

# ADMINISTRACION ORACLE

## INFORMACION

### FORMATO

Presencial

En Sitio

A partir de 3 participantes

### DURACIÓN:

80 Horas

10 días

Lunes a Viernes

### DIRIGIDO A

Usuarios de servidor de Bases de Datos Oracle que requieren realizar tareas administrativas

### REQUISITOS:

xxxxxxxxxxxxxxxxxxx

### MATERIAL:

Manual Oficial

DVD de la distribución

Linux más reciente

### DOCUMENTO

Diploma expedido por

PLCT S.A. DE C.V.

## DESCRIPCION 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.

## OBJETIVOS

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



## Chapter 1: Installing the Oracle Binaries

Understanding the OFA

Oracle Inventory Directory

Oracle Base Directory

Oracle Home Directory

Oracle Network Files Directory

Automatic Diagnostic Repository

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

Step 7. Troubleshoot Any Issues

Installing with a Copy of an Existing Installation

Step 1. Copy the Binaries, Using an OS Utility

Step 2. Attach the Oracle Home

Upgrading Oracle Software

Reinstalling After Failed Installation

Applying Interim Patches

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

A Manually Intensive Approach

Oracle's Approach to Setting OS Variables

My Approach to Setting OS Variables

Creating a Database

Step 1. Set the OS Variables

Step 2: Configure the Initialization File

Step 3: Create the Required Directories

Step 4: Create the Database

Step 5. Create a Data Dictionary

Configuring and Implementing the Listener

Manually Configuring a Listener

Implementing a Listener with the Net Configuration Assistant

Connecting to a Database through the Network

Creating a Password File

Starting and Stopping the Database

Understanding OS Authentication

Starting the Database

Stopping the Database

Using the dbca to Create a Database

Dropping a Database

How Many Databases on One Server?

Understanding Oracle Architecture

## Chapter 3: Configuring an Efficient Environment

Customizing Your OS Command Prompt

Customizing Your SQL Prompt

Creating Shortcuts for Frequently Used Commands

Using Aliases

Using a Function

Rerunning Commands Quickly

Rerunning Commands Quickly Scrolling with the Up and Down Arrow Keys	Altering Tablespace Size
Using Ctrl+P and Ctrl+N	toggling Data Files Offline and Online
Listing the Command History	Renaming or Relocating a Data File
Searching in Reverse	Performing Online Data File Operations
Setting the Command Editor	Performing Offline Data File Operations
Developing Standard Scripts	
dba_setup	Chapter 5: Managing Control Files, Online Redo Logs, and Archiving
dba_fcns	Managing Control Files
tbsp_chk.bsh	Viewing Control File Names and Locations
conn.bsh	Adding a Control File
filesp.bsh	Moving a Control File
login.sql	Removing a Control File
top.sql	Managing Online Redo Logs
lock.sql	Displaying Online Redo Log Information
users.sql	Determining the Optimal Size of Online Redo Log Groups
Organizing Scripts	Determining the Optimal Number of Redo Log Groups
Step 1: Create Directories	Adding Online Redo Log Groups
Step 2: Copy Files to Directories	Resizing and Dropping Online Redo Log Groups
Step 3: Configure the Startup File	Adding Online Redo Log Files to a Group
	Removing Online Redo Log Files from a Group
Chapter 4: Tablespaces and Data Files	Moving or Renaming Redo Log Files
Understanding the First Five	Implementing Archivelog Mode
Understanding the Need for More	Making Architectural Decisions
Creating Tablespaces	Setting the Archive Redo File Location
Renaming a Tablespace	Thinking “Un-Oraclethodox” FRA Thoughts
Controlling the Generation of Redo	Enabling Archivelog Mode
Changing a Tablespace’s Write Mode	Disabling Archivelog Mode
Dropping a Tablespace	Reacting to a Lack of Disk Space in Your Archive Log Destination
Using Oracle Managed Files	Backing Up Archive Redo Log Files
Creating a Bigfile Tablespace	
Enabling Default Table Compression within a Tablespace	
Displaying Tablespace Size	

