

SQLskills Immersion Event

IE1: Internals and Performance

Module 0: Overview



SQLskills Immersion Event

- Agenda (see inside front page)
- **Monday: Modules 1-2-3**
 - 6pm – Meet and Greet
- **Tuesday: Modules 3-4-5**
 - Tips for Query Tuning with SQL Sentry Plan Explorer and More!
- **Wednesday: Modules 6-7-8**
 - Free evening!
- **Thursday: Modules 9-10**
 - 6pm – Attendee mini-session free4all!
- **Friday: Modules 11-12**
 - 5:30pm – Class ends!

Exclusive Sponsor




2

© SQLskills, All rights reserved.
<http://www.SQLskills.com>



Author/Instructor: Kimberly L. Tripp

- Consultant/Trainer/Speaker/Writer
- President/Founder, SYSolutions, Inc.
 - e-mail: Kimberly@SQLskills.com
 - blog: <http://www.SQLskills.com/blogs/Kimberly>
 - Twitter: @KimberlyLTripp
- Author/Instructor for SQL Server Immersion Events
- Instructor for multiple rotations of both the SQL MCM & Sharepoint MCM
- Author/Manager of SQL Server 2005 & 2008 Launch Content
- Author/Speaker at Microsoft TechEd, SQL Connections, SQLPASS, ITForum, TechDays
- Author of several SQL Server Whitepapers on MSDN/TechNet: Partitioning, Snapshot Isolation, Manageability, SQLCLR for DBAs
- Author/Presenter for more than 25 online webcasts on MSDN and TechNet
- Co-author MSPress Title: SQL Server 2008 Internals, the SQL Server MVP Project (1 & 2), and SQL Server 2000 High Availability
- Presenter/Technical Manager for SQL Server HOL DVDs
- Owner and Manager of the SQLIntersection conference




3

© SQLskills. All rights reserved.
<http://www.SQLskills.com>

Author/Instructor: Paul S. Randal

- Consultant/Trainer/Speaker/Author
- CEO, SQLskills.com
 - Email: Paul@SQLskills.com
 - Blog: <http://www.SQLskills.com/blogs/Paul>
 - Twitter: @PaulRandal
- Contributing Editor of TechNet Magazine, author of the SQL Q&A column, articles on DBA topics, multiple whitepapers for Microsoft
- 5 years at DEC responsible for the VMS file-system and chkdsk equivalent
- Almost 9 years as developer/manager in the SQL Storage Engine team through August 2007, ultimately responsible for Core Storage Engine
 - Wrote DBCC component, other Storage Engine code, DBCC CHECKDB/repair for SQL 2005
- Regular presenter at worldwide conferences on HA/DR, administration, performance, and internals
- Course author/instructor for Microsoft Certified Master qualifications
- Wrote and presented SQL Server 2008 training for Microsoft
- Owner and Manager of the SQLIntersection conference



4

© SQLskills. All rights reserved.
<http://www.SQLskills.com>

M1: Database Structures

- Records
- Pages
- Extents
- Allocation bitmaps
- IAM chains and allocation units
- 2008 new structures (and feature overview)
 - Data compression
 - FILESTREAM

M2: Data File Internals/Maintenance

- Physical layout considerations
- Allocation algorithms
- Instant initialization
- Auto-grow
- To shrink or not to shrink?
- Tempdb

M3: Locking and Blocking

- The anatomy of a data modification
- Locking and blocking
 - Granularity
 - Escalation
 - Duration
- Troubleshooting locking behavior
 - Blocking situations
 - Detecting and avoiding
 - Deadlock situations
 - Detecting and avoiding

M4: Versioning

- Understanding isolation levels
- Isolation in SQL Server
 - By default, uses locking
 - Optionally, can use versioning (and locking)
- Controlling isolation levels
- Statement-level read consistency
- Transaction-level read consistency
- Isolation summary
- Overhead/Monitoring

M5: Logging, Recovery, and the Transaction Log

- Transaction log architecture
- Log records
- Checkpoints and recovery
- Transaction log operations
- Recovery models
- Log file provisioning and maintenance

M6: Table Design

- Data type best practices
- Understanding row width (vertical partitioning)
- Application inconsistencies in types

M7: Index Internals

- Index concepts
- Table structure
- Index internals
 - Heaps
 - Why cluster
 - Table usage
 - Employee table case study
- Clustering key columns in nonclustered indexes
- Summary
 - What do we know?
 - What should we do?
 - Suggestions for the clustering key!

M8: Index Fragmentation

- Data access methods
- What is fragmentation?
- How does fragmentation happen?
- Detecting, mitigating, removing fragmentation
- 2010 PASS presentation on GUIDs
 - **Tales from the Trenches: GUIDs - Use, Abuse and How to Move Forward**
<mms://passmedia.sqlpass.org/SummitOnDemand/2010/Top10/AD372S.wmv>

M9: Internals and Data Access

- **Data access patterns**
- **Covering**
 - Understanding selectivity
 - Understanding the “tipping point”
- **What methods exist for covering?**
 - Nonclustered indexes (all releases)
 - Using indexed views (SQL Server 2000+)
 - Using INCLUDE (SQL Server 2005+)
 - Using filtered indexes (SQL Server 2008+)
 - Using filtered statistics (SQL Server 2008+)
- **Too many cooks in the kitchen...**
- **Index consolidation**

M10: Statistics

- **Cost-based optimization**
- **Data access patterns**
- **Statistics**
 - What do they look like?
 - What are they telling us?
 - How do you see them
 - When/how do they get created
 - When/how do they get updated
- **Problems/solutions with statistics**
 - The histogram
 - Filtered stats
 - Uneven distribution

M11: Indexing Strategies

- Indexing for performance
- Indexing for AND
- Indexing for OR
- Indexing for joins
- Indexing for aggregates
- Indexed views (for aggregates)

M12: Partitioning

- Horizontal partitioning strategies
 - Partitioned views
 - Partitioned tables
- Implementing the sliding window scenario
- Partitioning design techniques combined
- The case for columnstore indexes