

SPR201
**The Foundations of a Healthy
SQL Server Database**
Links and Hidden Slides




Kimberly L. Tripp & Paul S. Randal
SQLskills.com

Kimberly@SQLskills.com

Paul@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 and SQL Server Magazines, author of the SQL Q&A columns, articles on DBA topics, multiple 2008 whitepapers
- 5 years at DEC responsible for the VMS file-system and check/repair
- 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
- Regular presenter at worldwide TechEds, SQL PASS and other conferences on disaster recovery, HA, maintenance, and internals
- Course author/instructor for Microsoft Certified Master qualifications
- Wrote + presented SQL Server 2008 training for Microsoft
- Co-Chair of the SQL Connections conference






SQL Server
CONNECTIONS

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
- Writer/Editor for SQL Magazine www.sqlmag.com
- Author/Instructor for SQL Server Immersion Events
- Instructor for multiple rotations of both the SQL MCM and Sharepoint MCM
- Author/Manager of SQL Server 2005 Launch Content, co-author/Manager for SQL Server 2008 Jumpstart Content, Author/Speaker at TechEd, SQL Connections, SQLPASS, ITForum, Conference Co-chair for SQL Connections
- 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 (two series and other individual webcasts)
- 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
- We love this stuff... feel free to ask questions!



LOB Data

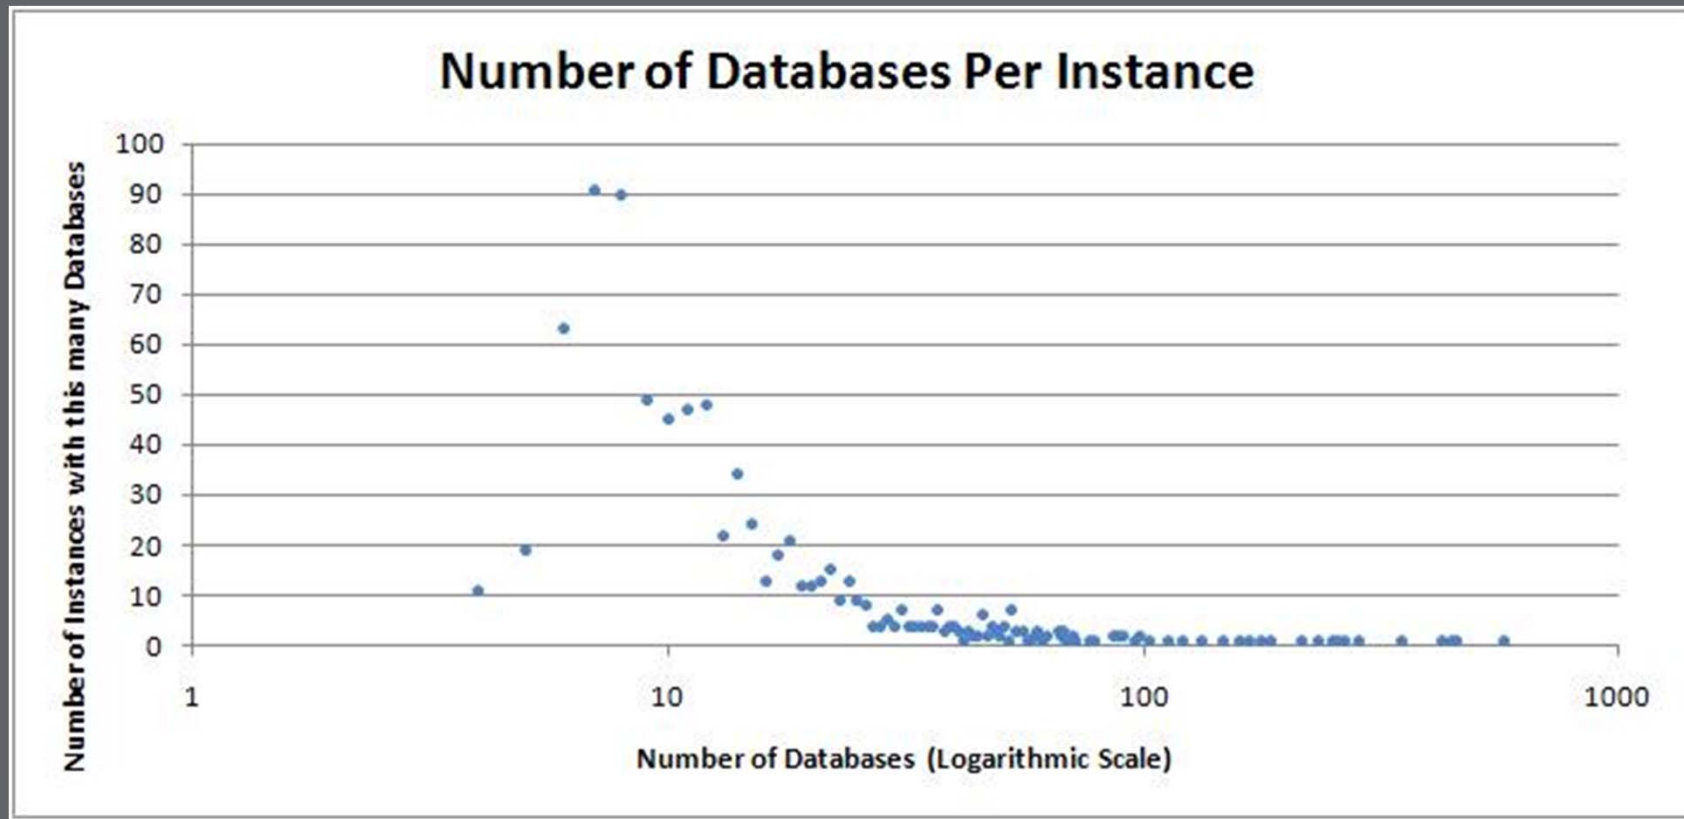
Changing the Storage Setting

- New and legacy types differ for default on/off-row storage
 - Legacy types (n/text, image) off-row by default
 - New types (n/varchar(max), varbinary(max), XML) on-row by default
- For legacy LOB data types:
 - Use the 'text in row' table option
 - `sp_tableoption 'ProdTable','text in row','ON'`
