

# ADMINISTRACIÓN ORACLE

## Descripción General

Curso orientado a usuarios de servidor de bases de datos Oracle que requiere aprender y dominar los procesos de instalación, configuración, puesta a punto, respaldo y recuperación de cluster de servidores Oracle

## Objetivo

Que el estudiante adquiera los conocimientos y habilidades necesaria para instalar, configurar y mantener un cluster de servidores Oracle con la tecnología RAC

## Duración

80 Horas

## Requisitos

## Chapter 1: Installing the Oracle Binaries

- 1.1. Understanding the OFA
- 1.2. Oracle Inventory Directory
- 1.3. Oracle Base Directory
- 1.4. Oracle Home Directory
- 1.5. Oracle Network Files Directory
- 1.6. Automatic Diagnostic Repository
- 1.7. Installing Oracle

Step 1. Create the OS Groups and User

Step 2. Ensure That the OS Is Adequately Configured

Step 3. Obtain the Oracle Installation Software

Step 4. Unzip the Files

Step 5: Creating oraInst.loc File

Step 6. Configure the Response File, and Run the Installer

1. Step 7. Troubleshoot Any Issues
  - 1.1. Installing with a Copy of an Existing Installation

Step 1. Copy the Binaries, Using an OS Utility

1. Step 2. Attach the Oracle Home Upgrading Oracle Software
  - 1.1. Reinstalling After Failed Installation
  - 1.2. Applying Interim Patches
  - 1.3. Installing Remotely with the Graphical Installer

Step 1. Install X Software and Networking Utilities on the Local PC

Step 2. Start an X Session on the Local Computer

Step 3. Copy the Oracle Installation Media to the Remote Server

Step 4. Run the xhost Command

Step 5. Log In to the Remote Computer from X

Step 6. Ensure that the DISPLAY Variable Is Set Correctly on the Remote Computer

Step 7. Execute the runInstaller Utility

Step 8. Troubleshoot

Chapter 2: Implementing a Database Setting OS Variables

- 1.1. A Manually Intensive Approach
- 1.2. Oracle's Approach to Setting OS Variables
- 1.3. My Approach to Setting OS Variables

- |   |  |
|---|--|
| 1.4. Creating a Database  | 1.14. Chapter 3: Configuring an Efficient Environment                      |
| Step 1. Set the OS Variables                                      | 1.15. Customizing Your OS Command Prompt                                   |
| Step 2: Configure the Initialization File                         | 1.16. Customizing Your SQL Prompt  |
| Step 3: Create the Required Directories                           | 1.17. Creating Shortcuts for Frequently Used Commands                      |
| Step 4: Create the Database                                       | 1.18. Using Aliases  |
| 1. Step 5. Create a Data Dictionary                               | 1.19. Using a Function   |
| 1.1. Configuring and Implementing the Listener                    | 1.20. Rerunning Commands Quickly   |
| 1.2. Manually Configuring a Listener                              | 1.21. Rerunning Commands Quickly Scrolling with the Up and Down Arrow Keys |
| 1.3. Implementing a Listener with the Net Configuration Assistant | 1.22. Using Ctrl+P and Ctrl+N  |
| 1.4. Connecting to a Database through the Network                 | 1.23. Listing the Command History  |
| 1.5. Creating a Password File                                     | 1.24. Searching in Reverse   |
| 1.6. Starting and Stopping the Database                           | 1.25. Setting the Command Editor   |
| 1.7. Understanding OS Authentication                              | 1.26. Developing Standard Scripts  |
| 1.8. Starting the Database  | 1.27. dba_setup  |
| 1.9. Stopping the Database  | 1.28. dba_fcns   |
| 1.10. Using the dbca to Create a Database                         | 1.29. tbsp_chk.bsh   |
| 1.11. Dropping a Database   | 1.30. conn.bsh   |
| 1.12. How Many Databases on One Server?                           | 1.31. filesp.bsh   |
| 1.13. Understanding Oracle Architecture                           | 1.32. login.sql  |

- 1.33. top.sql
- 1.34. lock.sql
- 1.35. users.sql
- 1.36. Organizing Scripts

Step 1: Create Directories

Step 2: Copy Files to Directories

Step 3: Configure the Startup File

Chapter 4: Tablespaces and Data Files

- 1.1. Understanding the First Five
- 1.2. Understanding the Need for More
- 1.3. Creating Tablespaces
- 1.4. Renaming a Tablespace
- 1.5. Controlling the Generation of Redo
- 1.6. Changing a Tablespace's Write Mode
- 1.7. Dropping a Tablespace
- 1.8. Using Oracle Managed Files
- 1.9. Creating a Bigfile Tablespace
- 1.10. Enabling Default Table Compression within a Tablespace
- 1.11. Displaying Tablespace Size

- 1.12. Altering Tablespace Size
- 1.13. Toggling Data Files Offline and Online
- 1.14. Renaming or Relocating a Data File
- 1.15. Performing Online Data File Operations
- 1.16. Performing Offline Data File Operations

