

# SQL SERVER MASTER-CLASS

by SARPEDON QUALITY LAB

## Performance Analyse & Tuning in Zeiten von Azure SQL Database und SQL Server 2017

Andreas Wolter

Microsoft Certified Solutions Master: Data Platform  
MCM, MCSM, MCT, MCSE, MCITP:BD, MCITP:DD, MCITP:DA, MCDBA, MCSA



## About: Andreas Wolter



Consultant, Trainer & Speaker  
Microsoft Certified Master SQL Server 2008  
+ Solutions Master Data Platform (SQL Server 2012)

- Datawarehouse & OLTP-System Architecture
- Performance Tuning
- Security

Email: a.wolter AT SarpedonQualityLab.com  
Blog: [www.andreas-wolter.com/en/blog/](http://www.andreas-wolter.com/en/blog/)  
Facebook: [www.facebook.com/SarpedonQualityLab](http://www.facebook.com/SarpedonQualityLab)  
LinkedIn: [www.linkedin.com/in/andreaswolter](http://www.linkedin.com/in/andreaswolter)  
Twitter: @AndreasWolter

# SQL SERVER MASTER-CLASS

by SARPEDON QUALITY LAB



SQL Server

SSAS, SSRS, SSIS





## Agenda

- Tools
- Adaptive Query Processing
- Automatic Tuning



6

## Tools

- DMVs
  - XEvents
  - Dashboards
  - „Intelligent“ Insights
- 
- SQL Trace & Profiler are deprecated since SQL Server 2008 and do not even work on SQL Azure DB at all!

7

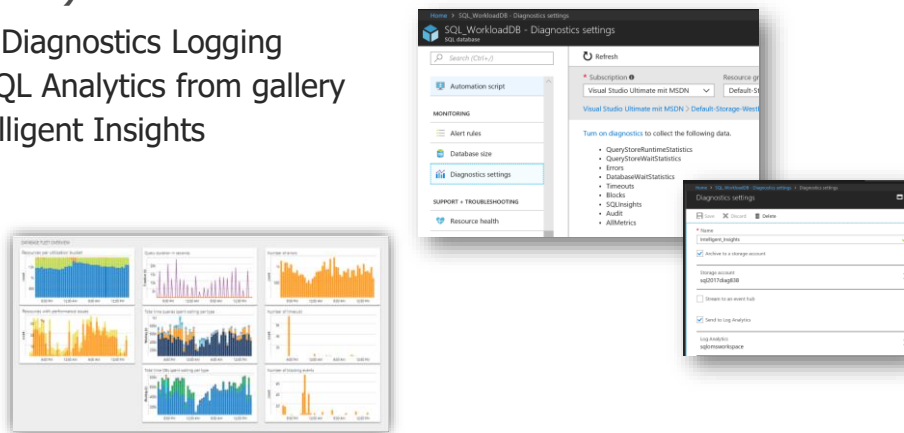
# DMVs & XEvents

- Most DMVs in SQL Azure DB limited to Database scope
  - sys.dm\_db\_wait\_stats
- for checking resource consumption
  - sys.dm\_db\_resource\_stats
  - sys.resource\_stats (master)
- Extended Events similar but less events & actions available
  - Only 3 targets in Azure: ring\_buffer, event\_counter, event\_file (Blob Storage)

	XEvents	SQLTrace Events	Xevent Actions
Microsoft SQL Server 2005 (SP3) - 9.0.4035	0	171	
Microsoft SQL Server 2008 (SP3) - 10.0.5500	243	180	
Microsoft SQL Server 2008 R2 (SP2) - 10.50.4000	262	180	
Microsoft SQL Server 2012 (SP1) - 11.0.3000	625	180	
Microsoft SQL Server 2012 (SP3) - 11.0.6020	644	180	
Microsoft SQL Server 2014 (RTM) - 12.0.2000	870	180	
Microsoft SQL Server 2014 (SP1) - 12.0.4100	872	180	
Microsoft SQL Server 2014 (SP2) - 12.0.5000		180	
Microsoft SQL Server 2016 (SP1) - 13.0.4001	1324	180	53
Microsoft SQL Server 2017 (CU4) - 14.0.3022	1503	180	55
SQL Azure (12.0.2000, 30 Jan 2018)	303	0	22

# Intelligent Insights with Azure SQL Analytics (in Preview)


- Activate Diagnostics Logging
- Install SQL Analytics from gallery
- Use Intelligent Insights



<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-metrics-diag-logging>  
<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-intelligent-insights>

