

Course name: DB2 10 for LUW Advanced Database Recovery (CL493G)

Course code: IBM-DB2-RCV-EN

Introduction

Gain a deeper understanding of the advanced recovery features of DB2 10 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 Advanced Enterprise Edition 10.5 for Linux.

Participant profile

This advanced course is for experienced database administrators who plan and implement recovery and high availability plans for DB2 10 for Linux, UNIX and Windows databases using single or multiple partition databases. The lab exercises can be performed using a single partition or a multiple partition database using DB2 Advanced Enterprise Edition 10.5 for Linux. This course is appropriate for those using DB2 in a z/Linux environment.

Goal description

- Describe the unique recovery planning requirements for DB2 10 single partition and 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
- List the benefits and limitations of disaster recovery alternatives including log shipping to a standby database or using DB2 replication
- Describe the use of the db2haicu command to select options for the integrated high availability cluster support for DB2 LUW 10 databases
- List some of the advantages and disadvantages associated with various High Availability and Disaster Recovery configurations DB2 databases
- Describe some of the differences between DB2 recovery facilities for DB2 pureScale database clusters compared to non-pureScale databases, including the use of the Cluster Caching Facilities to support Member Crash Recovery
- Plan and implement the use of split mirror copies of DB2 databases to create snapshot database copies or use as an alternative to a standard DB2 database backup.
- Plan, implement, and manage the Primary and Standby databases using the High Availability Disaster Recovery (HADR) facilities of DB2 10.5 for Linux, UNIX and Windows databases
- Implement Read-Only application access to the Standby database in a HADR environment
- Manage a HADR environment with multiple standby databases

Course duration and form

- 28 hours (4 days x 7 hours), including lectures and exercises.

Course plan

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: DB2 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: DB2 Incremental Backup and Recovery
- Exercise 6: Incremental Backup and Restore
- Unit 7: DB2 Database Crash Recovery
- Exercise 7: Database Crash Recovery

Day 3

- 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: Using Split Mirror Database Copies
- Exercise 10: Working with using split mirror database copies
- Unit 11: High Availability Options for DB2 LUW

Day 4

- Unit 12: DB2 High Availability Disaster Recovery (HADR) – Introduction
- Exercise 11: Implementing a HADR Standby Database
- Unit 13: DB2 High Availability Disaster Recovery (HADR) – Standby Database options
- Exercise 12: Implementation of Multiple Standby databases and Reads on a Standby Database
- Unit 14: DB2 pureScale Database Recovery Considerations
- Unit 15: DB2 Partitioned Database Recovery Considerations