

DB2 9.8 pureScale Implementation and Control for DB2 LUW Administrators Information

Length: 2.0 Days
Ref: CL800G
Delivery method: Classroom
Price: EUR

Overview

This course is intended to help students prepare to plan, implement and provide on-going support for a DB2 9.8 pureScale database cluster. The course explains the differences between a DB2 9.8 pureScale database and the standard DB2 LUW database system running on a single host system. The students are expected to have experience working with DB2 LUW databases on UNIX, Microsoft Windows or Linux systems. The course is lecture-only, no lab exercises are performed during the class.

Public

This course is intended for experienced DB2 LUW database administrators. Basic DB2 LUW functions and administration tasks are not covered. For experienced database administrators that have not previously supported DB2 LUW databases, some pre-requisite training should be taken before attending this class.

Prerequisites

The student could attend any of the following training courses to obtain the database administration skills that are prerequisites for the CL800 course:

- *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 (CL283)***or**
- *DB2 9 for Linux, UNIX, and Windows Quickstart for Experienced Relational DBAs (CL483)***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 major functional components of a DB2 pureScale database cluster
- Explain how Remote Direct Memory Access (RDMA) is used to efficiently process requests within a

DB2 pureScale database

- Describe the use of local and group buffer pools for reading pages with pureScale
- Configure the Cluster Caching Facility options to support shared memory for buffer pool and global locking
- Describe the workload balancing and automatic client routing options used for application connections to a pureScale database
- Plan the installation of DB2 pureScale software for a database cluster
- Estimate Cluster Caching Facility (CF) CPU and memory resources that will be needed to support your databases in the cluster
- Migrate an existing DB2 9.7 database system to a DB2 pureScale clustered environment
- Configure database manager and database member options for a DB2 pureScale cluster
- List the steps required to migrate an existing DB2 9.7 database system to a DB2 9.8 pureScale cluster
- Invoke the db2iupdt command to add or remove database members or cluster caching facilities from the cluster
- Monitor the buffer pool efficiency of local and group buffer pools
- Implement transaction-level workload balancing for application client systems to improve cluster resource usage
- Explain the difference between Member Crash Recovery and Group Crash Recovery
- Describe the use of multiple log streams for database members of a pureScale database
- Use the db2cluster command to add and drop disk devices with DB2 managed GPFS file systems
- Use the db2instance and db2cluster commands to show the current status and alter information for database members and cluster hosts
- Query the DB2 cluster status using SQL functions and views

Topics

DB2 pureScale Database Concepts

- List the basic components and functions in a DB2 pureScale database system
- Describe the types of information managed in the Cluster Caching Facility for a DB2 pureScale database
- Explain how Remote Direct Memory Access (RDMA) is used to efficiently process requests within a DB2 pureScale database
- Describe the workload balancing and automatic client routing options used for application connections to a pureScale database
- List the steps used to perform restart light processing when a database member fails
- Compare some of the characteristics of a DB2 partitioned database to the cluster implementation used by a DB2 pureScale database

Inside a DB2 pureScale database

- Describe the use of local and global locking to support application concurrency in a DB2 pureScale database
- Explain the use of page level locks with a pureScale database
- Monitor lock related statistics for a a DB2 pureScale clustered database
- Explain how multiple log streams are used and managed for a pureScale database
- Describe the use of local and group buffer pools for reading pages with pureScale
- Monitor the buffer pool efficiency of local and group buffer pools
- Explain how page reclaims are used in pureScale databases to access pages being changed by multiple database members
- Describe how castout processing is used to perform page writes with DB2 pureScale databases

Managing a DB2 pureScale Database Cluster

- Use db2start and db2stop commands to manage the components of a DB2 pureScale cluster
- Examine the db2instance command report to check the current status of the pureScale cluster
- Configure database manager and database member options for a DB2 pureScale cluster
- Invoke the db2iupdt command to add or remove database members or cluster caching facilities from the cluster
- Quiesce a host system in a pureScale cluster to perform system maintenance
- Explain how Automatic Client Reroute is used to maintain high availability for application client systems
- Implement transaction-level workload balancing for application client systems to improve cluster resource usage
- Configure client affinity for application systems that need to control connections to specific database member host systems

DB2 pureScale Database Backup and Recovery

- Explain the difference between Member Crash Recovery and Group Crash Recovery
- Plan for online or offline database BACKUP for a DB2 pureScale database
- Use the RESTORE utility or the RECOVER DATABASE command to perform database recovery or to make a database copy for a pureScale database
- Describe the impact of database cluster topology changes on database backup and recovery operations
- Explain the impact of multiple log streams for database roll forward processing with a standby database copy

Planning and Installation for a DB2 pureScale Database Cluster

- Select a set of supported system hardware and software components for a DB2 pureScale cluster
- Perform the pre-installation planning to insure the system software includes all necessary functions and service levels

- Choose the cluster interconnect network options for the pureScale cluster
- Estimate Cluster Caching Facility (CF) CPU and memory resources that will be needed to support your databases in the cluster
- Set the DBM and database configuration options and DB2 registry variables used to allow multiple active databases
- Perform the DB2 pureScale feature installation for the DB2 cluster

Migrating a DB2 LUW database to a DB2 pureScale Database Cluster

- List the steps required to migrate an existing DB2 9.7 database system to a DB2 9.8 pureScale cluster
- Select the best method to migrate any non-automatic storage table spaces into automatic storage table spaces based on the application requirements and current database structure
- Implement the GPFS file systems planned to support the DB2 pureScale database

Storage Management for a DB2 pureScale Database Cluster

- Use the db2cluster command to create the GPFS file systems needed to support DB2 pureScale databases
- Describe the difference between DB2 managed and user managed GPFS file systems
- Select the number of automatic storage paths that will be used to support the DB2 databases
- Create a new DB2 pureScale database using GPFS based file systems
- Use the db2cluster command to add and drop disk devices with DB2 managed GPFS file systems
- Utilize the - rebalance option of db2cluster to improve efficiency of DB2 managed GPFS file systems

Monitoring and Problem Determination for a DB2 pureScale Database Cluster

- Monitor application activity in a DB2 pureScale database cluster
- Utilize db2pd commands to display Cluster Caching Facility memory allocations and connection pool information
- Use the db2instance and db2cluster commands to show the current status and alter information for database members and cluster hosts
- Query the DB2 cluster status using SQL functions and views
- Perform problem determination using the diagnostic messages generated by the DB2 instance and Cluster Caching facilities
- Use TSA Issam report to view status of cluster resources
- Examine the message logs maintained by DB2 Cluster Services components

Agenda

Day 1

- Unit 1: DB2 pureScale Database Concepts

- Unit 2: Inside a DB2 pureScale database
- Unit 3: Managing a DB2 pureScale Database Cluster
- Unit 4: DB2 pureScale Database Backup and Recovery

Day 2

- Unit 5: Planning and Installation for a DB2 pureScale Database Cluster
- Unit 6: Migrating a DB2 LUW database to a DB2 pureScale Database Cluster
- Unit 7: Storage Management for a DB2 pureScale Database Cluster
- Unit 8: Monitoring and Problem Determination for a DB2 pureScale Database Cluster

□