

□

DB2 9 for LUW Advanced Database Recovery
Information

Length: 4.0 Days
Ref: CL492G-X
Delivery method: Classroom
Price: AUD

Overview

Gain a deeper understanding of the advanced recovery features of DB2 9 for Linux, UNIX, and Windows database environments with single and multiple partition databases. Get practical experience in the planning and utilization of a wide variety of DB2 recovery facilities in a series of database recovery scenarios you complete during lab exercises using DB2 Enterprise 9.7 for Linux.

The purpose of this course is to enable the learning to gain a deeper understanding of the advanced recovery features of DB2 9 for Linux, UNIX, and Windows database environments with single and multiple partition databases. In this course, the learning will receive practical experience in the planning and utilization of a wide variety of DB2 recovery facilities in a series of database recovery scenarios you complete during lab exercises using DB2 Enterprise 9.7 for Linux.

Public

This advanced course is for experienced database administrators who plan and implement recovery and high availability plans for DB2 9 for Linux, UNIX and Windows databases using single or multiple partition DPF databases. The lab exercises can be performed using a single partition or a multiple partition DPF database using DB2 Enterprise 9.7 for Linux.

This course is appropriate for those using DB2 in a z/Linux environment.

Prerequisites

You should complete:

- *DB2 9 Database Administration Workshop for Linux, UNIX, and Windows (CL2X2)*
- *DB2 9 Database Administration Workshop for Linux (CL202)***or**
- *DB2 9 Database Administration Workshop for UNIX (CL212)***or**
- *DB2 9 Database Administration Workshop for Windows (CL232)***or**
- *Fast Path to DB2 9 for Experienced Relational DBAs (CL282)***or**
- *DB2 9 for Linux, UNIX, and Windows Quickstart for Experienced Relational DBAs (CL482)***or**
- *DB2 9 for LUW Multiple Partition DBA Workshop (CL24)***or**

- *DB2 9 for LUW Multiple Partition Environment for Single Partition DBAs (CL25)*
- **or** have equivalent experience.

Objective

- Describe the unique recovery planning requirements for DB2 9 single partition and DPF multiple partition databases
- Explore the DB2 for Linux, UNIX and Windows recovery facilities and database configuration options
- Plan the implementation of automated archival and retrieval of database logs
- Recover a DB2 table following a DROP TABLE command issued in error
- Utilize the REBUILD option of the RESTORE Utility to recover a full or partial database copy using either database or table space backups
- Plan and execute the recovery of table spaces to a selected point in time
- Effectively utilize incremental backup and restore to reduce the size and duration of DB2 database backups
- Describe the database crash recovery processing performed when there is an unplanned outage of a DB2 database server and select database configuration options to minimize the restart time.
- Utilize the redirected restore option to recover DB2 data to alternate disk configurations and invoke the db2relocatedb command to alter the configuration of a DB2 database
- Execute recovery scenarios, including loss of DB2 log data using the DB2 log mirroring option, and configure a database for automatic backups
- Utilize the information in the DB2 recovery history file to plan and execute various DB2 utilities
- Configure a DB2 database to automatically remove recovery objects that are no longer needed, including backups and archived logs
- List the benefits and limitations of disaster recovery alternatives including log shipping to a standby database or using DB2 replication
- Utilize the db2haicu command to select options for the integrated high availability cluster support for DB2 databases
- Explore the options for operation of DB2 databases in high availability environments including the use of a split mirror copy of the database
- Plan, implement, and manage the Primary and Standby databases using the High Availability Disaster Recovery (HADR) facilities of DB2 9 for Linux, UNIX and Windows databases
- Implement Read-Only application access to the Standby database in a HADR environment

Topics

DB2 Database Recovery Review

- State the three types of recovery support
- Explain the basic principles DB2 uses for database recovery and restart
- Describe the capabilities and functional requirements of DB2 logging
- Select options for some of the DB2 database parameters that support recovery

DB2 Recovery Log Management

- Select the DB2 Database parameters that support the DB2 logs
- Configure the database options to automate archive log management
- Implement infinite active logging to support applications that require large amounts of logged database changes
- Describe the options provided by DB2 9 for Linux, UNIX, and Windows to reduce log volume, including the NOT LOGGED INITIALLY option

DB2 Dropped Table Recovery

- Explain the DB2 facility to recover from an SQL DROP TABLE statement issued in error
- Implement dropped table recovery for selected table spaces
- Use the DB2 LIST HISTORY command to gather the information needed to recover a dropped table
- Plan and invoke the DB2 utilities involved in recovering a dropped table

Database Rebuild Support

- Review the considerations of using standard DB2 database recovery options
- Explain the capabilities of the REBUILD option for the RESTORE command
- List the types of information included in each DB2 backup image and describe how it is used to support rebuilding a database
- Plan for supporting database and disaster recovery scenarios using DB2 database and table space backups using the RESTORE command with a REBUILD option
- Utilize LIST UTILITIES SHOW DETAIL to monitor progress of a RESTORE Utility during database rebuilding

Table Space Recovery

- Plan for point-in-time recovery of a table spaces to support application processing errors
- Explain how the minimum recovery time for a table space may impact application point in time recovery
- Use the ONLINE option for RESTORE and ROLLFORWARD commands to recover selected table spaces with an active database

DB2 Incremental Backup and Recovery

- Plan a database recovery strategy that includes both full and Incremental backups to reduce the duration and size of database backups
- Implement a physical database design to take advantage of Incremental backups of selected table spaces
- Utilize an Incremental restore to recover a DB2 database or table space from Incremental backup images

- Use the LIST UTILITIES command to track the processing of an Incremental backup or restore process

