

IE1 & IE2 Homework/Review Strategy

A schedule for pacing your review

Overview:

- IE1: Internals and Performance – 12 modules
- IE2: Performance Tuning – 12 modules

Use your time effectively:

By studying 2-3 modules per week (depending on module comfort, complexity and your available time), you will be able to better absorb and feel comfortable with the content. You should get through IE1/IE2 review in 11 weeks.

Week 1

IE1: Module 1 – Database Structures

(1) Review the MCM Online Videos for recap

- [Database Structures](#)
- [Database Structures - Demo](#)
- [New Database Structures in SQL Server 2008](#)
- [New Database Structures in SQL Server 2008 - Demo](#)

(2) Review the slides + the whiteboard + demo scripts for this module

(3) Go through this “optional exercise” (if you have any questions, contact Paul)

- a. One of the storage structures we did not have time to cover in depth is the way LOB values are stored.
 - i. Create a table with a LOB column value of something like 30KB (enough that it will require 4 text pages to store). Use DBCC IND (or the sp_allocationMetadata script) and DBCC PAGE to investigate how the LOB value is stored and how the pages are linked together.
 - ii. Create a table with a varchar(max) column and set the column to be stored off-row by default. Create multiple small values and then examine how they are stored on the same shared page.
 - iii. Create a nonclustered index with an INCLUDED LOB column and convince yourself that the INCLUDED column value does not

share the same physical storage as the column value in the table.

- b. Create two identical tables and enable page compression on one of them. Examine the page structures for each to see how they differ.
- c. Create a table with allocations from multiple data files. Examine the IAM chain for the table to see that it has one IAM page per data file. Truncate the table. Examine the IAM chain again. What happens to it when the table is truncate?
- d. Create a table with sparse columns and an XML COLUMN_SET and some data
 - i. Convince yourself with DBCC PAGE that the XML COLUMN_SET is not persisted
 - ii. Examine the sparse column array itself using DBCC PAGE
- e. Create a table with a FILESTREAM column
 - i. Create some FILESTREAM values and update some of them. Investigate garbage collection after reading through my blog post on the subject (FILESTREAM blog post category)
 - ii. Look at some of the filenames of the FILESTREAM files in the data container. Use the fn_dblog log exploration tool to find the log records where the files were created.

(4) Review the whitepaper: [FILESTREAM Storage in SQL Server 2008](#)

IE1: Module 2 – Data File Internals / Maintenance

(5) Review the MCM Online Videos for recap

- [Data File Internals and Maintenance](#)
- [Data File Internals and Maintenance - Demo](#)

(6) Review the slides + the whiteboard + demo scripts for this module

(7) Go through the HOL lab from your DVD: **SQL Server 2008 Instant Initialization**

- (8) Go through this “optional exercise” (if you have any questions, contact Paul)
- a. Create a large data file with and without instant initialization enabled to convince yourself of the performance difference.
 - b. If you have remote access to any servers, use the diskpart tool to investigate the partition offsets of some of the volumes.
 - c. Create a partitioned table with 10000 rows in each of 3 or more partitions

- i. Using the default lock escalation, cause lock escalation from a query updating one partition, and show that accessing another partition is blocked
 - ii. Change the lock escalation to AUTO and try exercise a) again
 - iii. Cause lock escalation with two queries updating separate partitions, then have each connection try to access the other locked partition. What happens?
- d. If you have remote access to any servers, check out the data file sizes and auto-growth settings for some of the databases. Are they optimal do you think?
- e. Enable snapshot isolation and create some long-running transactions that will stop the version store cleanup from happening
 - i. Examine the output from the sys.dm_db_file_space_usage DMV.
- f. Grow tempdb on a server to 100MB. Create an active workload using tempdb. Try shrinking tempdb. Did you see any corruption messages or were you lucky?
- g. Create a database snapshot of a database.
 - i. Examine the snapshot file sizes. Make various updates to the source database and examine the snapshot file sizes to see how they've changed.
 - ii. Shutdown and restart the instance to prove to yourself that a snapshot survives a shutdown

IE1: Module 5 – M5: Logging/Recovery, Log File Internals/ Maintenance

(9) Review the MCM Online Videos for recap

- [Log File Internals and Maintenance](#)
- [Log File Internals and Maintenance - Demo](#)

(10) Review the slides + the whiteboard + demo scripts for this module

(11) Go through this “optional exercise” (if you have any questions, contact Paul)

- a. Create a database in the SIMPLE recovery model and use it to investigate various log records. Doing a checkpoint between exercises will clear out the log as long as there aren't any active transactions.
 - i. Create a table and insert a row. Checkpoint. Insert some more rows and examine them with fn_dblog. Notice the implicit transactions have LOP_BEGIN_XACT and LOP_COMMIT_XACT log record.