Chapter 5: Managing Control Files, Online Redo Logs, and Archiving

- 1.1. Managing Control Files
- 1.2. Viewing Control File Names and Locations
- 1.3. Adding a Control File
- 1.4. Moving a Control File
- 1.5. Removing a Control File
- 1.6. Managing Online Redo Logs
- 1.7. Displaying Online Redo Log Information
- 1.8. Determining the Optimal Size of Online Redo Log Groups
- 1.9. Determining the Optimal Number of Redo Log Groups
- 1.10. Adding Online Redo Log Groups
- 1.11. Resizing and Dropping Online Redo Log Groups
- 1.12. Adding Online Redo Log Files to a Group
- 1.13. Removing Online Redo Log Files from a Group
- 1.14. Moving or Renaming Redo Log Files

[www.cursoslinux.com.mx](http://www.cursoslinux.com.mx)

[ventas@plct.com.mx](mailto:ventas@plct.com.mx)

PLCT S.A. de C.V.

Tel.: 55 4522 7839/55 1800 7696/7224447684

- 1.15. Implementing Archivelog Mode
- 1.16. Making Architectural Decisions
- 1.17. Setting the Archive Redo File Location
- 1.18. Thinking “Un-Oraclethodox” FRA Thoughts
- 1.19. Enabling Archivelog Mode
- 1.20. Disabling Archivelog Mode
- 1.21. Reacting to a Lack of Disk Space in Your Archive Log Destination
- 1.22. Backing Up Archive Redo Log Files

#### Chapter 6: Users and Basic Security

- 1.1. Managing Default Users
- 1.2. Locking Accounts and Expiring Passwords
- 1.3. Identifying DBA-Created Accounts
- 1.4. Checking Default Passwords
- 1.5. Creating Users
- 1.6. Choosing a Username and Authentication Method
- 1.7. Assigning Default Permanent and Temporary Tablespaces
- 1.8. Modifying Passwords
- 1.9. Logging In as a Different User
- 1.10. Modifying Users

- 1.11. Dropping Users
- 1.12. Enforcing Password Security and Resource Limits
- 1.13. Basic Password Security
- 1.14. Password Strength
- 1.15. Limiting Database Resource Usage
- 1.16. Managing Privileges
- 1.17. Assigning Database System Privileges
- 1.18. Assigning Database Object Privileges
- 1.19. Grouping and Assigning Privileges

#### Chapter 7: Tables and Constraints

- 1.1. Understanding Table Types
- 1.2. Understanding Data Types
- 1.3. Character
- 1.4. Numeric
- 1.5. Date/Time
- 1.6. RAW
- 1.7. ROWID
- 1.8. LOB
- 1.9. Extended Character Types
- 1.10. Creating a Table



- 1.11. Creating a Heap-Organized Table
- 1.12. Implementing Virtual Columns
- 1.13. Implementing Invisible Columns
- 1.14. Making Read-Only Tables
- 1.15. Understanding Deferred Segment Creation
- 1.16. Creating a Table with an Autoincrementing (Identity) Column
- 1.17. Allowing for Default Parallel SQL Execution
- 1.18. Compressing Table Data
- 1.19. Avoiding Redo Creation
- 1.20. Creating a Table from a Query
- 1.21. Modifying a Table
- 1.22. Obtaining the Needed Lock
- 1.23. Renaming a Table
- 1.24. Adding a Column
- 1.25. Altering a Column
- 1.26. Renaming a Column
- 1.27. Dropping a Column
- 1.28. Displaying Table DDL
- 1.29. Dropping a Table
- 1.30. Undropping a Table
- 1.31. Removing Data from a Table
- 1.32. Using DELETE
- 1.33. Using TRUNCATE
- 1.34. Viewing and Adjusting the High-Water Mark
- 1.35. You need to be aware of a couple of performance related issues regarding the high-water mark
- 1.36. Tracing to Detect Space Below the High-Water Mark
- 1.37. Using DBMS\_SPACE to Detect Space Below the High-Water Mark
- 1.38. Selecting from Data Dictionary Extents View
- 1.39. Lowering the High-Water Mark
- 1.40. Creating a Temporary Table
- 1.41. Creating an Index-Organized Table
- 1.42. Managing Constraints
- 1.43. Creating Primary Key Constraints
- 1.44. Enforcing Unique Key Values
- 1.45. Creating Foreign Key Constraints
- 1.46. Checking for Specific Data Conditions
- 1.47. Enforcing Not Null Conditions

