

Delivering High Availability and Performance with SQL Server 2014

Silviu Niculita
CTO @ RIA Solutions Group
www.riasolutionsgroup.com
ro.linkedin.com/in/silviuniculita

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Huge thanks to our sponsors & partners!



SESSION OBJECTIVES & TAKEAWAYS

Objective

To help you decide if there are any compelling reasons to upgrade to Microsoft's latest SQL Server version, in the context of offering increased availability and enhancing data warehousing scenarios.

Agenda

1. What Has NOT Been Improved
2. What Has Been Improved
3. What Else Has Been Improved

What Has NOT Been Improved

- Integration Services
- Reporting Services
- Replication
- Analysis Services *

* Added support in Power View Reports for connecting to Multidimensional Models;

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Have not looked into

- Master Data Services;
- Data Quality Services;
- StreamInsight;

The majority of the SQL Server team is now working on the SQL Azure and cloud related feature.

A much smaller group is focused strictly on the on-premises offering.

Instead of spreading out a few features across the entire suite, they went deep in a few key/strategic areas.

What Has Been Improved

- Resource Governor Enhancements for Physical IO Control
- Memory-Optimized Tables (Hekaton)
- Columnstore Indexes
- Buffer Pool Extension
- High Availability

What Else Has Been Improved

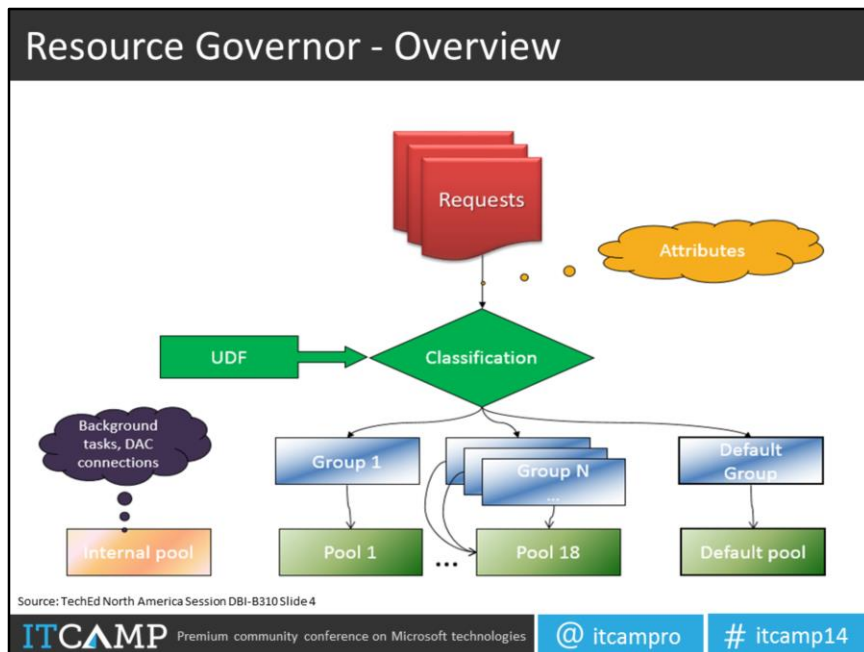
- SQL Server Data Files in Windows Azure
- Backup and Restore Enhancements
- New Design for Cardinality Estimation
- Delayed Durability
- Partition Switching and Indexing
- Incremental Statistics

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14



Added in SQL 2008.

Only Available in Enterprise Edition.

In SQL 2012 processor affinity was added, as well as cap percent.

Also, you can create up to 64 pools.

RG only works on the DB engine, not on SSAS, SSRS, SSIS, etc.

RG only works at the instance level, it is not multi-instance aware.

A workload group serves as a container for session requests that have similar classification criteria.

A resource pool, represents the physical resources of the server. You can think of a pool as a virtual SQL Server instance inside of a SQL Server instance.

Resource Governor – Why IO governance?

- Super-critical for SQL Server performance
 - IO is a member of the Holy Trinity
- Performance predictability
 - Rogue workloads
 - Throw away queries
 - Maintenance operations

Resource Governor – What Can It Control?