## Chapter 6: Users and Basic Security

Managing Default Users  
Locking Accounts and Expiring Passwords  
Identifying DBA-Created Accounts  
Checking Default Passwords  
Creating Users  
Choosing a Username and Authentication Method  
Assigning Default Permanent and Temporary Tablespace  
Modifying Passwords  
Logging In as a Different User  
Modifying Users  
Dropping Users  
Enforcing Password Security and Resource Limits  
Basic Password Security  
Password Strength  
Limiting Database Resource Usage  
Managing Privileges  
Assigning Database System Privileges  
Assigning Database Object Privileges  
Grouping and Assigning Privileges

## Chapter 7: Tables and Constraints

Understanding Table Types  
Understanding Data Types  
Character  
Numeric  
Date/Time  
RAW  
ROWID  
LOB  
Extended Character Types  
Creating a Table  
Creating a Heap-Organized Table

Implementing Virtual Columns  
Implementing Invisible Columns  
Making Read-Only Tables  
Understanding Deferred Segment Creation  
Creating a Table with an Autoincrementing (Identity) Column  
Allowing for Default Parallel SQL Execution  
Compressing Table Data  
Avoiding Redo Creation  
Creating a Table from a Query  
Modifying a Table  
Obtaining the Needed Lock  
Renaming a Table  
Adding a Column  
Altering a Column  
Renaming a Column  
Dropping a Column  
Displaying Table DDL  
Dropping a Table  
Undropping a Table  
Removing Data from a Table  
Using DELETE  
Using TRUNCATE  
Viewing and Adjusting the High-Water Mark  
You need to be aware of a couple of performance-related issues regarding the high-water mark  
Tracing to Detect Space Below the High-Water Mark  
Using DBMS\_SPACE to Detect Space Below the High-Water Mark  
Selecting from Data Dictionary Extents View  
Lowering the High-Water Mark  
Creating a Temporary Table  
Creating an Index-Organized Table  
Managing Constraints

Creating Primary Key Constraints	Making Indexes Unusable
Enforcing Unique Key Values	Monitoring Index Usage
Creating Foreign Key Constraints	Dropping an Index
Checking for Specific Data Conditions	Indexing Foreign Key Columns
Enforcing Not Null Conditions	Implementing an Index on a Foreign Key Column
Disabling Constraints	Determining if Foreign Key Columns Are Indexed
Enabling Constraints	Chapter 9: Views, Synonyms, and Sequences
Chapter 8: Indexes	Implementing Views
Deciding When to Create an Index	Creating a View
Proactively Creating Indexes	Checking Updates
Reactively Creating Indexes	Creating Read-Only Views
Planning for Robustness	Updatable Join Views
Determining Which Type of Index to Use	Creating an INSTEAD OF Trigger
Estimating the Size of an Index Before Creation	Implementing an Invisible Column
Creating Separate Tablespaces for Indexes	Modifying a View Definition
Creating Portable Scripts	Displaying the SQL Used to Create a View
Establishing Naming Standards	Renaming a View
Creating Indexes	Dropping a View
Creating B-tree Indexes	Managing Synonyms
Creating Concatenated Indexes	Creating a Synonym
Implementing Function-Based Indexes	Creating Public Synonyms
Creating Unique Indexes	Dynamically Generating Synonyms
Implementing Bitmap Indexes	Displaying Synonym Metadata
Creating Bitmap Join Indexes	Renaming a Synonym
Implementing Reverse-Key Indexes	Dropping a Synonym
Creating Key-Compressed Indexes	Managing Sequences
Parallelizing Index Creation	Creating a Sequence
Avoiding Redo Generation When Creating an Index	Using Sequence Pseudocolumns
Implementing Invisible Indexes	Autoincrementing Columns
Maintaining Indexes	Implementing Multiple Sequences That Generate Unique Values
Renaming an Index	Creating One Sequence or Many
Displaying Code to Recreate an Index	Viewing Sequence Metadata
Rebuilding an Index	Renaming a Sequence

