

SQL 2012 Admin, Part 5 of 5: Disaster Recovery and Backups

page 1

Meet the expert: Chris Bell, MCITP, is an 18-year SQL Server veteran for both business intelligence and application development providing solutions for businesses, organizations, and individuals. He is the founder and CEO of WaterOx Consulting, Inc., a provider of remote SQL Server consulting and services. Chris is also the founder and current President of the Washington DC chapter of PASS and a member of the Board of Directors for CPCUG. Chris also frequently attends and presents at PASS events around the country, sharing his passion for all things SQL Server. In 2012, Chris was one of 5 finalists in the world for Red Gate's Exceptional DBA Award.

Prerequisites: This course assumes that students have working experience with SQL Server 2005 or 2008; basic relational database concepts (e.g., tables, queries, and indexing); general knowledge of XML, Transact-SQL, and a fundamental understanding of networking and security concepts.

Runtime: 02:02:42

Course description: SQL Server 2012 offers a variety of High Availability and Disaster recovery options. The biggest change from previous versions is the introduction of AlwaysOn. In this course we will look at how AlwaysOn integrates high-availability and disaster recovery into a total solution. We will see how AlwaysOn is based on Windows Failover Clusters to protect the entire Instance, while Availability groups provide options to improve database availability and resource use. Then we'll cover log shipping as well as mirroring and replication, which are still useful in some scenarios for disaster recovery. We will go through each of these, and even configure separate log shipping, mirroring and transactional replication. Next we will go through the various recovery models SQL supports and we will discuss the various backup scopes and types available. We will look into how these all work together to allow you to create a complete disaster recovery strategy. We will also show how having proper backups allow us to perform actions such as restoring a database to a point in time. We even fix a corrupted page by referencing a proper backup all while keeping the database online to users.

Course outline:

Mirroring

- Introduction
- Mirroring
- Demo: Mirroring
- Demo: Mirroring Prep
- Summary

Log Shipping

- Introduction
- Log Shipping
- Demo: Log Shipping
- Demo: Restore Transaction Log
- Demo: Update Transaction Log
- Summary

Failover Clustering

- Introduction
- Failover Clustering
- Demo: Failover Clustering
- AlwaysOn Failover Cluster
- AlwaysOn Availability Groups
- Demo: AlwaysOn Availability Groups
- Demo: Creating Availability Groups
- Demo: Using Availability Groups
- Demo: Starting Failover
- Summary

Replication

- Introduction

- Replication - Overview
- Replication - Types
- Replication - Snapshot
- Replication - Transactional
- Replication - Peer-to-Peer
- Replication - Merge
- Demo: Replication
- Demo: Replicated Database
- Demo: Add a Subscriber
- Demo: Replication and Transactions
- Demo: Replication Monitor
- Summary

Recovery Models

- Introduction
- Recovery Models - Simple
- Recovery Models - Full
- Recovery Models - Bulk Logged
- Demo: Recovery Models
- Demo: Bulk Logged
- Summary

Backups

- Introduction
- Backup Scopes - Database
- Backup Scopes - Partial
- Backup Scopes - File

- Backup Types - Full
- Backup Types - Differential
- Backup Types - Log
- Backup Types - Tail Log
- Backup Compression
- Restoring Databases
- Demo: Backups
- Demo: Corrupted Pages
- Summary