connections
 - Auditing User activity
 - Identifying System and Object Privileges
 - Granting and Revoking Privileges
 - Creating and Modifying Roles
 - Displaying user security Information from the Data Dictionary

2. Oracle database management

- 2-1: Overview of instance management
 - Parameter files (init.ora, listener.ora, tnsnames.ora)
 - Rules for sizing SGA components
 - Automated Oracle memory management (AMM)
- 2-2: Initialization file management
 - Creating the init.ora file
 - Using spfile
 - Displaying init.ora values with v\$parameter
- 2-3: Oracle*Net configuration
 - Creating the listener.ora file
 - Creating the tnsnames.ora file
- 2-4: Data buffer configuration & sizing
 - Inside the Oracle data buffers
 - Using the KEEP pool
 - Monitoring buffer effectiveness
 - Using multiple blocksizes (multiple buffer pools)
- 2-5: Shared pool and PGA configuration & Sizing
 - Shared pool concepts and components
 - Understanding the library cache
 - Relieving shared pool contention
 - Overview of PGA for sorting and hash joins
 - Using sort_area_size, hash_area_size and pga_aggregate_target
- 2-6: Troubleshooting network connectivity
 - Verifying network connectivity with ping and tns ping
 - Testing database links

3. Oracle object management

- 3-1: Oracle tables, views and materialized views
 - Types of Oracle tables (regular, IOT, sorted hash clusters, nested tables)
 - Oracle Views
 - Oracle materialized views
- 3-2: Oracle indexes
 - Types of Oracle indexes (b-tree, bitmap, bitmap join index)
 - Creating B*-Tree, bitmap and function-based Indexes
 - Function-based indexes
 - Finding indexing opportunities
 - Index maintenance
- 3-3: Oracle constraints
 - Costs & benefits of constraints
 - Types of Oracle indexes constraints (check, not null, unique, PK, FK)
 - Cascading constraints
- 3-4: Schema, File & tablespace management
 - Describing the relationship between data files, Tablespaces and table
 - Understanding Oracle segments
 - Creating Tablespaces using the autoextend option
 - Changing the Size of Tablespaces alter database datafile command
 - Defining a TEMP tablespace
 - Changing the default storage Settings for a tablespace
 - Review of the storage parameters in DBA views (ASM, ASSM, pctfree, pctused and freelists).
 - Monitoring Chained rows (fetch continued rows)
 - Monitoring Insert and Update performance (pctused, APPEND)
- 3-5: Database Maintenance
 - Reason for reorgs chained rows, imbalanced freelists
 - Reorganizing Tables using Export and Import
 - Using CTAS to reorganize data
 - Index rebuilding
 - Backup & Recovery overview (hot & cold Backups, RMAN, block change tracking)
- 3-6: Oracle DBA Utilities
 - Data pump (Imp and exp utilities)
 - SQL *Loader
 - LogMiner
 - Flashback
 - DataGuard
 - Oracle DBA utilities Oracle dbms packages (dbms_redefinition)
 - Replication (Streams, multimaster, materialized views)

4. Monitoring Oracle

This section explores the methods used for monitoring all active components of the Oracle database.

- 4-1: Dictionary and v\$ views
 - The dba_, all_ and user_ structures
 - Querying the tables, indexes, and segments views
 - Querying the AWR (STATSPACK) tables
- 4-2: Table & index monitoring
 - Monitoring table extents and fragmentation
 - Using the dba_tables and dba_segments views
 - Monitoring table CBO statistics
 - Monitoring table extents and fragmentation

- Locating chained rows
- Monitoring table & index growth
- Monitoring index usage
- Monitoring index fragmentation
- Locating un-used indexes
- Identifying IOT candidates
- Reorganizing Indexes with alter index rebuild
- Dropping Indexes
- Getting Index Information from the Data Dictionary
- 4-3: workload & trend monitoring
 - Oracle automated workload tools
 - Using v\$dbh to monitor buffer activity
 - Using v\$sql and v\$sql_plan
- 4-4: Instance monitoring
 - Monitoring with the AWR and STATSPACK
 - Creating a time-series performance report
 - Using www.statspackanalyzer.com
 - Scripts for AWR and STATSPACK
 - Plotting performance data (WISE, Excel)
 - Finding performance trends and signatures
- 4-5: Oracle environment monitoring=
 - Displaying and managing Oracle sessions (v\$session, v\$process)
 - Using AWR to monitor disk, network and CPU consumption
 - Monitoring the alert log
 - Oracle trace/dump files
- 4-6: STATSPACK and AWR performance management
 - Installing STATSPACK
 - Running STATSPACK reports
 - Interpreting a STATSPACK report
 - Getting time series reports with STATSPACK
 - Finding performance signatures with STATSPACK

5. Performance Management and Various Tools

This section explores the methods used for performance management in Oracle and shows tips and scripts for monitoring all components of any Oracle database. You will also learn the proper action to take when any area of Oracle becomes a bottleneck.

- 5-1: Bottleneck performance analysis
 - Drill-down into AWR reports
 - Top-5 timed events
 - External Server Bottlenecks (Network, I/O, RAM, CPU)
 - Network troubleshooting
- 5-2: Instance Tuning
 - Changing init.ora optimizer parameters (index_optimizer_cost_adj, optimizer_mode)
 - Managing region parameters (shared_pool_size, db_cache_size)
 - Understanding instance contention (e.g. Buffer busy waits, library cache contention)
- 5-3: SQL and CBO behavior
 - Introduction to cost-based optimization
 - Changing the default optimizer modes

- Optimizer parameters
- Dynamic sampling
- Collecting table and index statistics (dbms_stats)
- Using column histograms and skewonly
- 5-4: Tracing SQL Execution
 - Using EXPLAIN PLAN
 - Using set autotrace
 - Interpreting EXPLAIN PLAN Output
 - Using TKPROF / SQL*Trace
- 5-5: SQL Execution Internals
 - Review of Basic joining methods
 - Merge join
 - Hash Join
 - Nested Loop join
 - Advanced SQL operators
 - Between operator
- 5-6: SQL Tuning
 - Using hints to improve SQL performance
 - Using parallel query to improve performance
 - SQL reusability within the library cache
 - Table high-water mark
 - Table striping and table partitions
 - Using indexes to improve performance
 - Identifying full-table scans
 - Re-writing SQL queries
 - Tuning sub-queries
- 6-1: Oracle High Availability tools
 - Continuous availability and disaster recovery
 - Quantifying the cost of unplanned downtime
 - Oracle multi-master replication
 - DataGuard
 - Oracle Streams
 - Real Application Clusters
- 6-2: Backup & Recovery
 - OS-level backups
 - Hardware-level backup & recovery
 - Block-level change tracking
 - Disk mirroring
 - Backup & recovery and RAID level
 - Oracle-level backups (expdp & RMAN)
 - Hot vs. Cold backups

Some of the Performance Course Slides.....

