VEER NARMAD SOUTH GUJARAT UNIVERSITY, SURAT

M.C.A. (5 th Semester)
Proposed Syllabus
Paper : 503 : Distributed Databases
Effective From June 2008

1. **Distributed Database Concepts**
   - 1.1 Introduction & Recent Trends in Distributed Databases
   - 1.2 Distributed Database Architecture
   - 1.3 Database Links
   - 1.4 Distributed Database Administration
   - 1.5 Transaction Processing in a Distributed System
   - 1.6 Distributed Database Application Development
   - 1.7 Character Set Support for Distributed Environments
2. **Managing a Distributed Database**
   - 2.1 Managing Global Names in a Distributed System
   - 2.2 Creating Database Links
   - 2.3 Using Shared Database Links
   - 2.4 Managing Database Links
   - 2.5 Viewing Information About Database Links
   - 2.6 Creating Location Transparency
   - 2.7 Managing Statement Transparency
   - 2.8 Managing a Distributed Database: Examples
3. **Developing Applications for a Distributed Database**
   - 3.1 Managing the Distribution of Application Data
   - 3.2 Controlling Connections Established by Database Links
   - 3.3 Maintaining Referential Integrity in a Distributed System
   - 3.4 Tuning Distributed Queries
   - 3.5 Characterization of Query Processors
   - 3.6 Layers of Query Processing
   - 3.7 Handling Errors in Remote Procedures
4. **Replication**
   - 4.1 Synchronous Replication
   - 4.2 Asynchronous Replication
   - 4.3 Peer-to-Peer Replication
   - 4.4 Primary Site Replication
   - 4.5 Implementing the Capture & Apply steps
5. **Distributed Transaction Concepts**
   - 5.1 What Are Distributed Transactions?
   - 5.2 Transaction Properties
   - 5.3 Session Trees for Distributed Transactions
   - 5.4 Two-Phase Commit Mechanism
   - 5.5 In-Doubt Transactions
   - 5.6 Distributed Transaction Processing: Case Study
6. **Managing Distributed Transactions**
   - 6.1 Specifying the Commit Point Strength of a Node
   - 6.2 Naming Transactions

6.3 Types of Transactions - Flat, Nested & Workflows
6.4 Viewing Information About Distributed Transactions
6.5 Deciding How to Handle In-Doubt Transactions
6.6 Manually Overriding In-Doubt Transactions
6.7 Purging Pending Rows from the Data Dictionary
6.8 Manually Committing an In-Doubt Transaction: Example
6.9 Data Access Failures Due to Locks
6.10 Simulating Distributed Transaction Failure
6.11 Managing Read Consistency

7 Database Security

7.1 Site Autonomy
7.2 Distributed Database Security
- 7.2.1 Authentication through database links
- 7.2.2 Authentication without passwords
- 7.2.3 Supporting user accounts and roles
- 7.2.4 Centralized user and privilege management
- 7.2.5 Data encryption

7.3 Auditing Database Links
7.4 Administration Tools

References :

1 Oracle Database Administration Guide 10G Version 1
2 Principles of Distributed Database Systems : M. Tamer Ozsu, Patrick Valduriez, Pearson Education
3 Database Management Systems - 2nd Edition : Raghu Ramakrishnan, Johannes Gehrke, MGH
4 Managing Distributed Databases - Building Bridges between Database Islands: Donald K Burleson

VEER NARMAD SOUTH GUJARAT UNIVERSITY, SURAT

M.C.A. (5 th Semester)
Proposed Syllabus
Paper : 504 : Advanced Database Administration
Effective From June 2008

1. Oracle10g Instance creation and management

1.1 Oracle Instance
1.2 Installing Oracle
1.3 Oracle Optimal Flexible Architecture (OFA)
1.4 Locating initialization, listener.ora & sqlnet.ora files
1.5 Finding the alert log
1.6 Common environment variables
1.7 Structures in an Oracle Instance
1.8 Oracle Memory Structures, SGA and PGA
1.9 Oracle Processes and their purposes
1.10 Startup nomount, mount and open database commands

2. Oracle10g Database Architecture

2.1 Oraclde10g management framework
2.1 Using the Database Creation Assistant (DBA)
2.3 Creating and dropping a database
2.4 Tablespaces
2.5 Tables and Indexes
2.6 Clusters
2.7 Partitioning of Tables and Indexes
2.8 Gathering and applying patches

3. Concurrency Management

3.1 Transactions, serialization, locks and latches
3.2 Lock modes
3.3 Detecting and resolving lock conflicts
3.4 Managing deadlocks

4. Interfacing with Oracle

4.1 Oracle transaction management
4.2 Using SQL*Plus and iSQL*Plus
4.3 Using embedded Oracle with Pro*C & Java
4.4 PL/SQL & Triggers
4.5 Pining PL/SQL packages & compiling PL/SQL
4.6 System-level triggers – startup trigger, logon trigger, PL/SQL error trigger

5. Oracle*Net

5.1 Basic Network structure
5.2 Oracle*Net Files

5.3 Multi-threaded server
5.4 Create additional listeners
5.5 Create Oracle Net service aliases
5.6 Configure connect time failover
5.7 Oracle*Net names resolution

6. Tablespace Management Overview

6.1 Dictionary Managed Tablespaces
6.2 Locally Managed Tablespaces
6.3 Automatic Segment Space Management
6.4 Moving tablespaces online and offline