- 1.48. Disabling Constraints
- 1.49. Enabling Constraints
- 1. Chapter 8: Indexes
  - 1.1. Deciding When to Create an Index
  - 1.2. Proactively Creating Indexes
  - 1.3. Reactively Creating Indexes
  - 1.4. Planning for Robustness
  - 1.5. Determining Which Type of Index to Use
  - 1.6. Estimating the Size of an Index Before Creation
  - 1.7. Creating Separate Tablespaces for Indexes
  - 1.8. Creating Portable Scripts
  - 1.9. Establishing Naming Standards
  - 1.10. Creating Indexes
  - 1.11. Creating B-tree Indexes
  - 1.12. Creating Concatenated Indexes
  - 1.13. Implementing Function-Based Indexes
  - 1.14. Creating Unique Indexes
  - 1.15. Implementing Bitmap Indexes
  - 1.16. Creating Bitmap Join Indexes
  - 1.17. Implementing Reverse-Key Indexes
  - 1.18. Creating Key-Compressed Indexes
  - 1.19. Parallelizing Index Creation
  - 1.20. Avoiding Redo Generation When Creating an Index
  - 1.21. Implementing Invisible Indexes
  - 1.22. Maintaining Indexes
  - 1.23. Renaming an Index
  - 1.24. Displaying Code to Recreate an Index
  - 1.25. Rebuilding an Index
  - 1.26. Making Indexes Unusable
  - 1.27. Monitoring Index Usage
  - 1.28. Dropping an Index
  - 1.29. Indexing Foreign Key Columns
  - 1.30. Implementing an Index on a Foreign Key Column
  - 1.31. Determining if Foreign Key Columns Are Indexed
- Chapter 9: Views, Synonyms, and Sequences
  - 1.1. Implementing Views
  - 1.2. Creating a View
  - 1.3. Checking Updates
  - 1.4. Creating Read-Only Views
  - 1.5. Updatable Join Views

- 1.6. Creating an INSTEAD OF Trigger
- 1.7. Implementing an Invisible Column
- 1.8. Modifying a View Definition
- 1.9. Displaying the SQL Used to Create a View
- 1.10. Renaming a View
- 1.11. Dropping a View
- 1.12. Managing Synonyms
- 1.13. Creating a Synonym
- 1.14. Creating Public Synonyms
- 1.15. Dynamically Generating Synonyms
- 1.16. Displaying Synonym Metadata
- 1.17. Renaming a Synonym
- 1.18. Dropping a Synonym
- 1.19. Managing Sequences
- 1.20. Creating a Sequence
- 1.21. Using Sequence Pseudocolumns
- 1.22. Autoincrementing Columns
- 1.23. Implementing Multiple Sequences That Generate Unique Values
- 1.24. Creating One Sequence or Many

- 1.25. Viewing Sequence Metadata
- 1.26. Renaming a Sequence

#### Chapter 10: Data Dictionary Fundamentals

- 1.1. Data Dictionary Architecture
- 1.2. Static Views
- 1.3. Dynamic Performance Views
- 1.4. A Different View of Metadata
- 1.5. A Few Creative Uses of the Data Dictionary
- 1.6. Derivable Documentation
- 1.7. Displaying User Information
- 1.8. Displaying Table Row Counts
- 1.9. Showing Primary Key and Foreign Key Relationships
- 1.10. Displaying Object Dependencies

#### Chapter 11: Large Objects

- 1.1. Describing LOB Types
- 1.2. Illustrating LOB Locators, Indexes, and Chunks
- 1.3. Distinguishing Between BasicFiles and SecureFiles
- 1.4. BasicFiles
- 1.5. SecureFiles
- 1.6. Creating a Table with a LOB Column



- 1.7. Creating a BasicFiles LOB Column
  - 1.8. Implementing a LOB in a Specific Tablespace
  - 1.9. Creating a SecureFiles LOB Column
  - 1.10. Implementing a Partitioned LOB
  - 1.11. Maintaining LOB Columns
  - 1.12. Moving a LOB Column
  - 1.13. Adding a LOB Column
  - 1.14. Removing a LOB Column
  - 1.15. Caching LOBs
  - 1.16. Storing LOBs In- and Out of Line
  - 1.17. Implementing SecureFiles Advanced Features
  - 1.18. Compressing LOBs
  - 1.19. Deduplicating LOBs
  - 1.20. Encrypting LOBs
  - 1.21. Migrating BasicFiles to SecureFiles
  - 1.22. Loading LOBs
  - 1.23. Loading a CLOB
  - 1.24. Loading a BLOB
  - 1.25. Measuring LOB Space Consumed
  - 1.26. BasicFiles Space Used
  - 1.27. SecureFiles Space Used
  - 1.28. Reading BFILEs
1. Chapter 12: Partitioning: Divide and Conquer
- 1.1. What Tables Should Be Partitioned?
  - 1.2. Creating Partitioned Tables
  - 1.3. Partitioning by Range
  - 1.4. Placing Partitions in Tablespaces
  - 1.5. Partitioning by List
  - 1.6. Partitioning by Hash
  - 1.7. Blending Different Partitioning Methods
  - 1.8. Creating Partitions on Demand
  - 1.9. Partitioning to Match a Parent Table
  - 1.10. Partitioning on a Virtual Column
  - 1.11. Giving an Application Control over Partitioning
  - 1.12. Maintaining Partitions
  - 1.13. Viewing Partition Metadata
  - 1.14. Moving a Partition
  - 1.15. Automatically Moving Updated Rows
  - 1.16. Partitioning an Existing Table
  - 1.17. Adding a Partition