# DTUs

- blended measure of CPU, memory, and read and write rates offered by each performance level
- Based on Azure SQL Database Benchmark (ASDB)
  - mix of basic operations found in all OLTP workloads
    - ▶ Transactions, Workload Mix read/write ratio ~ 2:1 in average
  - throughput achieved for databases running in each performance level
- DTU Calculator
  - <http://dtucalculator.azurewebsites.net/>



# Dienstebenen

	Basic	Standard	Premium	
Zielworkload		Entwicklung und Produktion	Entwicklung und Produktion	Entwicklung und Produktion
Betriebszeit-SLA		99,99 %	99,99 %	99,99 %
Sicherungsaufbewahrung		7 Tage	35 Tage	35 Tage
CPU		Niedrig	Niedrig, Mittel, Hoch	Mittel, Hoch
E/A-Durchsatz (ungefähr)		2,5 IOPS pro DTU	2,5 IOPS pro DTU	48 IOPS pro DTU
E/A-Wartezeit (ungefähr)		5 ms (Lesen), 10 ms (Schreiben)	5 ms (Lesen), 10 ms (Schreiben)	2 ms (Lesen/Schreiben)
Columnstore-Indizierung und In-Memory-OLTP		N/V	N/V	Unterstützt

# Auto-Skaliermöglichkeiten


- Scale Up
  - Azure Automation
    - ▶ <https://gallery.technet.microsoft.com/scriptcenter/Azure-SQL-Database-e957354f>
  - T-SQL
    - ▶ ALTER DATABASE MODIFY EDITION
      - ▶ <https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-azure-sql-database>
- Scale Out
  - Elastic scale
    - ▶ sharding data across multiple databases
- Scale Up & Out
  - Elastic Database Pools
    - ▶ Pool of multiple databases using shared elastic Database Transaction Units (eDTU)
    - ▶ Databases auto-scale within set limits based on eDTUs
      - ▶ <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-pool>
      - ▶ <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-elastic-scale-introduction>

Danke für den Hinweis Kai G.

# AUTOMATIC TUNING & ADAPTABILITY

# Automatic Tuning Features

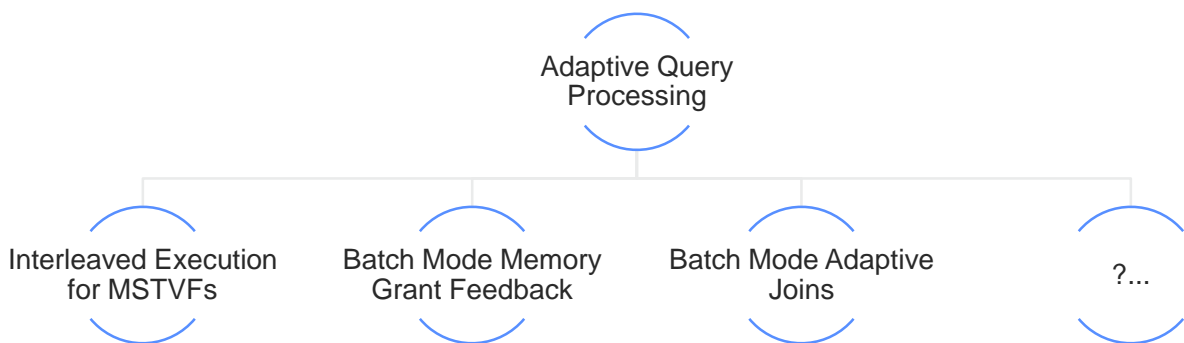
- Automatic Plan Correction
  - requires Query Store
  - „Last Known good“
  - Can also UNDO if this is problematic
  
- Automatic index management
  - in Azure SQL database



- [https://blogs.msdn.microsoft.com/sqlservers\\_torageengine/2017/05/17/automatic-plan-correction-in-sql-server-2017/](https://blogs.msdn.microsoft.com/sqlservers_torageengine/2017/05/17/automatic-plan-correction-in-sql-server-2017/)
- <https://docs.microsoft.com/en-us/sql/relational-databases/automatic-tuning/automatic-tuning>

# Adaptability in SQL Server

- Requires Database Compatibility Level = 140 in SQL Server and Azure SQL Database






## Interleaved Execution for MSTVFs

Problem: Multi-statement table valued functions (MSTVFs) are treated as a black box by QP and we use a fixed optimization guess

Interleaved Execution will materialize and use row counts for MSTVFs

Downstream operations will benefit from the corrected MSTVF cardinality estimate



16


## Batch Mode Memory Grant Feedback (MGF)

Problem: Queries may spill to disk or take too much memory based on poor cardinality estimates