DB2 Database Crash Recovery

- Describe the concepts and processing for DB2 database crash recovery
- Select an appropriate value for SOFTMAX to reduce database crash recovery time
- Configure DB2 database options to support crash recovery for applications
- Describe the impact of a database crash and restart for DB2 Utilities like LOAD, REORG and IMPORT

DB2 Database and Table Space Relocation

- Explain the facility of the DB2 RESTORE command to recover table spaces to different containers
- Use the SET TABLESPACE CONTAINERS command to define new containers during a redirected restore
- Utilize the RESTORE utility to change the database storage paths for AUTOMATIC STORAGE table spaces
- Plan the use of redirected restore as part of a disaster recovery
- Describe two methods that can be used to convert a DMS table space to utilize automatic storage
- Use the GENERATE SCRIPT option of RESTORE to set up a command script for a redirected restore operation
- Copy schemas from one database to another using the TRANSPORT option of the RESTORE utility
- Use db2relocatedb when moving or copying DB2 databases with non-DB2 utilities

DB2 Additional Recovery Facilities

- Explain the LOAD Utility options COPY YES, COPY NO and NONRECOVERABLE impact on database recovery
- Select the appropriate method for protecting the database from loss of log file
- Select the BACKUP Utility option to include the required logs in an online backup image
- Use the DB2 ARCHIVE LOG command to make the current DB2 log available for archival
- Implement the DB2 option for blocking update transactions when the log disk becomes full
- Implement the DB2 Utility Throttle using the UTIL_IMPACT_LIM Database Manager configuration option
- Configure a DB2 database for automatic backups
- Explain the self optimizing features of the Backup Utility

DB2 Recovery History File

- Explain the information recorded in the Recovery History file
- Use the DB2 LIST HISTORY command to display selected information needed to plan database

recovery

- Retrieve information from the database recovery history using SQL
- Configure a DB2 database for automatic removal of database recovery objects including backups and archived logs
- Use the PRUNE HISTORY command to remove outdated information in the Recovery History
- Restore a damaged Recovery History file

High Availability and Split Mirror Database Copy

- Describe various techniques for supporting High Availability of DB2 databases
- Explain the concepts involved when implementing the automated failover of a DB2 database
- Describe the integrated cluster configuration and management features of DB2
- Discuss the concepts used by DB2 pureScale to support High Availability for a DB2 database
- Plan for using split mirror copies to create a standby database, a database snapshot or a database mirror
- Utilize the SET WRITE SUSPEND and DBD2INIDB commands with split mirror database copies
- Use the AS SNAPSHOT option of the DB2 BACKUP and RESTORE utilities

Basic Database Disaster Recovery Techniques

- Describe the benefits and limitations of several methods for supporting database disaster recovery
- Explain setup for a Standby database at a remote location
- Plan for database disaster recovery using remotely mirrored disks
- Describe the use of DB2 Replication to recover selected DB2 tables in a remote database

DB2 High Availability Disaster Recovery (HADR)

- Describe the benefits and limitations of implementing High Availability Disaster Recovery (HADR)
- Implement HADR, including the setup of a Primary and Standby database
- Use the commands START HADR, STOP HADR, and TAKEOVER to manage an operational HADR system
- Plan for the use of the DB2 Utilities like LOAD, REORG, Backup and Restore in an HADR environment
- Select appropriate HADR database configuration options in order to meet the requirements of a specific application environment including the HADR_PEER_WINDOW
- Monitor the status of the Primary and Standby database using GET SNAPSHOT FOR DATABASE or db2pd - hadr
- Implement Read-only application access to the Standby Database of an HADR database environment

DB2 DPF Partitioned Database Recovery Considerations

- Explain database recovery facilities in a DB2 DPF partitioned database
- Describe the use of the Single System View options for Backup of a partitioned database
- Recover a partitioned database using the TO END OF LOGS option of the ROLLFORWARD command
- Create an online Backup of a DPF partitioned database that includes all of the required active logs
- Implement automated archival of DB2 logs in a DPF partitioned database
- Perform dropped table recovery in a DPF partitioned database
- Use the REBUILD option of RESTORE to build a partial copy of a DPF Partitioned database
- Describe the crash recovery processing performed for a DPF partitioned database
- Explain the use of redirected restores and db2relocatedb with DPF partitioned databases
- Use db2inidb for split mirror copies with a DPF partitioned databases

Agenda

Day 1

- Unit 1: DB2 Database Recovery Review
- Exercise 1: Configuring DB2 for Recovery
- Unit 2: DB2 Recovery Log Management
- Exercise 2: Managing DB2 Log Files
- Unit 3: Dropped Table Recovery
- Exercise 3: DB2 Dropped Table Recovery

Day 2

- Unit 4: Database Rebuild Support
- Exercise 4: Database Rebuild Support
- Unit 5: Table Space Recovery
- Exercise 5: Table Space Point-in-Time Recovery
- Unit 6: Incremental Backup and Recovery
- Exercise 6: Incremental Backup and Restore
- Unit 7: DB2 Database Crash Recovery

Day 3

- Exercise 7: Database Crash Recovery
- Unit 8: DB2 Database and Table Space Relocation
- Exercise 8: DB2 Table space Relocation
- Unit 9: DB2 Additional Recovery Facilities
- Exercise 9: Additional Recovery Facilities
- Unit 10: DB2 Recovery History File
- Unit 11: High Availability and Split Mirror Database Copy

Day 4

- Exercise 10: DB2 High Availability using split mirror copies
- Unit 12: Basic Database Disaster Recovery Techniques
- Exercise 11: DB2 Database Disaster Recovery using Log Shipping
- Unit 13: DB2 High Availability Disaster Recovery (HADR)
- Exercise 12: DB2 High Availability Disaster Recovery (HADR)
- Unit 14: DB2 DPF Partitioned Database Recovery

□