 - Defaults the in-row limit to 256 bytes
 - `sp_tableoption 'ProdTable','text in row','1000'`
 - Beware! Turning the option off is an immediate size-of-data operation
- For new LOB data types:
 - Use the 'large value types out of row' option
 - `sp_tableoption N'MyTable', 'large value types out of row', 'ON'`
 - `sp_tableoption N'MyTable', 'large value types out of row', 'OFF'`
 - Sets the in-row limit to 8000 bytes
 - Existing values are migrated on their next change

For more info









- **Blog posts on internals:**
 - [Inside the Storage Engine: Anatomy of a record](http://www.sqlskills.com/BLOGS/PAUL/post/Inside-the-Storage-Engine-Anatomy-of-a-record.aspx)
(http://www.sqlskills.com/BLOGS/PAUL/post/Inside-the-Storage-Engine-Anatomy-of-a-record.aspx)
 - [Inside the Storage Engine: Anatomy of a page](http://www.sqlskills.com/BLOGS/PAUL/post/Inside-the-Storage-Engine-Anatomy-of-a-page.aspx)
(http://www.sqlskills.com/BLOGS/PAUL/post/Inside-the-Storage-Engine-Anatomy-of-a-page.aspx)
 - [Inside the Storage Engine: Anatomy of an extent](http://www.sqlskills.com/BLOGS/PAUL/post/Inside-the-Storage-Engine-Anatomy-of-an-extent.aspx)
(http://www.sqlskills.com/BLOGS/PAUL/post/Inside-the-Storage-Engine-Anatomy-of-an-extent.aspx)
 - [Inside the Storage Engine: GAM, SGAM, PFS, and Other Allocation Maps](http://www.sqlskills.com/BLOGS/PAUL/post/Inside-The-Storage-Engine-GAM-SGAM-PFS-and-other-allocation-maps.aspx)
(http://www.sqlskills.com/BLOGS/PAUL/post/Inside-The-Storage-Engine-GAM-SGAM-PFS-and-other-allocation-maps.aspx)
 - [Inside the Storage Engine: IAM pages, IAM chains, and allocation units](http://www.sqlskills.com/BLOGS/PAUL/post/Inside-the-Storage-Engine-IAM-pages-IAM-chains-and-allocation-units.aspx)
(http://www.sqlskills.com/BLOGS/PAUL/post/Inside-the-Storage-Engine-IAM-pages-IAM-chains-and-allocation-units.aspx)