## Chapter 10: Data Dictionary Fundamentals

Data Dictionary Architecture

Static Views

Dynamic Performance Views

A Different View of Metadata

A Few Creative Uses of the Data Dictionary

Derivable Documentation

Displaying User Information

Displaying Table Row Counts

Showing Primary Key and Foreign Key Relationships

Displaying Object Dependencies

## Chapter 11: Large Objects

Describing LOB Types

Illustrating LOB Locators, Indexes, and Chunks

Distinguishing Between BasicFiles and SecureFiles

BasicFiles

SecureFiles

Creating a Table with a LOB Column

Creating a BasicFiles LOB Column

Implementing a LOB in a Specific Tablespace

Creating a SecureFiles LOB Column

Implementing a Partitioned LOB

Maintaining LOB Columns

Moving a LOB Column

Adding a LOB Column

Removing a LOB Column

Caching LOBs

Storing LOBs In- and Out of Line

Implementing SecureFiles Advanced Features

Compressing LOBs

Deduplicating LOBs

Encrypting LOBs

Migrating BasicFiles to SecureFiles

Loading LOBs

Loading a CLOB

Loading a BLOB

Measuring LOB Space Consumed

BasicFiles Space Used

SecureFiles Space Used

Reading BFILEs

## Chapter 12: Partitioning: Divide and Conquer

What Tables Should Be Partitioned?

Creating Partitioned Tables

Partitioning by Range

Placing Partitions in Tablespaces

Partitioning by List

Partitioning by Hash

Blending Different Partitioning Methods

Creating Partitions on Demand

Partitioning to Match a Parent Table

Partitioning on a Virtual Column

Giving an Application Control over Partitioning

Maintaining Partitions

Viewing Partition Metadata

Moving a Partition

Automatically Moving Updated Rows

Partitioning an Existing Table

Adding a Partition

Exchanging a Partition with an Existing Table

Renaming a Partition

Splitting a Partition

Merging Partitions

Dropping a Partition

Generating Statistics for a Partition

Removing Rows from a Partition

Manipulating Data Within a Partition

Partitioning Indexes	Exporting Table, Index, Constraint, and Trigger DDL
Partitioning an Index to Follow Its Table	Excluding Objects from Import
Partitioning an Index Differently from Its Table	Including Objects in Import
Partial Indexes	Common Data Pump Tasks
Partition Pruning	Estimating the Size of Export Jobs
	Listing the Contents of Dump Files
	Cloning a User
Chapter 13: Data Pump	Creating a Consistent Export
Data Pump Architecture	Importing When Objects Already Exist
Getting Started	Renaming a Table
Taking an Export	Remapping Data
Importing a Table	Suppressing a Log File
Using a Parameter File	Using Parallelism
Exporting and Importing with Granularity	Specifying Additional Dump Files
Exporting and Importing an Entire Database	Reusing Output File Names
Schema Level	Creating a Daily DDL File
Table Level	Compressing Output
Tablespace Level	Changing Table Compression Characteristics on Import
Transferring Data	Encrypting Data
Exporting and Importing Directly Across the Network	Exporting Views As Tables
Copying Data Files	Disabling Logging of Redo on Import
Features for Manipulating Storage	Interactive Command Mode
Exporting Tablespace Metadata	Entering Interactive Command Mode
Specifying Different Data File Paths and Names	Attaching to a Running Job
Importing into a Tablespace Different from the Original	Stopping and Restarting a Job
Changing the Size of Data Files	Terminating a Data Pump Job
Changing Segment and Storage Attributes	Monitoring Data Pump Jobs
Filtering Data and Objects	Data Pump Log File
Specifying a Query	Data Dictionary Views
Exporting a Percentage of the Data	Database Alert Log
Excluding Objects from the Export File	Status Table
Excluding Statistics	Interactive Command Mode Status
Including Only Specific Objects in an Export File	