- ii. Do the same as in a) but inside a transaction which you then rollback. Look at the Context column for the COMPENSATION log records. Look how the Previous LSN fields for the COMPENSATION log records point to the log record *before* the one being compensated for. Can you remember why?
 - iii. Create an active transaction and do a checkpoint. Investigate the checkpoint log record. Can you find the LSN of the start of the active transaction in the LOP_XACT_CKPT log record? Can you explain how crash recovery would know to start examining the log from there?
 - iv. Try allocating another page to the table. Examine the log records. Is the log record that alters the PFS page part of a transaction?
- b. Create a database in the FULL recovery model. Take a full backup and a log backup. Do some operations in the database. Switch to the SIMPLE recovery model, checkpoint, and switch back to the FULL recovery model. Try to take a log backup. Why do you get the error?
 - c. Create a database with some active transactions. Use the script I blogged in <http://www.sqlskills.com/BLOGS/PAUL/post/Script-open-transactions-with-text-and-plans.aspx> to investigate how much log has been generated for each transaction. Try doing more operations and seeing how the values change.
 - d. Create a database with default log growth. Put it into the FULL recovery model (full backup) and then generate lots of transaction log records. Use DBCC LOGINFO to investigate the VLFs. Follow the steps in the lecture to shrink the log down and grow it manually to an appropriate size.

Week 2

IE1: Module 3 – Locking and Blocking

(1) Review the MCM Online Videos for recap

- [Locking](#)
- [Locking - Demo](#)

(2) Review the slides + the whiteboard + demo scripts for this module

IE1: Module 4 – Snapshot Isolation

(3) Review the MCM Online Videos for recap

- [Snapshot Isolation](#)

(4) Review the slides + the whiteboard + demo scripts for this module

(5) Go through the HOL lab from your DVD: **SQL Server 2008 Snapshot Isolation**

IE1: Module 6 – Table Design

(6) Review the MCM Online Videos for recap

- [Sparse Columns](#)
- [Sparse Columns - Demo](#)

(7) Review the slides + the whiteboard + demo scripts for this module

Week 3

IE1: Module 7 – Index Internals

(1) Review the MCM Online Videos for recap

- [Index Internals](#)
- [Index Internals - Demo](#)
- [The Clustered Index Debate](#)

(2) Review the slides + the whiteboard + demo scripts for this module

IE1: Module 8 – Index Fragmentation

(3) Review the MCM Online Videos for recap

- [Index Fragmentation](#)
- [Index Fragmentation - Demo](#)

(4) Review the slides + the whiteboard + demo scripts for this module

- (5) Go through this “optional exercise” (if you have any questions, contact Paul)
- Create your own scenario where page splits occur. Investigate the log records generated.
 - Take the fragmentation-generating script I demo'd and alter it to insert 100000 or 1 million rows. Look at the effect on the fragmentation levels of the table with the GUID populated from NEWSEQUENTIALID. Try changing the nonclustered indexes to be on one of the 2008 datatypes with 100ns precision. What effect does it have on the fragmentation of both nonclustered indexes?
 - Create a large index and then perform an online index rebuild of it while there's a concurrent DML workload on the index. Examine the version store behavior in tempdb.
 - Try using SQL Profiler to get progress reporting of online index rebuild to work.

IE1: Module 9 – Internals & Data Access

- (6) Review the slides + the whiteboard + demo scripts for this module
- (7) Review the whitepaper: [Improving Performance with SQL Server 2008 Indexed Views](#) (NOTE: This applies more to the Indexing Strategies module; however, there aren't a lot of external resources for this module.)

Week 4

IE1: Module 10 – Statistics

- (1) Review the MCM Online Videos for recap
- [Statistics](#)
- (2) Review the slides + the whiteboard + demo scripts for this module
- (3) Review the statistics whitepaper: [Statistics Used by the Query Optimizer in Microsoft SQL Server 2008](#)

IE1: Module 11 – Indexing Strategies

- (4) Review the MCM Online Videos for recap
- [Indexing Strategies](#)
 - [Indexing Strategies - Demo](#)

(5) Review the slides + the whiteboard + demo scripts for this module

Week 5

IE1: Module 12 – Partitioning

(1) Review the MCM Online Videos for recap

- [Partitioning](#)

(2) Review the slides + the whiteboard + demo scripts for this module

(3) Go through the HOL lab from your DVD: **SQL Server 2008 Table Index & Partitioning**

(4) Read the SQL Server 2005 whitepaper on Partitioning:

[http://msdn.microsoft.com/en-US/library/ms345146\(v=SQL.90\).aspx](http://msdn.microsoft.com/en-US/library/ms345146(v=SQL.90).aspx)

(5) Read the SQL Server 2008 whitepaper on Partitioning:

<http://download.microsoft.com/download/D/B/D/DBDE7972-1EB9-470A-BA18-58849DB3EB3B/PartTableAndIndexStrat.docx>

(6) Go through the HOL lab from your DVD: **SQL Server 2008 Online Operations Hands-On Lab**

Week 6

IE2: Module 1: SQL Server IO – Storage Needs

(1) Review the slides + the whiteboard + demo scripts for this module

NOTE: This module is really an overview/review of IE1. Take some time to review any/all concepts from IE1 on which you feel a bit lighter/weak.