7. UNDO Tablespace Management

7.1 Use of undo segments
7.2 Creating an undo tablespace
7.3 User managed undo tablespaces
7.4 Automatic undo management
7.5 Monitor & Configure undo retention
7.6 Use the Undo Advisor
7.7 Size the undo tablespace

8. Oracle Utilities

8.1 Datapump - Import/export
8.2 SQL*Loader
8.3 Oracle Streams
8.4 Automatic Database Diagnostic Monitor
8.5 Automatic Tuning Optimizer
8.6 Automatic Shared Memory Tuning

9. Oracle Performance Tuning

9.1 Locate invalid and unusable objects
9.2 Gather SQL optimizer statistics with dbms_stats
9.3 Basic Oracle performance metrics
9.4 Use OEM and dbms_alert to set warning and critical alert thresholds
9.5 The SQL Tuning Advisor
9.6 The SQL Access Advisor
9.6 Interpreting server generated alerts
9.7 Oracle advisory utilities vdb_cache_advice, vshared_pool_advice v$pga_aggregate_target_advice
9.8 Using OEM performance screens
9.9 Fixing performance issues

10. User Management

10.1 Creating Users
10.2 Altering users
10.3 User Profiles
10.4 User resource groups
10.5 Granting privileges & roles
10.6 Auditing user activity with dbms_audit

11. Oracle Security

11.1 Password use in Oracle

Password encryption and password aging
External authentication
Using Single sign-on (SSO)

11.2 Object security
11.3 Virtual Private Databases (VPD) in Oracle
11.4 Oracle “grant execute” security
11.5 Use of Roles in Oracle
11.6 Register for security updates

12. Backup & Recovery

12.1 Oracle backup & recovery planning
12.2 Parallel instance recovery
12.3 Basics of checkpoints, redo log files, and archived log files
12.4 Using ARCHIVELOG mode
12.5 Creating consistent Oracle backups
12.6 Online hot backups
12.7 Incremental Oracle backups
12.8 Automating database backups with dbms_scheduler
12.9 Monitor the flash recovery area
12.10 Recovering from loss of a Control file
12.11 Recovering from loss of a Redo log file
12.12 Recovering from loss of a system-critical data file
12.13 Recovering from loss of a non system-critical data file

REFRENCES

1. Rick Greenwald, Robert Stackowiak, Jonathan Stern : Oracle Essentials : Oracle Database 10g – O’Reilly
2. Gavin Powell : Oracle High Performance Tuning for 9i and 10g – Digital Press
3. Loney, Kevin, Bryla, Bob : Oracle Database 10g, DBA Handbook – Oracle Press
4. Loney, Kevin : Oracle Database 10g The Complete Reference – Oracle Press
5. Micheal Abbey, Ian Abramson Oracle Database 10g: A Beginner’s Guide – Osborne Oracle Press Series

VEER NARMAD SOUTH GUJARAT UNIVERSITY, SURAT

M.C.A. (5 th Semester)
Proposed Syllabus
Paper : 505 : Data ware Housing and Data mining
Effec`tive From June 2008

1. Introduction
- 1.1 Data Warehouse characteristics
- 1.2 Data Marts

2. Online Analytical Processing
- 2.1 OLTP and OLAP systems
- 2.2 Star schema for multidimensional view
- 2.3 Multifact star schema or snow flake schema
- 1.4 OLAP Tools

3. Developing A Data Warehouse
- 3.1 Building a Data Warehouse
- 3.2 Architectural strategies & Design issues
- 3.3 Data Content
- 3.4 Metadata
- 3.5 Distribution of data
- 3.6 Tools for Data Warehousing
- 3.7 Performance considerations

4. Data Mining
- 4.1 Introduction
- 4.2 Data Description
  - 4.2.1 Clustering
  - 4.2.2 Link Analysis
- 4.3 Predictive Data Mining
  - 4.3.1 Classification
  - 4.3.2 Regression
  - 4.3.3 Time Series
- 4.4 Models & Patterns
  - 4.4.1 Decision Trees
  - 4.4.2 Multivariate adaptive regression splines
  - 4.4.3 Rule Induction
  - 4.4.4 K-nearest neighbour and memory based reasoning
  - 4.4.5 Logistic regression
  - 4.4.6 Discriminant Analysis
  - 4.4.7 Generalized Adaptive models
  - 4.4.8 Genetic algorithms
  - 4.4.9 Pattern Structures

4.4.9.1 Patterns in Data Matrices
4.4.9.2 Patterns for Strings

5 Applications of Data Warehousing and Data Mining

6. At least two case studies of Data Mining

REFRENCES

1. R. Kinball: Data Warehouse Toolkit – John Wiley & Sons
2. Efrem G. Mallach : Decision Support and Data Warehouse Systems – TMH
3. Paulraj Pulliah : Data Warehousing Fundamentals – Wiley
4. S. Anahory & D. Murray: Data Warehousing in the real world – Addison Wesley
5. R. Kinball, L.Reeves : The Data Warehouse Lifecycle Toolkit – John Wiley & Sons
6. David Hand, Heikki Mannila,Padhraic Smyth : Principles of Data Mining-PHI
7. C.S.R.Prabhu : Data Warehousing – PHI
8. Hillol Kargupta, Anupam Joshi, Yelena Yesha, Krishnamoorthy Sivakumar : Data Mining Next Generation Challenges & Future Directions – PHI
9. Jiawei Han, Micheline Kamber : Data Mining Concepts & Techniques
10. Dunham : Data Mining Introductory and Advanced Topics - Pearson