OS Utilities

Data Pump Legacy Mode

Data Pump Mapping to the exp Utility

Data Pump Mapping to the imp Utility

Chapter 14: External Tables

SQL\*Loader vs. External Tables

Loading CSV Files into the Database

Creating a Directory Object and Granting Access

Creating an External Table

Generating SQL to Create an External Table

Viewing External Table Metadata

Loading a Regular Table from the External Table

Performing Advanced Transformations

Viewing Text Files from SQL

Unloading and Loading Data Using an External Table

Enabling Parallelism to Reduce Elapsed Time

Compressing a Dump File

Encrypting a Dump File

Chapter 15: Materialized Views

Understanding MVs

MV Terminology

Referencing Useful Views

Creating Basic Materialized Views

Creating a Complete Refreshable MV

Creating a Fast Refreshable MV

Going Beyond the Basics

Creating MVs and Specifying Tablespace for MVs and Indexes

Creating Indexes on MVs

Partitioning MVs

Compressing an MV

Encrypting MV Columns

Building an MV on a Prebuilt Table

Creating an Unpopulated MV

Creating an MV Refreshed on Commit

Creating a Never Refreshable MV

Creating MVs for Query Rewrite

Creating a Fast Refreshable MV Based on a Complex Query

Viewing MV DDL

Dropping an MV

Modifying MVs

Modifying Base Table DDL and Propagating to MVs

Toggling Redo Logging on an MV

Altering Parallelism

Moving an MV

Managing MV Logs

Creating an MV Log

Indexing MV Log Columns

Viewing Space Used by an MV Log

Shrinking the Space in an MV Log

Checking the Row Count of an MV Log

Moving an MV Log

Dropping an MV Log

Refreshing MVs

Manually Refreshing MVs from SQL\*Plus

Automating Refreshes, Using a Shell Script and Scheduling Utility

Creating an MV with a Refresh Interval

Efficiently Performing a Complete Refresh

Handling the ORA-12034 Error

Monitoring MV Refreshes

Viewing MVs' Last Refresh Times

Determining Whether a Refresh Is in Progress

Monitoring Real-Time Refresh Progress



Checking Whether MVs Are Refreshing Within a Time Period	Restoring and Recovering with a Database Online
Creating Remote MV Refreshes	Restoring Control Files
Understanding Remote-Refresh Architectures	Performing an Incomplete Recovery of an Archivelog Mode Database
Viewing MV Base Table Information	
Determining How Many MVs Reference a Central MV Log	Chapter 17: Configuring RMAN
Managing MVs in Groups	Understanding RMAN
Creating an MV Group	Starting RMAN
Altering an MV Refresh Group	RMAN Architectural Decisions
Refreshing an MV Group	Running the RMAN Client Remotely or Locally
DBMS_MVIEW vs. DBMS_REFRESH	Specifying the Backup User
Determining MVs in a Group	Using Online or Offline Backups
Adding an MV to a Refresh Group	Setting the Archive Redo Log Destination and File Format
Removing MVs from a Refresh Group	Configuring the RMAN Backup Location and File Format
Dropping an MV Refresh Group	Setting the Autobackup of the Control File
Chapter 16: User-Managed Backup and Recovery	Specifying the Location of the Autobackup of the Control File
Implementing a Cold-Backup Strategy for a Noarchivelog Mode Database	Backing Up Archive Redo Logs
Making a Cold Backup of a Noarchivelog Mode Database	Determining the Location for the Snapshot Control File
Restoring a Cold Backup in Noarchivelog Mode with Online Redo Logs	Using a Recovery Catalog
Restoring a Cold Backup in Noarchivelog Mode Without Online Redo Logs	Using a Media Manager
Scripting a Cold Backup and Restore	Setting the CONTROL_FILE_RECORD_KEEP_TIME Initialization Parameter
Making a Cold Backup of an Archivelog Mode Database	Configuring RMAN's Backup Retention Policy
Implementing a Hot Backup Strategy	Configuring the Archive Redo Logs' Deletion Policy
Making a Hot Backup	Setting the Degree of Parallelism
Scripting Hot Backups	Using Backup Sets or Image Copies
Understanding the Split-Block Issue	Using Incremental Backups
Understanding the Need for Redo Generated During Backup	Using Incrementally Updated Backups
Understanding that Data Files are Updated	Using Block Change Tracking
Performing a Complete Recovery of an Archivelog Mode Database	Configuring Binary Compression
Restoring and Recovering with the Database Offline	Configuring Encryption
	Configuring Miscellaneous Settings