- Yes
 - Read I/O
 - Write I/O
 - Physical Reads
 - Data Files
- No
 - Internal Pool
 - Logical Reads
 - Log Files

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Read I/O – physical reads, not logical reads;
Write I/O – Inserts? Updates?

Resource Governor – How?

```
CREATE RESOURCE POOL LimitedIO
[ WITH
    (
        [ MIN_CPU_PERCENT = value ]
        [ [ , ] MAX_CPU_PERCENT = value ]
        [ [ , ] CAP_CPU_PERCENT = value ]
        [ [ , ] AFFINITY { SCHEDULER =
            AUTO | ( <scheduler_range_spec> )
            | NUMANODE = ( <NUMA_node_range_spec> ) } ]
        [ [ , ] MIN_MEMORY_PERCENT = value ]
        [ [ , ] MAX_MEMORY_PERCENT = value ]
        [ [ , ] MIN_IOPS_PER_VOLUME = value ]
        [ [ , ] MAX_IOPS_PER_VOLUME = value ]
    )
]
[;]
```

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Min Value Default = 0


Max Value Default = 2^{31}

Resource Governor Steps in when there is resource contention. If there is no server load, resource governor does not intervene.

If you set MAX values, Resource Governor steps in regardless of resource contention.

IOPS can't be set via the Management Studio UI in 2014, you have to do it via T-SQL.

Memory Optimized Tables



Source: TechEd North America Session DBI-8287 Slide 10

ITCAMP

Premium community conference on Microsoft technologies

@ itcamp

itcamp14

Hekaton adds

- Memory Area for tables and indexes
- FileStream type file group that is tagged as containing memory optimized data.

Index updates are not logged in Hekaton.

We have the option of making a table non-durable, so no logging is done at all.

I also hear a specialized ASP.Net session state provider will be making its way to CodePlex sometime soon. This will be based on the lessons that Microsoft learned when working with Bwin (which has been running 2014 since before Xmas).

Memory Optimized Tables - Limitations

- Optimized for high-throughput OLTP
 - No XML, CLR data types and sqlvariant
- Optimized for in-memory
 - Rows are at most 8060 bytes – no off row data
 - No Large Object (LOB) types like varchar(max)
- Other limitations
 - No FOREIGN KEY and no CHECK constraints
 - No schema changes (ALTER TABLE, add/drop index) – need to drop/recreate table
 - No DML triggers
 - Compression is not supported
 - Running Out of Memory

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Most of these limitations will go away in future versions.

Demo: In Memory OLTP

(script in the notes)

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

```
checkpoint  
go
```

```
dbcc dropcleanbuffers  
go
```

```
-----  
CREATE TABLE DiskConsumers  
(  
    Sequence int not null,  
    FirstName varchar(100),  
    LastName varchar(100)  
    CONSTRAINT PK_DiskSequence PRIMARY KEY NONCLUSTERED (Sequence),  
)  
GO
```

```
-----  
CREATE TABLE MemConsumers  
(  
    Sequence int not null,
```

```

FirstName varchar(100),
LastName varchar(100)
    CONSTRAINT PK_MemSequence PRIMARY KEY NONCLUSTERED HASH (Sequence)
with (bucket_count = 750000),
) WITH (MEMORY_OPTIMIZED = ON)
GO

```

```

CREATE TABLE MemVolatileConsumers
(
    Sequence int not null,
    FirstName varchar(100),
    LastName varchar(100)
    CONSTRAINT PK_VolatileMemSequence PRIMARY KEY NONCLUSTERED HASH
(Sequence) with (bucket_count = 750000),
) WITH (MEMORY_OPTIMIZED = ON, DURABILITY = SCHEMA_ONLY)
GO

```

```

begin tran
declare @counter int = 0
while @counter <= 1025282
begin
insert into DiskConsumers values (@counter, 'John', 'Jones')
set @counter += 1
end
commit
go;

```

```

begin tran
declare @counter int = 0
while @counter <= 1025282
begin
insert into MemConsumers values (@counter, 'John', 'Jones')
set @counter += 1
end
commit
go;