- 1.18. Exchanging a Partition with an Existing Table
- 1.19. Renaming a Partition
- 1.20. Splitting a Partition
- 1.21. Merging Partitions
- 1.22. Dropping a Partition
- 1.23. Generating Statistics for a Partition
- 1.24. Removing Rows from a Partition
- 1.25. Manipulating Data Within a Partition
- 1.26. Partitioning Indexes
- 1.27. Partitioning an Index to Follow Its Table
- 1.28. Partitioning an Index Differently from Its Table
- 1.29. Partial Indexes
- 1.30. Partition Pruning
- 1. Chapter 13: Data Pump
  - 1.1. Data Pump Architecture
  - 1.2. Getting Started
  - 1.3. Taking an Export
  - 1.4. Importing a Table
  - 1.5. Using a Parameter File
  - 1.6. Exporting and Importing with Granularity
  - 1.7. Exporting and Importing an Entire Database
  - 1.8. Schema Level
  - 1.9. Table Level
  - 1.10. Tablespace Level
  - 1.11. Transferring Data
  - 1.12. Exporting and Importing Directly Across the Network
  - 1.13. Copying Data Files
  - 1.14. Features for Manipulating Storage
  - 1.15. Exporting Tablespace Metadata
  - 1.16. Specifying Different Data File Paths and Names
  - 1.17. Importing into a Tablespace Different from the Original
  - 1.18. Changing the Size of Data Files
  - 1.19. Changing Segment and Storage Attributes
  - 1.20. Filtering Data and Objects
  - 1.21. Specifying a Query
  - 1.22. Exporting a Percentage of the Data
  - 1.23. Excluding Objects from the Export File
  - 1.24. Excluding Statistics
  - 1.25. Including Only Specific Objects in an Export File



- |       |  |       |   |
|-------|--|-------|---|
| 1.26. | Exporting Table, Index, Constraint, and Trigger DDL  | 1.45. | Disabling Logging of Redo on Import             |
| 1.27. | Excluding Objects from Import                        | 1.46. | Interactive Command Mode                        |
| 1.28. | Including Objects in Import                          | 1.47. | Entering Interactive Command Mode               |
| 1.29. | Common Data Pump Tasks                               | 1.48. | Attaching to a Running Job                      |
| 1.30. | Estimating the Size of Export Jobs                   | 1.49. | Stopping and Restarting a Job                   |
| 1.31. | Listing the Contents of Dump Files                   | 1.50. | Terminating a Data Pump Job                     |
| 1.32. | Cloning a User                                       | 1.51. | Monitoring Data Pump Jobs                       |
| 1.33. | Creating a Consistent Export                         | 1.52. | Data Pump Log File                              |
| 1.34. | Importing When Objects Already Exist                 | 1.53. | Data Dictionary Views                           |
| 1.35. | Renaming a Table                                     | 1.54. | Database Alert Log                              |
| 1.36. | Remapping Data Suppressing a Log File                | 1.55. | Status Table                                    |
| 1.37. | Using Parallelism                                    | 1.56. | Interactive Command Mode Status                 |
| 1.38. | Specifying Additional Dump Files                     | 1.57. | OS Utilities                                    |
| 1.39. | Reusing Output File Names                            | 1.58. | Data Pump Legacy Mode                           |
| 1.40. | Creating a Daily DDL File                            | 1.59. | Data Pump Mapping to the exp Utility            |
| 1.41. | Compressing Output                                   | 1.60. | Data Pump Mapping to the imp Utility            |
| 1.42. | Changing Table Compression Characteristics on Import | 1.    | Chapter 14: External Tables                     |
| 1.43. | Encrypting Data                                      | 1.1.  | SQL*Loader vs. External Tables                  |
| 1.44. | Exporting Views As Tables                            | 1.2.  | Loading CSV Files into the Database             |
|       |  | 1.3.  | Creating a Directory Object and Granting Access |

[www.cursoslinux.com.mx](http://www.cursoslinux.com.mx)

[ventas@plct.com.mx](mailto:ventas@plct.com.mx)

PLCT S.A. de C.V.