Databases per Instance



- Source: <http://www.sqlskills.com/BLOGS/PAUL/post/Log-file-configuration-metrics-for-17000-databases.aspx>









Data File Layout Survey: <10GB

What's the physical layout of your database, and why? (Databases under 10GB)			
Single filegroup, single file		71%	259
Single filegroup, multiple files on a single drive		5%	18
Single filegroup, multiple files spread over multiple drives/arrays/LUNs		13%	47
Multiple filegroups, one per drive/array/LUN, for performance		3%	11
Multiple filegroups, one per drive/array/LUN, for recoverability		0%	1
Multiple filegroups, one per drive/array/LUN, for manageability		1%	3
Multiple filegroups, one per drive/array/LUN, for multiple reasons		3%	11
Multiple filegroups, all on same drive/array/LUN, for any reason		4%	14
			Total: 364 responses

- Source: <http://www.sqlskills.com/BLOGS/PAUL/post/physical-database-layout-vs-database-size.aspx>









Data File Layout Survey: 10-100GB

What's the physical layout of your database, and why? (Databases between 10GB and 100GB)

Single filegroup, single file		43%	135
Single filegroup, multiple files on a single drive		9%	27
Single filegroup, multiple files spread over multiple drives/arrays/LUNs		19%	58
Multiple filegroups, one per drive/array/LUN, for performance		11%	34
Multiple filegroups, one per drive/array/LUN, for recoverability		0%	1
Multiple filegroups, one per drive/array/LUN, for manageability		2%	5
Multiple filegroups, one per drive/array/LUN, for multiple reasons		11%	34
Multiple filegroups, all on same drive/array/LUN, for any reason		5%	17
		Total: 311 responses	









- Source: <http://www.sqlskills.com/BLOGS/PAUL/post/physical-database-layout-vs-database-size.aspx>

Data File Layout Survey: 100GB-1TB

What's the physical layout of your database, and why? (Databases between 100GB and 1TB)			
Single filegroup, single file		20%	45
Single filegroup, multiple files on a single drive		5%	11
Single filegroup, multiple files spread over multiple drives/arrays/LUNs		17%	39
Multiple filegroups, one per drive/array/LUN, for performance		21%	48
Multiple filegroups, one per drive/array/LUN, for recoverability		1%	3
Multiple filegroups, one per drive/array/LUN, for manageability		1%	2
Multiple filegroups, one per drive/array/LUN, for multiple reasons		25%	57
Multiple filegroups, all on same drive/array/LUN, for any reason		11%	25
		Total: 230 responses	

- Source: <http://www.sqlskills.com/BLOGS/PAUL/post/physical-database-layout-vs-database-size.aspx>









Data File Layout Survey: 1TB+

What's the physical layout of your database, and why? (Databases over 1TB)			
Single filegroup, single file		11%	16
Single filegroup, multiple files on a single drive		2%	3
Single filegroup, multiple files spread over multiple drives/arrays/LUNs		8%	12
Multiple filegroups, one per drive/array/LUN, for performance		11%	17
Multiple filegroups, one per drive/array/LUN, for recoverability		2%	3
Multiple filegroups, one per drive/array/LUN, for manageability		3%	4
Multiple filegroups, one per drive/array/LUN, for multiple reasons		52%	78
Multiple filegroups, all on same drive/array/LUN, for any reason		11%	16
		Total: 149 responses	

- Source: <http://www.sqlskills.com/BLOGS/PAUL/post/physical-database-layout-vs-database-size.aspx>