```

```

begin tran
declare @counter int = 0
while @counter <= 1025282
begin
insert into MemVolatileConsumers values (@counter, 'John', 'Jones')

```

```
set @counter += 1
end
commit
go;
```

```
update DiskConsumers set LastName = 'Snow'
```

```
update MemConsumers set LastName = 'Snow'
```

```
update MemVolatileConsumers set LastName = 'Snow'
```

```
delete from DiskConsumers
go
```

```
delete from MemConsumers
go
```


```
delete from MemVolatileConsumers
go
```

```
CREATE PROCEDURE SP_InsertRows
WITH NATIVE_COMPILATION, SCHEMABINDING, EXECUTE AS OWNER
AS BEGIN ATOMIC WITH
(
    TRANSACTION ISOLATION LEVEL = SNAPSHOT, LANGUAGE = 'us_english'
)
declare @counter int = 0
while @counter <= 1025282
begin
insert into dbo.MemConsumers values (@counter, 'John', 'Smith')
set @counter += 1
end
END
GO
```

```
EXECUTE [dbo].[SP_InsertRows]
GO
```

```
select count (*) from MemConsumers
```

In Memory OLTP – Expected Performance Gains



Source: TechEd North America Session DBI-B287 Slide 18

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Your mileage may vary.

Columnstore Indexing

- NonClustered Columnstore Index
 - Unchanged since SQL Server 2012
 - Can index a subset of columns (like frequently used columns)
 - Requires extra storage
 - Not updateable
 - Can be combined with other indexes on the table.
 - Can be configured to use columnstore or columnstore archival compression.

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

According to TDWI:

41% of companies have a data warehouse under 1TB;

In 3 years, only 17% of companies are going to have a data warehouse under 1TB;

In 3 years, over 34% of companies are going to have a data warehouse over 10TB;

A columnstore index is a technology for storing, retrieving and managing data by using a columnar data format, called a columnstore.

A *columnstore* is data that is logically organized as a table with rows and columns, and physically stored in a column-wise data format.

Not updateable - Is updated by rebuilding the index or switching partitions in and out. It is not updateable by using the DML operations such as insert, update, and delete.

Columnstore Indexing

- Clustered Columnstore Index
 - Is updateable
 - Is the primary storage method for the entire table
 - Has no key columns - all columns are included
 - Is the only index on the table - it cannot be combined with any other indexes

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Primary storage method – translates into massive disk space savings;

Both index types can be configured to use columnstore or columnstore archival compression.

Demo: Columnar Indexing

(script in the notes)

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

```
checkpoint  
go
```

```
dbcc dropcleanbuffers  
go
```

```
-----  
set statistics io on  
-----
```

```
select  
count(*),  
marital_status,  
gender  
from  
[RealConsumers]  
where  
collector_stamps = 1  
and collector_dolls = 1  
and collector_figurines = 1
```

```
and state = 'TX'
group by
marital_status,
gender
```

```
select
count(*),
marital_status,
gender
from
[RealConsumersNCCS]
where
collector_stamps = 1
and collector_dolls = 1
and collector_figurines = 1
and state = 'NY'
group by
marital_status,
gender
```

```
update
[RealConsumersNCCS]
set
marital_status = 'G'
where
collector_stamps = 1
and collector_dolls = 1
and collector_figurines = 1
and state = 'NY'
and marital_status is null
```

```
select
count(*),
marital_status,
gender
from
[RealConsumersCCS]
where
collector_stamps = 1
and collector_dolls = 1
and collector_figurines = 1
and state = 'NY'
```

```
group by
marital_status,
gender
```

```
select
count(*),
marital_status,
gender
from
[RealConsumersCCSA]
where
collector_stamps = 1
and collector_dolls = 1
and collector_figurines = 1
and state = 'NY'
group by
marital_status,
gender
```

```
update
[RealConsumersCCS]
set
marital_status = 'S'
where
collector_stamps = 1
and collector_dolls = 1
and collector_figurines = 1
and state = 'NY'
and marital_status is null
```

```
--Notice the deltastore
select
o.name as TableName,
c.state,
c.state_description,
c.total_rows
from
sys.objects o
join sys.column_store_row_groups c on o.object_id = c.object_id
where
o.name = 'RealConsumersCCS'
```