IE2: Module 2: IO for DBAs

(2) Review the MCM Online Videos for recap

- [Storage Core Concepts](#)
- [Storage Guidelines](#)

- (3) Review the slides + the whiteboard + demo scripts for this module
- (4) Review these whitepapers:

- [SQL Server I/O Basics \(SQL Server 2000\)](#)
- [SQL Server I/O Basics, Chapter 2 \(SQL Server 2005\)](#)

IE2: Module 3: IO – SANs for DBAs

- (5) Review the MCM Online Videos for recap

- [DAS vs SAN Architecture](#)
- [Storage Testing and Monitoring](#)

- (6) Review the slides + the whiteboard + demo scripts for this module

Week 7

IE2: Module 4: SQLOS

- (1) Review the MCM Online Videos for recap

- [SQLOS](#) (NOTE: this video is fairly basic)

- (2) Review the slides + the whiteboard + demo scripts for this module

Misc / Related Topics

These topics were part of our original MCM Training. You should consider watching these and becoming familiar with the topics/features covered and discussed here.

- (1) Go through the HOL lab from your DVD: **SQL Server 2008: Understanding and Using DMVs** (note: this is fairly basic but good for quick review)
- (2) Go through the HOL lab from your DVD: **SQL Server 2008: Command-line Management with SQLCMD** (note: again, fairly basic but might show you something about which you were unaware for SQLCMD)

Week 8

IE2: Module 5: Waits and Queues

(3) Review the MCM Online Videos for recap

- [Waits and Queues](#)
- [Waits and Queues - Demo](#)

(4) Review the slides + the whiteboard + demo scripts for this module

(5) Review the whitepaper: [Performance Tuning Using Waits and Queues](#)

IE2: Module 6: Extended Events

(6) Review the MCM Online Videos for recap

- [Extended Events](#)
- [Extended Events - Demo](#)

(7) Review the slides + the whiteboard + demo scripts for this module

Week 9

IE2: Module 7: Query Plan Analysis

(1) Review the slides + the whiteboard + demo scripts for this module

IE2: Module 9: Plan Cache & Index Analysis

(2) Review the MCM Online Videos for recap

- [Extended Events](#)
- [Extended Events - Demo](#)

(3) Review the slides + the whiteboard + demo scripts for this module

Week 10

IE2: Module 8: CPU Performance Tuning

- (1) Review the slides + the whiteboard + demo scripts for this module
- (2) Review the MCM Online Videos for recap

- [Resource Governor](#)
- [Resource Governor - Demo](#)

IE2: Module 10: Memory Performance Tuning

- (3) Review the slides + the whiteboard + demo scripts for this module

Week 11

IE2: Module 11: Benchmarking / Baselining

- (1) Review the MCM Online Videos for recap

- [Baselining and Benchmarking](#) (NOTE: this video is very basic)

- (2) Review the slides + the whiteboard + demo scripts for this module

IE2: Module 12: Collecting Performance Data

- (3) Review the slides + the whiteboard + demo scripts for this module

Misc / Related Topics

These topics were part of our original MCM Training. You should consider watching these and becoming familiar with the topics/features covered and discussed here.

- (3) Review the MCM Online Videos for recap

- [Policy-Based Management](#)
- [Policy-Based Management - Demo](#)
- [Multi-Server Management](#)

- [Multi-Server Management - Demo](#)

(4) Go through the HOL lab from your DVD: **Using Performance Data Collection**

(5) Go through the HOL lab from your DVD: **Using Policy-based Management**

Week 12

IE3 online videos

- o [Designing an Availability Strategy](#)
- o [Virtualization](#)
- o [Consolidation](#)
- o [Backup Internals](#)
- o [Backup Internals - Demo](#)
- o [Restore Internals](#)
- o [Restore Internals - Demo](#)
- o [Partial Database Availability and Online Piecemeal Restore](#)
- o [Replication](#)
- o [Replication - Demo](#)
- o [Database Mirroring](#)
- o [Database Mirroring - Demo](#)
- o [Clustering Introduction](#)
- o [Clustering Setup and Troubleshooting](#)
- o [Database Snapshots](#)
- o [Database Snapshots - Demo](#)
- o [Consistency Checking \(Part 1\)](#)
- o [Consistency Checking \(Part 2\)](#)

Week 13

IE4 online videos

- o [Optimizing Procedural Code](#)
- o [Service Broker](#)
- o [SQL Azure](#)
- o [SQL Server Clients](#)
- o [Change Data Capture, Change Tracking and Integrated Fulltext Search](#)
- o [Security: Introduction and Server Setup](#)
- o [Security: Authentication and Authorization](#)
- o [Security: Keys and Encryption](#)
- o [Security: Auditing and Compliance](#)

- o [Security: EXECUTE AS and SQL Injection](#)
- o [XML](#)
- o [SQLCLR](#)
- o [SQL Server Spatial](#)