Log Size Management Survey

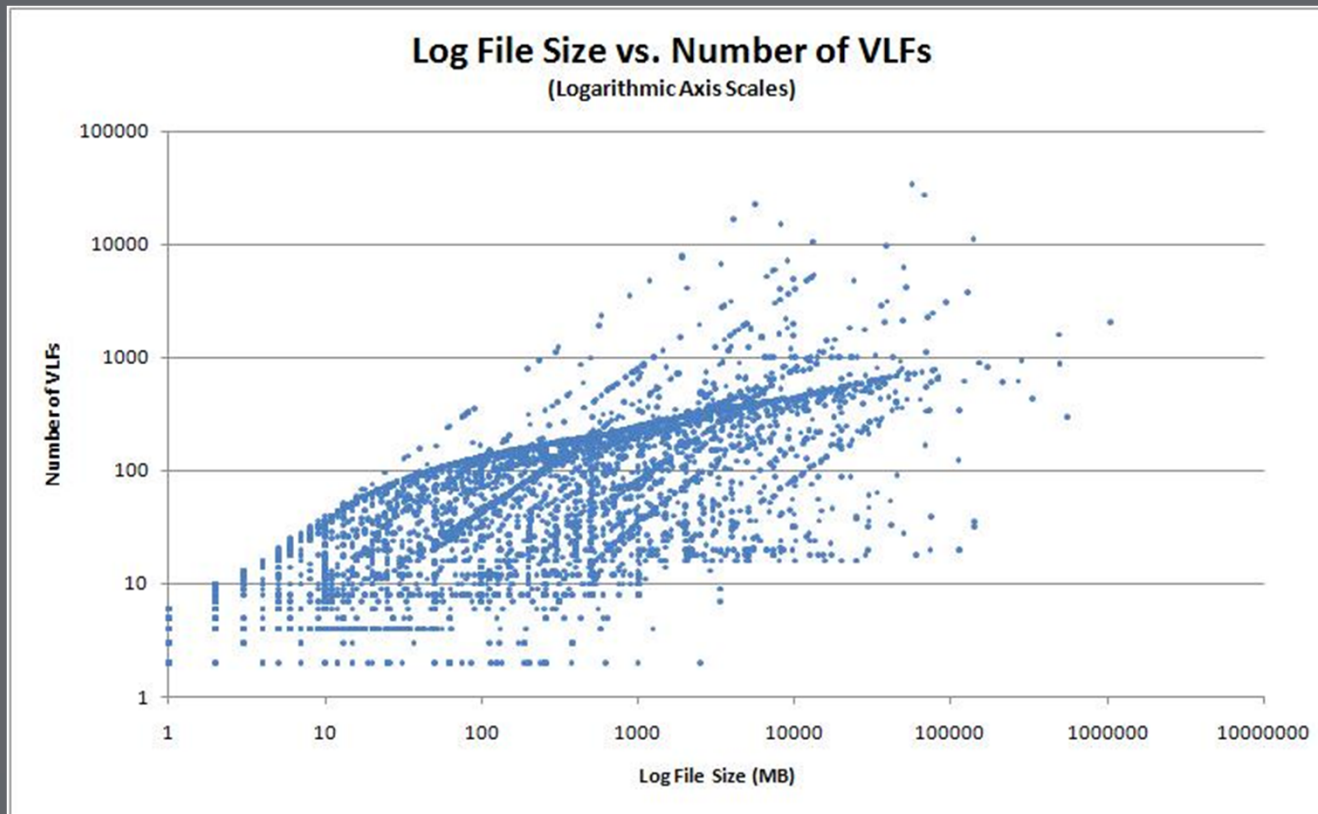
How do *you* manage the size of your transaction log?

Take regular log backups		84%	149
Use BACKUP LOG WITH NO_LOG/TRUNCATE_ONLY regularly		5%	9
Use BACKUP LOG WITH NO_LOG/TRUNCATE_ONLY when it fills up		1%	2
Run in the SIMPLE recovery model all the time		7%	13
Switch to SIMPLE when it fills up, then back to FULL		0%	0
Switch to SIMPLE when it fills up, shrink the log, then switch back to FULL		2%	3
Shutdown SQL Server and delete the transaction log file(s)		0%	0
Add more transaction log files		1%	1

Total: 177 responses

- <http://www.sqlskills.com/BLOGS/PAUL/post/Importance-of-proper-transaction-log-size-management.aspx>

Log File Size vs. VLFs



- Source: <http://www.sqlskills.com/BLOGS/PAUL/post/Log-file-configuration-metrics-for-17000-databases.aspx>

For more info

- Blog posts on the transaction log:

- [8 Steps to better Transaction Log throughput](#)

- (<http://www.sqlskills.com/BLOGS/KIMBERLY/post/8-Steps-to-better-Transaction-Log-throughput.aspx>)

- [Transaction Log VLFs - too many or too few?](#)

- (<http://www.sqlskills.com/BLOGS/KIMBERLY/post/Transaction-Log-VLFs-too-many-or-too-few.aspx>)