Tel.: 55 4522 7839/55 1800 7696/7224447684

- 1.4. Creating an External Table
- 1.5. Generating SQL to Create an External Table
- 1.6. Viewing External Table Metadata
- 1.7. Loading a Regular Table from the External Table
- 1.8. Performing Advanced Transformations
- 1.9. Viewing Text Files from SQL
- 1.10. Unloading and Loading Data Using an External Table
- 1.11. Enabling Parallelism to Reduce Elapsed Time
- 1.12. Compressing a Dump File
- 1.13. Encrypting a Dump File

#### Chapter 15: Materialized Views

- 1.1. Understanding MVs
- 1.2. MV Terminology
- 1.3. Referencing Useful Views
- 1.4. Creating Basic Materialized Views
- 1.5. Creating a Complete Refreshable MV
- 1.6. Creating a Fast Refreshable MV
- 1.7. Going Beyond the Basics
- 1.8. Creating MVs and Specifying Tablespace for MVs and Indexes

- 1.9. Creating Indexes on MVs
- 1.10. Partitioning MVs
- 1.11. Compressing an MV
- 1.12. Encrypting MV Columns
- 1.13. Building an MV on a Prebuilt Table
- 1.14. Creating an Unpopulated MV
- 1.15. Creating an MV Refreshed on Commit
- 1.16. Creating a Never Refreshable MV
- 1.17. Creating MVs for Query Rewrite
- 1.18. Creating a Fast Refreshable MV Based on a Complex Query
- 1.19. Viewing MV DDL
- 1.20. Dropping an MV
- 1.21. Modifying MVs
- 1.22. Modifying Base Table DDL and Propagating to MVs
- 1.23. Toggling Redo Logging on an MV
- 1.24. Altering Parallelism
- 1.25. Moving an MV
- 1.26. Managing MV Logs
- 1.27. Creating an MV Log

- 1.28. Indexing MV Log Columns
- 1.29. Viewing Space Used by an MV Log
- 1.30. Shrinking the Space in an MV Log
- 1.31. Checking the Row Count of an MV Log
- 1.32. Moving an MV Log
- 1.33. Dropping an MV Log
- 1.34. Refreshing MVs
- 1.35. Manually Refreshing MVs from SQL\*Plus
- 1.36. Automating Refreshes, Using a Shell Script and Scheduling Utility
- 1.37. Creating an MV with a Refresh Interval
- 1.38. Efficiently Performing a Complete Refresh
- 1.39. Handling the ORA-12034 Error
- 1.40. Monitoring MV Refreshes
- 1.41. Viewing MVs' Last Refresh Times
- 1.42. Determining Whether a Refresh Is in Progress
- 1.43. Monitoring Real-Time Refresh Progress
- 1.44. Checking Whether MVs Are Refreshing Within a Time Period
- 1.45. Creating Remote MV Refreshes

- 1.46. Understanding Remote-Refresh Architectures
- 1.47. Viewing MV Base Table Information
- 1.48. Determining How Many MVs Reference a Central MV Log
- 1.49. Managing MVs in Groups
- 1.50. Creating an MV Group
- 1.51. Altering an MV Refresh Group
- 1.52. Refreshing an MV Group
- 1.53. DBMS\_MVIEW vs. DBMS\_REFRESH
- 1.54. Determining MVs in a Group
- 1.55. Adding an MV to a Refresh Group
- 1.56. Removing MVs from a Refresh Group
- 1.57. Dropping an MV Refresh Group

#### Chapter 16: User-Managed Backup and Recovery

- 1.1. Implementing a Cold-Backup Strategy for a Noarchivelog Mode Database
- 1.2. Making a Cold Backup of a Noarchivelog Mode Database
- 1.3. Restoring a Cold Backup in Noarchivelog Mode with Online Redo Logs
- 1.4. Restoring a Cold Backup in Noarchivelog Mode Without Online Redo Logs

- 1.5. Scripting a Cold Backup and Restore
- 1.6. Making a Cold Backup of an Archivelog Mode Database
- 1.7. Implementing a Hot Backup Strategy
- 1.8. Making a Hot Backup
- 1.9. Scripting Hot Backups
- 1.10. Understanding the Split-Block Issue
- 1.11. Understanding the Need for Redo Generated During Backup
- 1.12. Understanding that Data Files are Updated
- 1.13. Performing a Complete Recovery of an Archivelog Mode Database
- 1.14. Restoring and Recovering with the Database Offline
- 1.15. Restoring and Recovering with a Database Online
- 1.16. Restoring Control Files
- 1.17. Performing an Incomplete Recovery of an Archivelog Mode Database

#### Chapter 17: Configuring RMAN

- 1.1. Understanding RMAN
- 1.2. Starting RMAN
- 1.3. RMAN Architectural Decisions
- 1.4. Running the RMAN Client Remotely or Locally

- 1.5. Specifying the Backup User
- 1.6. Using Online or Offline Backups
- 1.7. Setting the Archive Redo Log Destination and File Format
- 1.8. Configuring the RMAN Backup Location and File Format
- 1.9. Setting the Autobackup of the Control File
- 1.10. Specifying the Location of the Autobackup of the Control File
- 1.11. Backing Up Archive Redo Logs
- 1.12. Determining the Location for the Snapshot Control File
- 1.13. Using a Recovery Catalog
- 1.14. Using a Media Manager
- 1.15. Setting the CONTROL\_FILE\_RECORD\_KEEP\_TIME Initialization Parameter
- 1.16. Configuring RMAN's Backup Retention Policy
- 1.17. Configuring the Archive Redo Logs' Deletion Policy
- 1.18. Setting the Degree of Parallelism
- 1.19. Using Backup Sets or Image Copies
- 1.20. Using Incremental Backups
- 1.21. Using Incrementally Updated Backups
- 1.22. Using Block Change Tracking