Hash Joins  
-> Spill to Tempdb  
IO\_Completion Wait type

MGF will adjust memory grants based on execution feedback

MGF will remove spills and improve concurrency for **repeating** queries



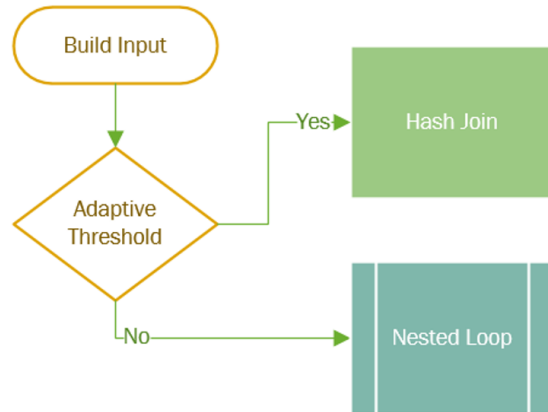
17

# Batch Mode Adaptive Joins (AJ)

Problem: If cardinality estimates are skewed, we may choose an inappropriate join algorithm

AJ will defer the choice of **hash join** or **nested loop** until after the first join input has been scanned

AJ uses nested loop for small inputs, hash joins for large inputs



```

SELECT [fo].[Order Key], [si].[Lead Time Days], [fo].[Quantity]
FROM [Fact].[Order] AS [fo]
INNER JOIN [Dimension].[Stock Item] AS [si]
ON [fo].[Stock Item Key] = [si].[Stock Item Key]
WHERE [fo].[Quantity] = 130;
          
```

Adaptive Join	
Misc	
Actual Execution Mode	Batch
Actual I/O Statistics	
Actual Join Type	NestedLoops
Actual Number of Batches	1
Actual Number of Rows	36
Actual Rebinds	0
Actual Rewinds	0
Actual Time Statistics	
Adaptive Threshold Rows	46.2882
BitmapCreator	true
Defined Values	[[WideWorldImportersDW],[Fact]
Description	Chooses dynamically between ha
Estimated CPU Cost	0.0000018
Estimated Execution Mode	Batch
Estimated I/O Cost	0
Estimated Join Type	HashMatch
Estimated Number of Execution	1
Estimated Number of Rows	87.7599
Estimated Operator Cost	0 (0%)
Estimated Rebinds	0
Estimated Rewinds	0
Estimated Row Size	23 B
Estimated Subtree Cost	0.1002
Hash Keys Build	[WideWorldImportersDW],[Fact]
Hash Keys Probe	[WideWorldImportersDW],[Dime
Is Adaptive	True

0% | Results | Messages | Execution plan

Query 1: Query cost (relative to the batch): 100%

```

SELECT [fo].[Order Key], [si].[Lead Time Days], [fo].[Quantity] FROM [Fact].[Order] AS [fo] INNER JOIN [Dimension].[Stock Item] AS [si] ON [fo].[Stock Item Key] = [si].[Stock Item Key] WHERE [fo].[Quantity] = 130;
          
```





# Vielen Dank

Andreas Wolter



Microsoft  
CERTIFIED  
Solutions Master  
Data Platform

Microsoft  
CERTIFIED  
Master  
SQL Server® 2008

Contact: [a.wolter@Sarpedon.de](mailto:a.wolter@Sarpedon.de)  
LinkedIn: [www.linkedin.com/in/AndreasWolter](http://www.linkedin.com/in/AndreasWolter)  
Twitter: [@AndreasWolter](https://twitter.com/AndreasWolter)

## Seminare zu Performance-Analyse 1.HJ 2018

powered by PASS Deutschland e.V.

### *PASS Essential:* Praktisches Performance Monitoring-Toolset für Microsoft SQL Server

12. April bei Microsoft in Bad  
Homburg

Dozent: Andreas Wolter, MCSM, MCM, MVP

<http://bit.ly/2GGMuew>

Mit Glück noch einer von 2 letzten Plätzen!



## Seminare zu Performance-Analyse 1.HJ 2018

**SentryOne**<sup>™</sup>

**SQL SERVER**   
**MASTER-CLASS**

by **SARPEDON** QUALITY LAB

**PERFORMANCE BOOTCAMP**

<http://bit.ly/SQLMCPSS>

15% Discount-Code für PASS'ler:  
**SQLPASSDEPerformanceMaster**

**SQL Server Master-Class:**  
**Performance Troubleshooting**  
**Bootcamp mit SQLSentry**

- 16.-19. April 2018 in Frankfurt a.M.  
PRAXIS pur: ein eigener Server in Azure mit Workload zum Analysieren ohne Ende!