- [Database Maintenance Best Practices Part III - TLog Maintenance](#)

- (<http://www.sqlskills.com/BLOGS/KIMBERLY/post/Database-Maintenance-Best-Practices-Part-III-Transaction-Log-Maintenance.aspx>)

- [Understanding backups and log-related Trace Flags in SQL 200x](#)

- (<http://www.sqlskills.com/BLOGS/KIMBERLY/post/Understanding-the-Transaction-log.aspx>)

- [Misconceptions around the log and log backups](#)

- (<http://www.sqlskills.com/BLOGS/PAUL/post/Misconceptions-around-the-log-and-log-backups-how-to-convince-yourself.aspx>)

- TechNet Magazine article

- [Understanding Logging and Recovery in SQL Server](#) (February 2009)

- (<http://technet.microsoft.com/en-us/magazine/2009.02.logging.aspx>)

For more info

- **Blog posts on tempdb:**

- [Understanding TempDB, table variables v. temp tables and Improving throughput for TempDB](http://www.sqlskills.com/BLOGS/KIMBERLY/post/Understanding-TempDB-table-variables-v-temp-tables-and-Improving-throughput-for-TempDB.aspx)

(<http://www.sqlskills.com/BLOGS/KIMBERLY/post/Understanding-TempDB-table-variables-v-temp-tables-and-Improving-throughput-for-TempDB.aspx>)

- [Changing Database Collation and dealing with TempDB Objects](http://www.sqlskills.com/blogs/kimberly/post/Changing-Database-Collation-and-dealing-with-TempDB-Objects.aspx)

(<http://www.sqlskills.com/blogs/kimberly/post/Changing-Database-Collation-and-dealing-with-TempDB-Objects.aspx>)

- [Comprehensive tempdb blog post series](http://www.sqlskills.com/BLOGS/PAUL/post/Comprehensive-tempdb-blog-post-series.aspx)

(<http://www.sqlskills.com/BLOGS/PAUL/post/Comprehensive-tempdb-blog-post-series.aspx>)

- [Misconceptions around TF 1118](http://www.sqlskills.com/BLOGS/PAUL/post/Misconceptions-around-TF-1118.aspx)

(<http://www.sqlskills.com/BLOGS/PAUL/post/Misconceptions-around-TF-1118.aspx>)

- **Whitepapers on tempdb:**

- [Working with tempdb in SQL Server 2005](http://www.microsoft.com/technet/prodtechnol/sql/2005/workingwithtempdb.msp)

(<http://www.microsoft.com/technet/prodtechnol/sql/2005/workingwithtempdb.msp>)

SQLConnections Sessions

For more information on Indexes

- For more details, see the following conference sessions:
 - SQL207: Precarious? Nefarious? Strategies that Work for Nonclustered Indexes!
 - SQL313: Filtered Indexes and Filtered Statistics: The Good, the Bad and the Ugly
 - SQL310: Taking the Sting Out of Statistics
 - SQL225: Follow the Rabbit: Wrap-up Q&A
- **And, there are quite a few other presenters delivering performance-related sessions!!**

For more info

- **Blog categories (each with multiple posts):**
 - [Indexes](http://www.sqlskills.com/blogs/kimberly/category/Indexes.aspx)
(http://www.sqlskills.com/blogs/kimberly/category/Indexes.aspx)
 - [Fillfactor](http://www.sqlskills.com/blogs/kimberly/category/FILLFACTOR.aspx)
(http://www.sqlskills.com/blogs/kimberly/category/FILLFACTOR.aspx)
 - [Indexes from Every Angle](http://www.sqlskills.com/blogs/paul/category/Indexes-From-Every-Angle.aspx)
(http://www.sqlskills.com/blogs/paul/category/Indexes-From-Every-Angle.aspx)
- **Whitepapers on indexes:**
 - [Online Indexing Operations in SQL Server 2005](http://www.microsoft.com/technet/prodtechnol/sql/2005/onlineindex.msp)
(http://www.microsoft.com/technet/prodtechnol/sql/2005/onlineindex.msp)
 - [Improving Performance with SQL Server 2005 Indexed Views](http://www.microsoft.com/technet/prodtechnol/sql/2005/impprfiv.msp)
(http://www.microsoft.com/technet/prodtechnol/sql/2005/impprfiv.msp)

sp_createstats

```
sp_createstats @indexonly = 'indexonly',  
              @fullscan = 'fullscan',  
              @norecompute = 'norecompute'
```