- 1.23. Configuring Binary Compression
- 1.24. Configuring Encryption
- 1.25. Configuring Miscellaneous Settings
- 1.26. Configuring Informational Output
- 1.27. Segueing from Decision to Action
- 1. Chapter 18: RMAN Backups and Reporting
  - 1.1. Preparing to Run RMAN Backup Commands
  - 1.2. Setting NLS\_DATE\_FORMAT
  - 1.3. Setting ECHO Setting ECHO
  - 1.4. Showing Variables
  - 1.5. Running Backups
  - 1.6. Backing Up the Entire Database
  - 1.7. Backing Up Tablespaces
  - 1.8. Backing Up Data Files
  - 1.9. Backing Up the Control File
  - 1.10. Backing up the spfile
  - 1.11. Backing Up Archive Redo Logs
  - 1.12. Backing Up FRA
  - 1.13. Excluding Tablespaces from Backups
  - 1.14. Backing Up Data Files Not Backed Up
  - 1.15. Skipping Read-Only Tablespaces
  - 1.16. Skipping Offline or Inaccessible Files
  - 1.17. Backing Up Large Files in Parallel
  - 1.18. Adding RMAN Backup Information to the Repository
  - 1.19. Taking Backups of Pluggable Databases
    - 1.20. While Connected to the Root Container
    - 1.21. While Connected to a Pluggable Database
  - 1.22. Creating Incremental Backups
    - 1.23. Taking Incremental-Level Backups
    - 1.24. Making Incrementally Updating Backups
    - 1.25. Using Block Change Tracking
    - 1.26. Checking for Corruption in Data Files and Backups
    - 1.27. Using VALIDATE
      - 1.28. Using BACKUP...VALIDATE
      - 1.29. Using RESTORE...VALIDATE
    - 1.30. Using a Recovery Catalog
      - 1.31. Creating a Recovery Catalog
      - 1.32. Registering a Target Database
      - 1.33. Backing Up the Recovery Catalog
      - 1.34. Synchronizing the Recovery Catalog

[www.cursoslinux.com.mx](http://www.cursoslinux.com.mx)

[ventas@plct.com.mx](mailto:ventas@plct.com.mx)

PLCT S.A. de C.V.

Tel.: 55 4522 7839/55 1800 7696/7224447684

- 1.35. Recovery Catalog Versions
- 1.36. Dropping a Recovery Catalog
- 1.37. Logging RMAN Output
- 1.38. Redirecting Output to a File
- 1.39. Capturing Output with Linux/Unix Logging Commands
- 1.40. Logging Output to a File
- 1.41. Querying for Output in the Data Dictionary
- 1.42. RMAN Reporting
- 1.43. Using LIST
- 1.44. Using REPORT
- 1.45. Using SQL

#### Chapter 19: RMAN Restore and Recovery

- 1.1. Determining if Media Recovery Is Required
- 1.2. Determining What to Restore
- 1.3. How the Process Works
- 1.4. Using Data Recovery Advisor
- 1.5. Using RMAN to Stop/Start Oracle
- 1.6. Shutting Down
- 1.7. Starting Up

- 1.8. Complete Recovery
- 1.9. Testing Restore and Recovery
- 1.10. Restoring and Recovering the Entire Database
- 1.11. Restoring and Recovering Tablespaces
- 1.12. Restoring Read-Only Tablespaces
- 1.13. Restoring Temporary Tablespaces
- 1.14. Restoring and Recovering Data Files
- 1.15. Restoring Data Files to Nondefault Locations
- 1.16. Performing Block-Level Recovery
- 1.17. Restoring a Container Database and Its Associated Pluggable Databases
- 1.18. Restoring Archive Redo Log Files
- 1.19. Restoring to the Default Location
- 1.20. Restoring to a Nondefault Location
- 1.21. Restoring a Control File
- 1.22. Using a Recovery Catalog
- 1.23. Using an Autobackup
- 1.24. Specifying a Backup File Name
- 1.25. Restoring the spfile
- 1.26. Incomplete Recovery



- 1.27. Determining the Type of Incomplete Recovery
- 1.28. Performing Time-Based Recovery
- 1.29. Performing Log Sequence-Based Recovery
- 1.30. Performing SCN-Based Recovery
- 1.31. Restoring to a Restore Point
- 1.32. Restoring Tables to a Previous Point
- 1.33. Flashing Back a Table
- 1.34. FLASHBACK TABLE TO BEFORE DROP
- 1.35. Flashing Back a Table to a Previous Point in Time
- 1.36. Flashing Back a Database
- 1.37. Restoring and Recovering to Different Server