Dozent: Andreas Wolter, MCSM, MCM, MVP

[www.sql-server-master-class.com/#PSS](http://www.sql-server-master-class.com/#PSS)

## Seminare zu Sicherheit 1.HJ 2018

**SQL SERVER**   
**MASTER-CLASS**

by **SARPEDON** QUALITY LAB

**SECURITY**

<http://bit.ly/MCSec18>

15% Discount-Code für PASS'ler:  
**SQLPASSDESecurityMaster**

**SQL Server Master-Class:**  
**Securityworkshop for SQL Server**  
**Administrators**

- In Planung: Mai/Juni 2018 in Frankfurt a.M.. Warteliste eröffnet. (5 Plätze bereits vergeben)
- 3 Tage, Praktische Lösungen und Deep Dive

Dozent: Andreas Wolter, MCSM, MCM, MVP

[www.sql-server-master-class.com/#SIA](http://www.sql-server-master-class.com/#SIA)



MICROSOFT® CERTIFIED SINCE 2000

## Sarpedon Quality Lab: Ihr Spezialist für Datenbank-Systeme basierend auf SQL Server Technologien



Eine von 2 Firmen in Europa,  
die die höchsten technischen Zertifizierungen von Microsoft für  
sowohl SQL Server 2008 als auch SQL Server 2012 erreicht hat.



Wir unterstützen und setzen unser Know-How gern zu Ihrem Vorteil ein.

Unsere **Dienste** umfassen:

- *SQL Server Supportverträge & Hotline*
- *SQL Server Health Checks*
- *Architektur-Planung, Beratung und Implementierung*
- *Performance Analyse & Tuning*
- *Disaster-Recovery & SLA-Compliance-Checks*
- *Security-Checks*
- *Datenrettung bei Korruption*

• **Training:** 

Fragen Sie uns: [info@Sarpedon.de](mailto:info@Sarpedon.de)

Partner von   
Fragen Sie uns nach speziellen Konditionen und Abo-Modellen



## Artikel zu Tracing & Monitoring von Sarpedon Quality Lab



- Where is that Preemptive Wait coming from? - Database Ownership and Performance: a journey through SQL Server internals with XEvents, Windbg and Wireshark
  - [andreas-wolter.com/en/where-is-that-preemptive-wait-coming-from/](http://andreas-wolter.com/en/where-is-that-preemptive-wait-coming-from/)
- How to import Extended Events session event\_file target and parse deadlock-graph
  - [andreas-wolter.com/en/how-to-import-extended-events-session-event\\_file-target/](http://andreas-wolter.com/en/how-to-import-extended-events-session-event_file-target/)
- Free SQL Deadlock Collector & Parser published at codeplex
  - [andreas-wolter.com/en/free-sql-deadlock-collector-parser-published/](http://andreas-wolter.com/en/free-sql-deadlock-collector-parser-published/)
- Performance/ Management Data Warehouse Data Collector & AlwaysOn Availability Groups
  - [andreas-wolter.com/en/performance-management-data-warehouse-data-collector-alwayson-availability-groups/](http://andreas-wolter.com/en/performance-management-data-warehouse-data-collector-alwayson-availability-groups/)
- Comparing Extended Events vs SQL Trace – or why SQL Trace & Profiler are just a thing of the past :-)
  - [andreas-wolter.com/en/extended-events-vs-sql-trace/](http://andreas-wolter.com/en/extended-events-vs-sql-trace/)
- Performance overhead of tracing with Extended Event targets vs SQL Trace under CPU Load
  - [andreas-wolter.com/en/performance-overhead-of-tracing-with-extended-event-targets-vs-sql-trace-under-cpu-load/](http://andreas-wolter.com/en/performance-overhead-of-tracing-with-extended-event-targets-vs-sql-trace-under-cpu-load/)
- Extended Event File Target size vs SQL Server Trace trace file – a comparison
  - [andreas-wolter.com/en/extended-event-file-target-size-vs-sql-server-trace-trace-file-a-comparison/](http://andreas-wolter.com/en/extended-event-file-target-size-vs-sql-server-trace-trace-file-a-comparison/)
- Tracing Analysis Services (SSAS) with Extended Events – Yes it works and this is how
  - [andreas-wolter.com/en/tracing-analysis-services-ssas-with-extended-events-yes-it-works-and-this-is-how/](http://andreas-wolter.com/en/tracing-analysis-services-ssas-with-extended-events-yes-it-works-and-this-is-how/)
- Locking & READONLY Filegroups vs READONLY Databases
  - [andreas-wolter.com/en/locking-readonly-filegroups-vs-readonly-databases/](http://andreas-wolter.com/en/locking-readonly-filegroups-vs-readonly-databases/)