- **@indexonly**: only create statistics for the secondary columns of indexes. This can help to make non-clustered indexes more useful.
- **@fullscan**: requires more time but will create more accurate statistics. If off-hours this is a good idea.
- **@norecompute**: the statistics will NOT get automatically updated as distribution of data changes. Generally, not recommended.
- **Recommendation**:
 `sp_createstats 'indexonly', 'fullscan'`

For more info

- Blog posts on statistics:

- [Auto update statistics and auto create statistics - should you leave them on and/or turn them on??](http://www.sqlskills.com/blogs/kimberly/post/Auto-update-statistics-and-auto-create-statistics-should-you-leave-them-on-and-or-turn-them-on??)

(<http://www.sqlskills.com/blogs/kimberly/post/Auto-update-statistics-and-auto-create-statistics-should-you-leave-them-on-and-or-turn-them-on.aspx>)

- [Index rebuilds depend on stats, which are updated by index rebuilds?!?](http://www.sqlskills.com/blogs/paul/post/Index-rebuilds-depend-on-stats-which-are-updated-by-index-rebuilds?!?)

(<http://www.sqlskills.com/blogs/paul/post/Index-rebuilds-depend-on-stats-which-are-updated-by-index-rebuilds!.aspx>)


- Whitepaper on statistics:

- [Statistics Used by the Query Optimizer in Microsoft SQL Server 2005](http://www.microsoft.com/technet/prodtechnol/sql/2005/qrystats.mspx)

(<http://www.microsoft.com/technet/prodtechnol/sql/2005/qrystats.mspx>)

- Chapter of interest:


- Chapter 6: Index Internals and Management
(Microsoft® SQL Server® 2008 Internals)



L Server
CONNECTIONS

For more info









- Blog categories of interest:
 - [Fragmentation](http://www.sqlskills.com/BLOGS/PAUL/category/Fragmentation.aspx)
(<http://www.sqlskills.com/BLOGS/PAUL/category/Fragmentation.aspx>)
- Webcasts on Index Fragmentation:
 - MSDN Webcast Series: [Part 5: New Features in Indexing and Index Maintenance Best Practices \(Level 200\)](http://msevents.microsoft.com/CUI/WebCastEventDetails.aspx?EventID=1032278599)
(<http://msevents.microsoft.com/CUI/WebCastEventDetails.aspx?EventID=1032278599>)
 - [TechEd 2005: Index Defragmentation with SQL Server 2005 \(Level 300\)](http://www.microsoft.com/emea/spotlight/sessionh.aspx?videoid=30)
(<http://www.microsoft.com/emea/spotlight/sessionh.aspx?videoid=30>)
- Whitepaper on Fragmentation (still relevant for concepts even though written for SQL Server 2000):
 - [Microsoft SQL Server 2000 Index Defragmentation Best Practices](http://www.microsoft.com/technet/prodtechnol/sql/2000/maintain/ss2kidbp.mspx)
(<http://www.microsoft.com/technet/prodtechnol/sql/2000/maintain/ss2kidbp.mspx>)
- Chapter of interest:
 - Chapter 6: Index Internals and Management (Microsoft® SQL Server® 2008 Internals)



SQL Server
CONNECTIONS

Consistency Checking Survey (1)

What method do you use to run consistency checks on your production database (s)?












Run DBCC CHECKDB with no options on the production database		61%	41
Run DBCC CHECKDB WITH PHYSICAL_ONLY on the production database		10%	7
Run DBCC CHECKDB on a restored backup on another server		15%	10
Run DBCC CHECKDB on a database snapshot on a mirror database		0%	0
Spread consistency checking over several days, some tables per day		1%	1
Spread consistency checking over several days using DBCC CHECKFILEGROUP and partitioning		0%	0
Use BACKUP WITH CHECKSUM to validate page checksums, no DBCCs		0%	0
Don't run consistency checks at all		12%	8

Total: 67 responses

- <http://www.sqlskills.com/BLOGS/PAUL/post/Importance-of-how-you-run-consistency-checks.aspx>

SQL Server
CONNECTIONS


Consistency Checking Survey (2)