Step 1. Create an RMAN Backup on the Originating Database

Step 2. Copy the RMAN Backup to the Destination Server

Step 3. Ensure That Oracle Is Installed

Step 4. Source the Required OS Variables

Step 5. Create an init.ora File for the Database to Be Restored

Step 6. Create Any Required Directories for Data Files, Control Files, and Dump/Trace Files Step

7. Start Up the Database in Nomount Mode

Step 8. Restore the Control File from the RMAN Backup

Step 9. Start Up the Database in Mount Mode Step

10. Make the Control File Aware of the Location of the RMAN Backups

Step 11. Rename and Restore the Data Files to Reflect New Directory Locations

Step 12. Recover the Database

Step 13. Set the New Location for the Online Redo Logs

Step 14. Open the Database

Step 15. Add the Temp File

Step 16. Rename the Database

Chapter 20: Oracle Secure Backup

1.1. OSB Editions and Features

1.2. OSB Terminology

1.3. OSB Administrative Domain and Servers

1.4. OSB Interfaces

1.5. OSB Users and Classes

1.6. OSB Daemons

1.7. Download and Installation Command-Line Access to OSB

1.8. OSB Configuration

1.9. Configuring Users and Classes

- 1.10. Configuring Media Families
  - 1.11. Configuring Database Backup Storage Selector
  - 1.12. Database Backup
  - 1.13. Database Restore
  - 1.14. Filesystem Backup
  - 1.15. Creating Data Set Files
  - 1.16. Configuring Backup Windows
  - 1.17. Configuring Backup Schedules and Triggers
  - 1.18. Performing On-Demand Filesystem Backups
  - 1.19. Filesystem Restore
  - 1.20. Performing Catalog-Based Restore
  - 1.21. Performing a Raw Restore
  - 1.22. Performing an obtar Restore
  - 1.23. Restoring a Container Database and Its Associated Pluggable Databases
  - 1.24. Restoring Archive Redo Log Files
  - 1.25. Restoring to the Default Location
  - 1.26. Restoring to a Nondefault Location
  - 1.27. Restoring a Control File
  - 1.28. Using a Recovery Catalog
  - 1.29. Using an Autobackup
  - 1.30. Specifying a Backup File Name
  - 1.31. Restoring the spfile
  - 1.32. Incomplete Recovery
  - 1.33. Determining the Type of Incomplete Recovery
  - 1.34. Performing Time-Based Recovery
  - 1.35. Performing Log Sequence-Based Recovery
  - 1.36. Performing SCN-Based Recovery
  - 1.37. Restoring to a Restore Point
  - 1.38. Restoring Tables to a Previous Point
  - 1.39. Flashing Back a Table
  - 1.40. FLASHBACK TABLE TO BEFORE DROP
  - 1.41. Flashing Back a Table to a Previous Point in Time
  - 1.42. Flashing Back a Database
  - 1.43. Restoring and Recovering to Different Server
- Step 1. Create an RMAN Backup on the Originating Database
- Step 2. Copy the RMAN Backup to the Destination Server
- Step 3. Ensure That Oracle Is Installed
- Step 4. Source the Required OS Variables
- Step 5. Create an init.ora File for the Database to Be Restored

[www.cursoslinux.com.mx](http://www.cursoslinux.com.mx)

[ventas@plct.com.mx](mailto:ventas@plct.com.mx)

PLCT S.A. de C.V.

Tel.: 55 4522 7839/55 1800 7696/7224447684

Step 6. Create Any Required Directories for Data Files, Control Files, and Dump/Trace Files

Step 7. Start Up the Database in Nomount Mode

Step 8. Restore the Control File from the RMAN Backup

Step 9. Start Up the Database in Mount Mode

Step 10. Make the Control File Aware of the Location of the RMAN Backups

Step 11. Rename and Restore the Data Files to Reflect New Directory Locations

Step 12. Recover the Database

Step 13. Set the New Location for the Online Redo Logs

Step 14. Open the Database

Step 15. Add the Temp File

Step 16. Rename the Database

Chapter 20: Oracle Secure Backup

1.1. OSB Editions and Features

1.2. OSB Terminology

1.3. OSB Administrative Domain and Servers

1.4. OSB Interfaces

1.5. OSB Users and Classes

1.6. OSB Daemons

1.7. Download and Installation

1.8. Command-Line Access to OSB

1.9. OSB Configuration

1.10. Configuring Users and Classes

1.11. Configuring Media Families

1.12. Configuring Database Backup Storage Selector

1.13. Database Backup

1.14. Database Restore

1.15. Filesystem Backup

1.16. Creating Data Set Files

1.17. Configuring Backup Windows

1.18. Configuring Backup Schedules and Triggers

1.19. Performing On-Demand Filesystem Backups

1.20. Filesystem Restore

1.21. Performing Catalog-Based Restore

1.22. Performing a Raw Restore

1.23. Performing an obtar Restore

1.24. OSB Job Monitoring

1.25. Listing Jobs

1.26. Showing Job Transcripts

[www.cursoslinux.com.mx](http://www.cursoslinux.com.mx)

[ventas@plct.com.mx](mailto:ventas@plct.com.mx)

PLCT S.A. de C.V.

