

DATABASES

ADVANCED

3 DAYS

SQL Server High Availability and Disaster Recovery: Always On, Backup and Recovery

Three days for teams accountable for SQL Server availability and recovery. Covers translating recovery objectives into architecture, backup and restore design, corruption response, Windows Server Failover Clustering, Always On Availability Group design and implementation, failover troubleshooting, and Azure disaster recovery. Build a multi-replica AG, fail it over under fault conditions and document a DR test.

Why SQL Server HA and DR Training?

Objectives first, technology second

Architecture decisions are driven from stated recovery time and recovery point targets rather than from a feature list.

Corruption handled properly

Detection, triage and page level recovery are practised, including the decisions that make corruption incidents worse.

Failover under fault conditions

Participants trigger real failures, including quorum loss and a corrupted replica, not just clean planned failovers.

Produces a testable DR plan

Teams leave with a runbook and a test procedure rather than a diagram nobody has exercised.

What you'll be able to do

- ✓ Translate business continuity requirements into recovery time and recovery point objectives
- ✓ Compare availability options and select the right one for a given constraint set
- ✓ Design a backup architecture that meets stated recovery objectives
- ✓ Execute complex restore scenarios including point-in-time and piecemeal recovery
- ✓ Detect, triage and recover from database corruption
- ✓ Configure Windows Server Failover Clustering and understand quorum behaviour
- ✓ Design Always On Availability Groups including replica roles and commit modes
- ✓ Implement availability groups with listeners and appropriate routing
- ✓ Configure read-scale, distributed and cross-platform availability groups
- ✓ Monitor synchronisation health and troubleshoot failed failovers
- ✓ Implement Azure based disaster recovery and hybrid topologies
- ✓ Write and execute a disaster recovery test plan

Who this is for

- Database administrators responsible for availability commitments
- Infrastructure and platform engineers supporting SQL Server clusters
- Site reliability engineers handling database incidents
- Solution architects designing resilient data tiers
- Teams preparing for audit, regulatory or contractual availability requirements

Curriculum

01 Defining Availability and Recovery Requirements

- Recovery time and recovery point objectives and how to elicit them honestly
- Availability targets, maintenance windows and the cost of each additional nine
- High availability versus disaster recovery versus backup, and why conflating them fails audits
- Failure domains: instance, host, storage, site and region
- Comparing options against cost, complexity and licensing
- **Hands-on:** produce a requirements matrix and shortlist an architecture

02 Recovery Models and Backup Architecture

- Recovery model selection and its effect on what is recoverable
- Full, differential and log backup chains and dependency
- Log growth, VLF management and the causes of a runaway log
- Backup compression, encryption and checksum verification
- Retention, offsite copies and immutability against ransomware
- **Hands-on:** design a backup schedule that meets a stated recovery point objective

03 Advanced Backup Strategy

- Filegroup and partial backups for very large databases
- Copy-only backups and avoiding chain breakage
- Backup to URL and Azure Blob Storage configuration
- Striping, throughput tuning and shrinking the backup window
- Automated verification and restore testing pipelines
- **Hands-on:** implement and verify a backup strategy for a multi-terabyte design

04 Restore and Point-in-Time Recovery

- Restore sequences, NORECOVERY, STANDBY and recovery states
- Point-in-time restore using STOPAT and marked transactions
- Piecemeal and filegroup restore to bring critical data online first
- Restoring to a different instance, version or environment
- Recovering a single object without restoring the whole database
- **Hands-on:** recover a database to a precise timestamp under a time limit

05 Corruption Detection and Response

- DBCC CHECKDB, its variants and an appropriate execution cadence
- Reading consistency check output and classifying severity
- Page level restore as the preferred recovery route
- Repair options, data loss implications and when to escalate instead
- Storage level causes and preventing recurrence
- **Hands-on:** detect injected corruption and recover with no data loss

06 Windows Server Failover Clustering Foundations

- Cluster architecture, nodes and resource management
- Quorum models, witnesses and vote configuration
- Cloud and file share witnesses for stretched clusters
- Cluster validation, networking and heartbeat configuration
- Failover cluster instances compared with availability groups
- **Hands-on:** build a cluster and test quorum loss behaviour

07 Always On Availability Group Design

- Replica roles, limits and edition constraints
- Synchronous and asynchronous commit and their latency cost
- Automatic, manual and forced failover behaviour
- Multi-subnet, multi-site and cross-region topologies
- Storage, network and latency requirements between replicas
- **Hands-on:** design a topology against a stated set of objectives

08 Implementing Availability Groups

- Prerequisites, endpoints and service account configuration
- Seeding options including automatic seeding for large databases
- Listener configuration, DNS and connection routing
- Adding and removing databases and replicas safely
- Contained availability groups and instance level object handling
- **Hands-on:** build a three-replica availability group with a working listener

09 Read-Scale, Distributed and Advanced Configurations

- Read-only routing lists and offloading reporting workloads
- Read-scale availability groups without a cluster
- Distributed availability groups for migration and cross-platform topologies
- Basic availability groups and Standard edition limitations
- Backups on secondary replicas and their constraints
- **Hands-on:** configure read-only routing and verify traffic separation

10 Monitoring, Failover and Troubleshooting

- Synchronisation state, send and redo queue monitoring
- Data loss exposure during asynchronous failover
- Failure conditions, health check timeouts and flexible failover policy
- Diagnosing a failed or unexpected failover from cluster and SQL logs
- Patching and maintenance with rolling upgrades
- **Hands-on:** force a failure, recover the group and account for any data loss

11 Disaster Recovery Beyond the Cluster

- Log shipping and replication as complementary DR layers
- Azure Site Recovery and hybrid disaster recovery topologies
- Azure SQL Managed Instance failover groups and geo-replication
- Runbook structure and the decisions that must be pre-authorised
- Scheduled DR testing, evidence capture and audit expectations
- Post-incident review and continuous improvement
- **Hands-on:** capstone disaster recovery test executed and documented end to end