How often do you run consistency checks? (regardless of [^] how [^] you run them)			
What are consistency checks?		9%	25
Never		5%	14
Only when corruption is detected some other way		8%	22
Only during an event like an upgrade or migration		1%	2
Regularly, but less than monthly		3%	8
Monthly		5%	14
Weekly		37%	103
Daily		25%	69
More frequently than daily		1%	3
Only after performing a restore, or after a failover		1%	3
Other? Enter here...		5%	13
		Total: 276 responses	

- <http://www.sqlskills.com/BLOGS/PAUL/post/Importance-of-running-regular-consistency-checks.aspx>

For more info

- Category on corruption and CHECKDB (*many* posts!):
 - [CHECKDB From Every Angle](#)
(<http://www.sqlskills.com/blogs/paul/category/CHECKDB-From-Every-Angle.aspx>)
 - Specific post of interest: CHECKDB From Every Angle: Complete description of all CHECKDB stages
 - [Page Checksums](#)
(<http://www.sqlskills.com/BLOGS/PAUL/category/Page-Checksums.aspx>)
 - Specific post of interest: How to tell if the IO subsystem is causing corruptions?
- Chapter of interest:
 - Chapter 11: DBCC Internals
(Microsoft® SQL Server® 2008 Internals)



For more info

- **Blog posts on alerts:**

- [Easy monitoring of high-severity errors: create Agent alerts](http://www.sqlskills.com/BLOGS/PAUL/post/Easy-monitoring-of-high-severity-errors-create-Agent-alerts.aspx)
(<http://www.sqlskills.com/BLOGS/PAUL/post/Easy-monitoring-of-high-severity-errors-create-Agent-alerts.aspx>)
- [A little-known sign of impending doom: error 825](http://www.sqlskills.com/BLOGS/PAUL/post/A-little-known-sign-of-impending-doom-error-825.aspx)
(<http://www.sqlskills.com/BLOGS/PAUL/post/A-little-known-sign-of-impending-doom-error-825.aspx>)
- [Conference Questions Pot-Pourri #7: How to create Agent alerts](http://www.sqlskills.com/BLOGS/PAUL/post/Conference-Questions-Pot-Pourri-How-to-create-Agent-alerts.aspx)
(<http://www.sqlskills.com/BLOGS/PAUL/post/Conference-Questions-Pot-Pourri-How-to-create-Agent-alerts.aspx>)

For more info

- **Blog posts on backup & restore:**
 - [Database Maintenance Best Practices Part III - Transaction Log Maintenance](http://www.sqlskills.com/BLOGS/KIMBERLY/post/Database-Maintenance-Best-Practices-Part-III-Transaction-Log-Maintenance.aspx)
(<http://www.sqlskills.com/BLOGS/KIMBERLY/post/Database-Maintenance-Best-Practices-Part-III-Transaction-Log-Maintenance.aspx>)
 - [Backup and Restore](http://www.sqlskills.com/BLOGS/PAUL/category/BackupRestore.aspx) (many posts!)
(<http://www.sqlskills.com/BLOGS/PAUL/category/BackupRestore.aspx>)
- **Webcasts on Backup/Restore:**
 - Part 1 of MSDN Webcast Series: [Creating a Recoverable Database \(Level 200\)](#)
 - Part 2 of MSDN Webcast Series: [Creating a Reliable and Automated Backup Strategy \(Level 200\)](#)
 - Both at:
<http://msevents.microsoft.com/CUI/WebCastEventDetails.aspx?EventID=1032278586>
- **TechNet Magazine article**
 - [Understanding SQL Server Backups](http://technet.microsoft.com/en-us/magazine/dd822915.aspx) (July 2009)
(<http://technet.microsoft.com/en-us/magazine/dd822915.aspx>)

For more info

- Check out our website (no advertising! It's only free content and links to more!)
 - www.SQLskills.com
- Check out our links to whitepapers:
 - www.SQLskills.com/whitepapers.asp
- Check out our blogs:
 - www.SQLskills.com/blogs/Kimberly
 - www.SQLskills.com/blogs/Paul
- Check out our online webcasts (6 individual webcasts, MCM webcasts, and 2 multi-part series for a total of 27 webcasts – over 1000 hours of FREE online/downloadable content):
 - www.SQLskills.com/webcasts.asp