Tel.: 55 4522 7839/55 1800 7696/7224447684

- 1.27. Monitoring OSB Logs
- 1.28. Virtual Test Devices
- 1.29. Oracle Database Backup in the Cloud
- 1.30. OSB Software Upgrades

#### Chapter 21: Automating Jobs

- 1.1. Automating Jobs with Oracle Scheduler
- 1.2. Creating and Scheduling a Job
- 1.3. Viewing Job Details
- 1.4. Modifying Job Logging History
- 1.5. Modifying a Job
- 1.6. Stopping a Job
- 1.7. Disabling a Job
- 1.8. Enabling a Job
- 1.9. Copying a Job
- 1.10. Running a Job Manually
- 1.11. Deleting a Job
- 1.12. Oracle Scheduler vs. cron
- 1.13. Automating Jobs via cron
- 1.14. How cron Works
- 1.15. Enabling Access to cron

- 1.16. Understanding cron Table Entries
- 1.17. Scheduling a Job to Run Automatically
- 1.18. Redirecting cron Output
- 1.19. Troubleshooting cron
- 1.20. Examples of Automated DBA Jobs
- 1.21. Starting and Stopping the Database and Listener
- 1.22. Checking for Archive Redo Destination Fullness
- 1.23. Truncating Large Log Files
- 1.24. Checking for Locked Production Accounts
- 1.25. Checking for Files Older Than a Certain Age
- 1.26. Checking for Too Many Processes
- 1.27. Verifying the Integrity of RMAN Backups

#### Chapter 22: Database Troubleshooting

- 1.1. Quickly Triaging
- 1.2. Checking Database Availability
- 1.3. Investigating Disk Fullness
- 1.4. Inspecting the Alert Log
- 1.5. Identifying Bottlenecks via OS Utilities
- 1.6. Identifying System Bottlenecks
- 1.7. Mapping an Operating System Process to an SQL Statement

[www.cursoslinux.com.mx](http://www.cursoslinux.com.mx)

[ventas@plct.com.mx](mailto:ventas@plct.com.mx)

PLCT S.A. de C.V.

Tel.: 55 4522 7839/55 1800 7696/7224447684

- 1.8. Finding Resource-Intensive SQL Statements
- 1.9. Monitoring Real-Time SQL Execution Statistics
- 1.10. Running Oracle Diagnostic Utilities
- 1.11. Detecting and Resolving Locking Issues
- 1.12. Resolving Open-Cursor Issues
- 1.13. Troubleshooting Undo Tablespace Issues
- 1.14. Determining if Undo Is Correctly Sized
- 1.15. Viewing SQL That Is Consuming Undo Space
- 1.16. Handling Temporary Tablespace Issues
- 1.17. Determining if Temporary Tablespace Is Sized Correctly
- 1.18. Viewing SQL That Is Consuming Temporary Space

Chapter 23: Pluggable Databases

- 1.1. Understanding Pluggable Architecture
- 1.2. Paradigm Shift
- 1.3. B&R Implications
- 1.4. Tuning Nuances
- 1.5. Creating a CDB
- 1.6. Creating Manually with SQL
- 1.7. Using the DBCA

- 1.8. Generating CDB Create Scripts via DBCA
- 1.9. Verifying that a CDB was Created
- 1.10. Administrating the Root Container
- 1.11. Connecting to the Root Container
- 1.12. Displaying Currently Connected Container Information
- 1.13. Starting/Stopping the Root Container
- 1.14. Creating Common Users
- 1.15. Creating Common Roles
- 1.16. Reporting on Container Space
- 1.17. Switching Containers
- 1.18. Creating a Pluggable Database within a CDB
- 1.19. Cloning the Seed Database
- 1.20. Cloning an Existing Pluggable Database
- 1.21. Cloning from a Non-CDB Database
- 1.22. Unplugging a Pluggable Database from a CDB
- 1.23. Plugging an Unplugged Pluggable Database into a CDB
- 1.24. Using the DBCA to Create a Pluggable Database from the Seed Database
- 1.25. Checking the Status of Pluggable Databases

[www.cursoslinux.com.mx](http://www.cursoslinux.com.mx)

[ventas@plct.com.mx](mailto:ventas@plct.com.mx)

PLCT S.A. de C.V.

Tel.: 55 4522 7839/55 1800 7696/7224447684

- 1.26. Administrating Pluggable Databases
- 1.27. Connecting to a Pluggable Database
- 1.28. Managing a Listener in a Pluggable Database Environment
- 1.29. Showing the Currently Connected Pluggable Database
- 1.30. Starting/Stopping a Pluggable Database
- 1.31. Modifying Initialization Parameters Specific to a Pluggable Database
- 1.32. Renaming a Pluggable Database
- 1.33. Limiting the Amount of Space Consumed by a Pluggable Database
- 1.34. Viewing Pluggable Database History
- 1.35. Dropping a Pluggable Database