

SQL Server® High Availability: Hands-On - 3 Days

Course 141 Overview

- You Will Learn How To**
- Leverage SQL Server technologies to achieve high-availability database solutions
 - Design a Service Level Agreement that matches business requirements
 - Plan, prepare and install a SQL Server failover cluster
 - Maintain a warm standby server with log shipping
 - Mirror a database to ensure instant failover
 - Employ peer-to-peer transactional replication as a high-availability solution
- Course Benefits** SQL Server provides a full range of technologies that allows organizations to reduce downtime and maintain high levels of availability. In order to achieve these goals, an enterprise database administrator must implement failover clustering, log shipping, database mirroring and peer-to-peer transactional replication. This course offers enterprise SQL Server administrators the skills to maintain a large number of constantly available database servers.
- Who Should Attend** Anyone involved in planning, supporting or implementing a high-availability solution with SQL Server. Course 138, "SQL Server 2008 Database Administration," or equivalent experience is assumed.
- Hands-On Training** Throughout this course, a series of extensive hands-on exercises provides you with practical experience implementing high-availability solutions. Exercises include:
- Identifying availability requirements
 - Installing a SQL Server Failover Cluster
 - Backing up and restoring databases with advanced features
 - Maintaining a warm standby server
 - Mirroring a database
 - Implementing replication
 - Integrating high-availability techniques

SQL Server® High Availability: Hands-On - 3 Days

Course 141 Outline

High-Availability Fundamentals

Characterizing high availability

- The "five nines"
- Site
- Instance
- Database
- Defining a Service Level Agreement

Planned vs. unplanned downtime

- Performing routine maintenance
- Recovering from disasters
- Handling hardware failures

Configuring hardware

- RAID
- SAN

Installing Windows 2008 Failover Clustering

Defining components of a cluster

- Single instance architecture
- Multi-instance architecture
- Shared storage
- Resource Group
- Heartbeat
- Quorum

Preparing for Windows failover clustering

- Setting up an iSCSI target
- Managing shared storage
- Configuring iSCSI initiators
- Validating a cluster configuration
- Installing the failover clustering feature
- Creating a Windows Failover cluster

Establishing SQL Server Failover Clustering

Planning for SQL Server failover clustering

- Upgrading from Windows Server 2003 to 2008
- Migrating to SQL Server 2008
- Leveraging virtual machines for clustering
- Geographically dispersed clusters

Creating a SQL Server failover cluster

- Installing a new failover cluster
- Adding a new node
- Failing over to a passive node

Employing Log Shipping

Preparing for a warm standby server

- Synchronizing logins
- Creating network shares

Implementing transaction log shipping

- Specifying log backup frequency
- Configuring the log shipping monitor

Switching to the standby server

- Redirecting applications
- Bringing the standby online

Configuring Database Mirroring

Comparing mirroring architectures

- High availability
- High protection
- High performance
- Full safety vs. safety off

Getting ready for mirroring

- Setting the recovery model
- Selecting the principal, mirror and witness servers

Mirroring a database

- Configuring the principal, mirror and witness
- Initiating the mirroring session

Administering mirroring

- Enabling and disabling mirroring with scripts
- Launching the Database Mirroring Monitor (DMM)

Employing a mirror for reporting

- Taking a database snapshot
- Querying a snapshot of the mirror

Leveraging Replication for High Availability

Laying out a replication strategy

- Enumerating types of replication
- Making the business case

Configuring Peer-to-Peer Transactional Replication

- Creating distributors
- Initializing databases
- Creating and subscribing to a publication
- Adding a node to a topology

Combining High-Availability Technologies

Selecting the appropriate strategies

- Determining the pros and cons of each HA technology

- Failover clustering with a mirror
- Log shipping with a failover cluster

Pulling it all together

- Creating a complete high-availability solution
- Evaluating the results