
IBM DB2 9 for LUW Multiple Partition Environment for Single Partition DBAs

Duration: 3 Days Course Code: CL250G

Overview:

This course teaches you how to perform database administration tasks using IBM DB2 Enterprise 9 with the Database Partitioning Feature. This course provides a fast path to DB2 Enterprise 9 partitioned database administration skills for experienced DB2 single partition database administrators. These tasks include customization for the DB2 Enterprise 9 with DPF system, creating and populating partitioned databases, designing a database for parallel use, and using partitioned DB2 utilities. There is no actual installation of DB2 Enterprise software in this workshop. A DB2 for Linux system is used to exercise various administrative functions.

Course Materials

The course materials address DB2 9.7 for Linux, UNIX, and Windows.

Hands-On Labs

Eleven labs are included to address DB2 9.7 for Linux, UNIX, and Windows.

Certification

Prepare for IBM Certification Test 731: DB2 9 for Linux, UNIX, and Windows Database Administration.

Training Path

This course is part of an IBM Training Path. Taking this course in the recommended sequence allows you to maximize the benefits from your education.

Target Audience:

DB2 for Linux, UNIX and Windows single partition experienced database administrators involved in planning, implementing, or maintaining DB2 multi partitioned DPF databases.

Objectives:

- Describe the steps to install and customize DB2 Enterprise 9 in a DPF partitioned environment
 - Load data on DB2 Enterprise 9 in a partitioned environment
 - Define a DB2 DPF partitioned recovery strategy and perform the tasks necessary to support the strategy
 - Identify how a database should be designed to take advantage of the parallel architecture
 - Use the DB2 utilities to manage data and maintain your DPF partitioned database
-

Prerequisites:

You should be able to:

- Use basic UNIX functions such as utilities, file permissions, hierarchical file system, commands, and the vi editor
 - Administer a DB2 for Linux, UNIX and Windows single partition database
-

Content:

DB2 DPF Partitioned Database Components and Concepts

- Identify the kinds of business applications that use parallel databases
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up

- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding

Agenda

Day 1

- Welcome
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs

- and recover a DPF partitioned database
 - Recover the database to a prior point in time using the RECOVER DATABASE command
 - State general considerations regarding disaster recovery and implementation of a server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command
 - Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
 - Describe the basic architecture of each of
- disaster recovery and implementation of a server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command
 - Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
 - Describe the basic architecture of each of
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
 - Recover the database to a prior point in time using the RECOVER DATABASE command
 - State general considerations regarding disaster recovery and implementation of a server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command
 - Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab

the hardware platforms on which parallel databases may run

- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group

the hardware platforms on which parallel databases may run

- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to

- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a

- Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled

- add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
- Partitioned Database Backup and Recovery
- Describe the three types of recovery support provided by DB2
 - Describe the basic architecture of each of

- server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command
 - Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture

- List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and loading data
 - Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
 - Explain the principles DB2 uses for its recovery/restart functions
 - Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
 - Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
 - Recover the database to a prior point in time using the RECOVER DATABASE command
 - State general considerations regarding disaster recovery and implementation of a server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command
 - Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
- the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and loading data
 - Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
 - Explain the principles DB2 uses for its recovery/restart functions
 - Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
 - Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
 - Recover the database to a prior point in time using the RECOVER DATABASE command
 - State general considerations regarding disaster recovery and implementation of a server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command
 - Utilize the REDISTRIBUTE command to
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and loading data
 - Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
 - Explain the principles DB2 uses for its recovery/restart functions
 - Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
 - Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
 - Recover the database to a prior point in time using the RECOVER DATABASE command
 - State general considerations regarding disaster recovery and implementation of a server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command
 - Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution

- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

Getting Started with DB2

- Describe the purpose of DAS

Installation

- Identify the steps to install and customize DB2 in a partitioned environment
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces

add or remove database partitions from a database partition group

- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a

- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage

- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility

- DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL

- Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index

- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

Creating a Partitioned DB2 Database

- Describe DB2 database partitioning
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory

- optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic

- reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database

- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to

- Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect

- directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace

- obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and

- detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces

- facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces

- db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the

- Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility

- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory

- database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment

- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab

- Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and loading data
 - Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
 - Explain the principles DB2 uses for its recovery/restart functions
 - Describe the configuration options for DB2 logging and explain the differences between circular and archive logs

- facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Scaling the Database
 - Identify how the partitioned database configuration can be scaled by adding new database partitions
 - Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database

- Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and

- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

- directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types

- loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2

- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add

- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and

- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

Day 2

- Data Placement on the Partitions
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and

- or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned

databases

- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab

load processing

- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab

- environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces

- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

Partitioning and Database Partition Groups

- Identify considerations when choosing a distribution key
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces

- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

- and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL

- database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Database and Application Performance
- Describe the steps used by the DB2 Optimizer to generate access plans
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces

- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on

- statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups

- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from

- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the

- database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status

- other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics,

- database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

Data Placement on the Partitions

- List the advantages and disadvantages for each type of table space management
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space

- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation

- database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table

- status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan
- spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan

- Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and
-
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and loading data
 - Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
 - Explain the principles DB2 uses for its recovery/restart functions
 - Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
 - Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
 - Recover the database to a prior point in time using the RECOVER DATABASE command
 - State general considerations regarding disaster recovery and implementation of a server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command
 - Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
-
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL

- loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool
- statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool

- Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and loading data
 - Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
 - Explain the principles DB2 uses for its recovery/restart functions
 - Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
 - Use the BACKUP, RESTORE and

- loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1

- to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

Day 3

- Partitioned Database Backup and Recovery
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases

- ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespace Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information

- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space

- database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespace Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation

- Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using

- reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and loading data
 - Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
 - Explain the principles DB2 uses for its recovery/restart functions
 - Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
 - Use the BACKUP, RESTORE and
- buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and loading data
 - Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
 - Explain the principles DB2 uses for its recovery/restart functions
 - Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
 - Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
 - Recover the database to a prior point in time using the RECOVER DATABASE command
 - State general considerations regarding disaster recovery and implementation of a server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command
 - Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab

- statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing

- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

Moving Data

- Utilize the EXPORT Utility to extract data from a DB2 database
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces

- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command

- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the db2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance

- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database

- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
 - Troubleshoot problems that may occur during redistribution
 - List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
 - Utilize the RUNSTATS utility to collect detailed table and index statistics
 - Explain the major goals for table and index reorganization
 - Use the REORGCHK report to plan effective use of the REORG utility
 - Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
 - Identify the different monitor types
 - Identify how to use the independent trace facility (db2trc)
 - Use the dp2pd problem determination tool to obtain statistics from a running instance
 - Identify additional commands for listing application information
 - Discuss available lock modes and their compatibility
 - Influence locking strategies used by the database manager
 - DB2 DPF Partitioned Database Components and Concepts
 - Getting Started with DB2
 - Installation
 - Installation Lab
 - Creating a Partitioned Database
 - Creating Database Lab
 - Partitioning and Database Partition Groups
 - Partitioning and Database Partition Groups Lab
 - Creating Tablespaces Lab
 - Create Objects Lab
 - Moving Data - Part 1
 - Moving Data Lab - Import and Load
 - Moving Data - Part 2
 - Moving Data Lab - Partitioned Load
 - Recovery Lab
 - Scaling the Database
 - Scaling the Database Lab
 - Database and Application Performance
 - Database and Application Performance Lab
 - Monitoring and Problem Determination
 - Locks and Concurrency
 - Monitoring and Problem Determination Lab
- Locks and Concurrency
- List objects that may be locked by the database manager
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
 - Define the strengths of the DB2 DPF partitioned architecture
 - Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
 - Use db2_all and rah commands
 - Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
 - List the three types of Storage Management for table spaces
 - Describe the three default system table spaces
 - Access and update the database manager configuration file and the database configuration files
 - Access and update the system database directory and list the local database directory
 - Use tools to issue commands and SQL statements
 - Describe join strategies in partitioned databases
 - Identify the syntax for creating database partition groups
 - Create SMS table spaces
 - Create DMS tablespaces
 - Create Automatic Storage managed table spaces
 - Use the GET SNAPSHOT commands and db2pd commands to check table space status
 - Use SQL functions specific to the partitioned environment
 - Identify catalog views that contain information about your partitioned environment
 - Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
 - Describe the process of partitioning and loading data
 - Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
 - Explain the principles DB2 uses for its recovery/restart functions
 - Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
 - Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
 - Recover the database to a prior point in time using the RECOVER DATABASE command
 - State general considerations regarding disaster recovery and implementation of a server cluster for high availability
 - Differentiate between the three options on the REDISTRIBUTE command

- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family

- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration,

- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS table spaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Describe the basic architecture of each of the hardware platforms on which parallel databases may run
- Define the strengths of the DB2 DPF partitioned architecture
- Identify the key features that set DB2 in a DPF partitioned environment apart from other members of the DB2 family
- Use db2_all and rah commands
- Create a database in a DB2 partitioned environment with or without Automatic Storage enabled
- List the three types of Storage Management for table spaces
- Describe the three default system table spaces

- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility
- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab
- Access and update the database manager configuration file and the database configuration files
- Access and update the system database directory and list the local database directory
- Use tools to issue commands and SQL statements
- Describe join strategies in partitioned databases
- Identify the syntax for creating database partition groups
- Create SMS tablespaces
- Create DMS tablespaces
- Create Automatic Storage managed table spaces
- Use the GET SNAPSHOT commands and db2pd commands to check table space status
- Use SQL functions specific to the partitioned environment
- Identify catalog views that contain information about your partitioned environment
- Identify the different methods for inputting data, including the Import Utility, using buffered and unbuffered SQL INSERT and the LOAD Utility
- Describe the process of partitioning and loading data
- Identify how the partitioned database options on the LOAD command can be used to control the data partitioning and load processing
- Explain the principles DB2 uses for its recovery/restart functions
- Describe the configuration options for DB2 logging and explain the differences between circular and archive logs
- Use the BACKUP, RESTORE and ROLLFORWARD commands to back up and recover a DPF partitioned database
- Recover the database to a prior point in time using the RECOVER DATABASE command
- State general considerations regarding disaster recovery and implementation of a server cluster for high availability
- Differentiate between the three options on the REDISTRIBUTE command
- Utilize the REDISTRIBUTE command to add or remove database partitions from a database partition group
- Troubleshoot problems that may occur during redistribution
- List the major influences for SQL optimization, including Catalog statistics, database memory configuration, optimization class selection and database partition groups
- Utilize the RUNSTATS utility to collect detailed table and index statistics
- Explain the major goals for table and index reorganization
- Use the REORGCHK report to plan effective use of the REORG utility

- Describe the use of the DB2 explain tools to analyze access strategies for SQL statements
- Identify the different monitor types
- Identify how to use the independent trace facility (db2trc)
- Use the dp2pd problem determination tool to obtain statistics from a running instance
- Identify additional commands for listing application information
- Discuss available lock modes and their compatibility
- Influence locking strategies used by the database manager
- DB2 DPF Partitioned Database Components and Concepts
- Getting Started with DB2
- Installation
- Installation Lab
- Creating a Partitioned Database
- Creating Database Lab
- Partitioning and Database Partition Groups
- Partitioning and Database Partition Groups Lab
- Creating Tablespaces Lab
- Create Objects Lab
- Moving Data - Part 1
- Moving Data Lab - Import and Load
- Moving Data - Part 2
- Moving Data Lab - Partitioned Load
- Recovery Lab
- Scaling the Database
- Scaling the Database Lab
- Database and Application Performance
- Database and Application Performance Lab
- Monitoring and Problem Determination
- Locks and Concurrency
- Monitoring and Problem Determination Lab

Further Information:

For More information, or to book your course, please call us on 353-1-814 8200

info@globalknowledge.ie

www.globalknowledge.ie

Global Knowledge, 3rd Floor Jervis House, Millennium Walkway, Dublin 1