Configuring Informational Output

Segueing from Decision to Action

Chapter 18: RMAN Backups and Reporting

Preparing to Run RMAN Backup Commands

Setting NLS\_DATE\_FORMAT

Setting ECHO Setting ECHO

Showing Variables

Running Backups

Backing Up the Entire Database

Backing Up Tablespaces

Backing Up Data Files

Backing Up the Control File

Backing up the spfile

Backing Up Archive Redo Logs

Backing Up FRA

Excluding Tablespaces from Backups

Backing Up Data Files Not Backed Up

Skipping Read-Only Tablespaces

Skipping Offline or Inaccessible Files

Backing Up Large Files in Parallel

Adding RMAN Backup Information to the Repository

Taking Backups of Pluggable Databases

While Connected to the Root Container

While Connected to a Pluggable Database

Creating Incremental Backups

Taking Incremental-Level Backups

Making Incrementally Updating Backups

Using Block Change Tracking

Checking for Corruption in Data Files and Backups

Using VALIDATE

Using BACKUP...VALIDATE

Using RESTORE...VALIDATE

Using a Recovery Catalog

Creating a Recovery Catalog

Registering a Target Database

Backing Up the Recovery Catalog

Synchronizing the Recovery Catalog

Recovery Catalog Versions

Dropping a Recovery Catalog

Logging RMAN Output

Redirecting Output to a File

Capturing Output with Linux/Unix Logging Commands

Logging Output to a File

Querying for Output in the Data Dictionary

RMAN Reporting

Using LIST

Using REPORT

Using SQL

Chapter 19: RMAN Restore and Recovery

Determining if Media Recovery Is Required

Determining What to Restore

How the Process Works

Using Data Recovery Advisor

Using RMAN to Stop/Start Oracle

Shutting Down

Starting Up

Complete Recovery

Testing Restore and Recovery

Restoring and Recovering the Entire Database

Restoring and Recovering Tablespaces

Restoring Read-Only Tablespaces

Restoring Temporary Tablespaces

Restoring and Recovering Data Files

Restoring Data Files to Nondefault Locations

Performing Block-Level Recovery

Restoring a Container Database and Its Associated Pluggable Databases	Step 10. Make the Control File Aware of the Location of the RMAN Backups
Restoring Archive Redo Log Files	Step 11. Rename and Restore the Data Files to Reflect New Directory Locations
Restoring to the Default Location	Step 12. Recover the Database
Restoring to a Nondefault Location	Step 13. Set the New Location for the Online Redo Logs
Restoring a Control File	Step 14. Open the Database
Using a Recovery Catalog	Step 15. Add the Temp File
Using an Autobackup	Step 16. Rename the Database
Specifying a Backup File Name	
Restoring the spfile	
Incomplete Recovery	Chapter 20: Oracle Secure Backup
Determining the Type of Incomplete Recovery	OSB Editions and Features
Performing Time-Based Recovery	OSB Terminology
Performing Log Sequence-Based Recovery	OSB Administrative Domain and Servers
Performing SCN-Based Recovery	OSB Interfaces
Restoring to a Restore Point	OSB Users and Classes
Restoring Tables to a Previous Point	OSB Daemons
Flashing Back a Table	Download and Installation
FLASHBACK TABLE TO BEFORE DROP	Command-Line Access to OSB
Flashing Back a Table to a Previous Point in Time	OSB Configuration
Flashing Back a Database	Configuring Users and Classes
Restoring and Recovering to Different Server	Configuring Media Families
Step 1. Create an RMAN Backup on the Originating Database	Configuring Database Backup Storage Selector
Step 2. Copy the RMAN Backup to the Destination Server	Database Backup
Step 3. Ensure That Oracle Is Installed	Database Restore
Step 4. Source the Required OS Variables	Filesystem Backup
Step 5. Create an init.ora File for the Database to Be Restored	Creating Data Set Files
Step 6. Create Any Required Directories for Data Files, Control Files, and Dump/Trace Files	Configuring Backup Windows
Step 7. Start Up the Database in Nomount Mode	Configuring Backup Schedules and Triggers
Step 8. Restore the Control File from the RMAN Backup	Performing On-Demand Filesystem Backups
Step 9. Start Up the Database in Mount Mode	Filesystem Restore
	Performing Catalog-Based Restore
	Performing a Raw Restore
	Performing an obtar Restore