--Do not run, takes FOREVER!!!

```
ALTER INDEX [ClusteredColumnStoreIndex] ON [dbo].[RealConsumersCCS]  
REBUILD PARTITION = ALL WITH (DATA_COMPRESSION = COLUMNSTORE)
```

--Connect to SQL4

```
select  
count(*),  
marital_status,  
gender  
from  
[RealConsumersCCS]  
where  
collector_stamps = 1  
and collector_dolls = 1  
and collector_figurines = 1  
and state = 'NY'  
group by  
marital_status,  
gender
```

Buffer Pool Extension

- Can extend the Buffer Pool to non-volatile storage
- Only “Clean” pages are written to BPE in order to prevent data loss
- Potential Performance Gains – Especially for Read-Intensive OLTP workloads
- No Application Changes, just enable the Feature and sit back
- Simple Syntax

Buffer Pool Extension 2

- Sweet Spot
 - Read Heavy OLTP workloads
 - Indexes in BPE
 - High Throughput SSD Storage (or 15K SAS drives in RAID0 type arrays)
 - Ratio between the size of the physical memory (`max_server_memory`) and the size of the buffer pool extension of 1:16 or less
 - A lower ratio in the range of 1:4 to 1:8 may be optimal
- Limitations
 - Data Warehouse workloads
 - Write Heavy OLTP workloads
- Available on Standard Edition!!!

Demo: Enabling/Disabling BPE

(script in the notes)

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

--Enable Server Configuration Syntax

```
ALTER SERVER CONFIGURATIONSET BUFFER POOL EXTENSION  
ON(FILENAME='C:\BufferPool\SSD.bpe', SIZE = 5 GB);
```

--Check Status

```
SELECT * FROM sys.dm_os_buffer_pool_extension_configuration
```

--Returning cached page count for each database

```
SELECT  
COUNT(*) AS cached_page_count,  
CASE database_id  
    WHEN 32767 THEN 'ResourceDb'  
    ELSE db_name(database_id)  
END  
AS database_name,  
is_in_bpool_extension  
FROM  
sys.dm_os_buffer_descriptors
```

```
GROUP BY
DB_NAME(database_id),
database_id,
is_in_bpool_extension
ORDER BY
database_name DESC
```

```
--Turn off BPE
```

```
ALTER SERVER CONFIGURATIONSET BUFFER POOL EXTENSION OFF;
```

```
select count(*), state, first_name from RealConsumers
where mayonnaise = 1
group by state, first_name order by 1 desc , 2
```

AlwaysOn Technologies

- AlwaysOn Failover Cluster Instances
 - Server Level Protection
 - Also available in Standard Edition (2 nodes maximum)
 - Shared storage is a requirement (SAN, NAS)
- AlwaysOn Availability Group
 - Database-set level protection
 - Directly attached storage is used
 - Single listener endpoint for clients
 - Load balancing read requests possible via DNS round-robin or specialized load balancers
 - Secondary Replicas can be readable
- SQL Server 2014 is the last version of SQL Server that will have database mirroring

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Database mirroring was deprecated in 2012, will no longer be available after 2014.

Both of these were introduced in 2012 and then enhanced in 2014.

AlwaysOn FCI Improvements in 2014

- Support for multi-subnet clustering
- Support for Cluster Shared Volumes
- Improved health monitoring through DMVs
- Improved fail-over scenarios (Failover takes 30s to couple of minutes)
- TempDB on local disk leading to improved performance & better SAN utilization

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

AlwaysOn AG Improvements in 2014

- 8 secondary replicas, up from 4;
- Readable secondaries remain available during “Resolving” state
- Hybrid (on-premise and Azure) deployments are facilitated
- Enhanced Diagnostics through DMVs

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14


Read workloads not impacted during network failures (or primary down, or cluster quorum loss)

Single DNS entry or IP address that clients connect to; (Virtual Network Name)