Restoring a Container Database and Its Associated Pluggable Databases	Step 10. Make the Control File Aware of the Location of the RMAN Backups
Restoring Archive Redo Log Files	Step 11. Rename and Restore the Data Files to Reflect New Directory Locations
Restoring to the Default Location	Step 12. Recover the Database
Restoring to a Nondefault Location	Step 13. Set the New Location for the Online Redo Logs
Restoring a Control File	Step 14. Open the Database
Using a Recovery Catalog	Step 15. Add the Temp File
Using an Autobackup	Step 16. Rename the Database
Specifying a Backup File Name	
Restoring the spfile	
Incomplete Recovery	Chapter 20: Oracle Secure Backup
Determining the Type of Incomplete Recovery	OSB Editions and Features
Performing Time-Based Recovery	OSB Terminology
Performing Log Sequence-Based Recovery	OSB Administrative Domain and Servers
Performing SCN-Based Recovery	OSB Interfaces
Restoring to a Restore Point	OSB Users and Classes
Restoring Tables to a Previous Point	OSB Daemons
Flashing Back a Table	Download and Installation
FLASHBACK TABLE TO BEFORE DROP	Command-Line Access to OSB
Flashing Back a Table to a Previous Point in Time	OSB Configuration
Flashing Back a Database	Configuring Users and Classes
Restoring and Recovering to Different Server	Configuring Media Families
Step 1. Create an RMAN Backup on the Originating Database	Configuring Database Backup Storage Selector
Step 2. Copy the RMAN Backup to the Destination Server	Database Backup
Step 3. Ensure That Oracle Is Installed	Database Restore
Step 4. Source the Required OS Variables	Filesystem Backup
Step 5. Create an init.ora File for the Database to Be Restored	Creating Data Set Files
Step 6. Create Any Required Directories for Data Files, Control Files, and Dump/Trace Files	Configuring Backup Windows
Step 7. Start Up the Database in Nomount Mode	Configuring Backup Schedules and Triggers
Step 8. Restore the Control File from the RMAN Backup	Performing On-Demand Filesystem Backups
Step 9. Start Up the Database in Mount Mode	Filesystem Restore
	Performing Catalog-Based Restore
	Performing a Raw Restore
	Performing an obtar Restore

OSB Job Monitoring	Checking for Too Many Processes
Listing Jobs	Verifying the Integrity of RMAN Backups
Showing Job Transcripts	
Monitoring OSB Logs	Chapter 22: Database Troubleshooting
Virtual Test Devices	Quickly Triaging
Oracle Database Backup in the Cloud	Checking Database Availability
OSB Software Upgrades	Investigating Disk Fullness
	Inspecting the Alert Log
Chapter 21: Automating Jobs	Identifying Bottlenecks via OS Utilities
Automating Jobs with Oracle Scheduler	Identifying System Bottlenecks
Creating and Scheduling a Job	Mapping an Operating System Process to an SQL Statement
Viewing Job Details	
Modifying Job Logging History	Finding Resource-Intensive SQL Statements
Modifying a Job	Monitoring Real-Time SQL Execution Statistics
Stopping a Job	Running Oracle Diagnostic Utilities
Disabling a Job	Detecting and Resolving Locking Issues
Enabling a Job	Resolving Open-Cursor Issues
Copying a Job	Troubleshooting Undo Tablespace Issues
Running a Job Manually	Determining if Undo Is Correctly Sized
Deleting a Job	Viewing SQL That Is Consuming Undo Space
Oracle Scheduler vs. cron	Handling Temporary Tablespace Issues
Automating Jobs via cron	Determining if Temporary Tablespace Is Sized Correctly
How cron Works	Viewing SQL That Is Consuming Temporary Space
Enabling Access to cron	
Understanding cron Table Entries	Chapter 23: Pluggable Databases
Scheduling a Job to Run Automatically	Understanding Pluggable Architecture
Redirecting cron Output	Paradigm Shift
Troubleshooting cron	B&R Implications
Examples of Automated DBA Jobs	Tuning Nuances
Starting and Stopping the Database and Listener	Creating a CDB
Checking for Archive Redo Destination Fullness	Creating Manually with SQL
Truncating Large Log Files	Using the DBCA
Checking for Locked Production Accounts	Generating CDB Create Scripts via DBCA
Checking for Files Older Than a Certain Age	Verifying that a CDB was Created

Administrating the Root Container

Connecting to the Root Container

Displaying Currently Connected Container Information

Starting/Stopping the Root Container

Creating Common Users

Creating Common Roles

Reporting on Container Space

Switching Containers

Creating a Pluggable Database within a CDB

Cloning the Seed Database

Cloning an Existing Pluggable Database

Cloning from a Non-CDB Database

Unplugging a Pluggable Database from a CDB

Plugging an Unplugged Pluggable Database into a CDB

Using the DBCA to Create a Pluggable Database from the Seed Database

Checking the Status of Pluggable Databases

Administrating Pluggable Databases

Connecting to a Pluggable Database

Managing a Listener in a Pluggable Database Environment

Showing the Currently Connected Pluggable Database

Starting/Stopping a Pluggable Database

Modifying Initialization Parameters Specific to a Pluggable Database

Renaming a Pluggable Database

Limiting the Amount of Space Consumed by a Pluggable Database

Viewing Pluggable Database History

Dropping a Pluggable Database