Still expensive to do Azure Deployments

Memory Intensive VM (8 x 1.6GHz CPU, 56GB RAM) \$3,124.80

Figure: Organization using AlwaysOn AG



Source: Introducing Microsoft SQL Server 2014 Technical Overview, Pg. 5

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Demo: AG Replica in Azure

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

What Else Has Been Improved (time permitting)

- SQL Server Data Files in Windows Azure
- Backup and Restore Enhancements
- New Design for Cardinality Estimation
- Delayed Durability
- Partition Switching and Indexing
- Incremental Statistics

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

SQL Server Data Files in Windows Azure

- Benefits
 - Cost and limitless storage *
 - High availability and disaster recovery
- Limitations
 - FileStream data is not supported, so no In-Memory OLTP
 - Upper limit on individual database data and log files at 1 TB
 - AlwaysOn Availability Groups are supported as long as you do not add new database files to the primary database
 - AlwaysOn Failover Cluster Instances is not supported
 - Not recommendable for on-premises installations;

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Although it is theoretically possible and officially supported, using an on-premises SQL Server 2014 installation and database files in Windows Azure blob storage is not recommended due to high network latency, which would hurt performance.

For this reason, the main target scenario for this white paper is SQL Server 2014 installed in Windows Azure Virtual Machines (IaaS).

Backup and Restore Enhancements

- SQL Server Managed Backups to Microsoft Azure
- Backup to URLs directly in Management Studio
- Encryption for backups
 - The industry standard encryption algorithms that are supported include AES 128, AES 192, AES 256, and Triple DES
 - Encrypted backups are supported in Windows Azure storage or on-premises

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

No longer need to use TDE in order to achieve encryption of database backups.

New Design for Cardinality Estimation

- Re-designed in SQL Server 2014 to improve the quality of query plans, and therefore to improve query performance
- To use it, you have to set DB compatibility level to 120
- Testing is recommended before moving up to the new Cardinality Estimator
- You can coerce use of new Cardinality Estimator in compatibility level 110 databases with trace flag 9481
- Reference: MSDN White Paper [SQL Server 2014's new cardinality estimator](#)

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Note: Compatibility level 90 is no longer supported in 2014. Level 90 is SQL Server 2005.

Run your workload with the new cardinality estimator, and then troubleshoot any new performance issues in the same manner you do today.

If your workload is running with database compatibility level 110 and you want to test or run a specific query with the new cardinality estimator, you can run the query with trace flag 9481 to use version 120 (the new version) of the cardinality estimator.

Delayed Durability

- Ability to reduce latency by designating some or all transactions as delayed durable
- A delayed durable transaction returns control to the client before the transaction log record is written to disk
- Durability can be controlled at the database level, COMMIT level, or ATOMIC block level

Delayed Durability 2

- When to use delayed transaction durability
 - You can tolerate some data loss
 - You are experiencing a bottleneck on transaction log writes
 - Your workloads have a high contention rate
- When can I lose data?
 - Catastrophic events you will lose the data for all committed transactions that have not been saved to disk
 - SQL Server shutdown and restart

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

For delayed durability, there is no difference between an unexpected shutdown and an expected shutdown/restart of SQL Server.

Like catastrophic events, you should plan for data loss. In a planned shutdown/restart some transactions that have not been written to disk may first be saved to disk, but you should not plan on it. Plan as though a shutdown/restart, whether planned or unplanned, loses the data the same as a catastrophic event.

Partition Switching and Indexing

- In 2012 you could rebuild the index for an entire table online or rebuild the index at a partition level offline
- In 2014 we can rebuild an index (or indexes) for a table at a partition level granularity online
- Benefits
 - Table will be accessible for DML and query operations during index rebuild
 - Resource saving – (CPU, memory and disk space)
 - Log space usage reduced

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Resource saving – (CPU, memory and disk space) by rebuilding only a single partition online instead of rebuilding the entire index online.

Incremental Statistics

- Updating statistics on large tables translated into a table scan (also true when using partitioning)
- Benefits
 - You can update only the partition or partitions that you need
 - The information on these partitions will be merged with the existing information to create the final statistics object
 - Percentage of data changes required to trigger an automatic update of statistics is at the partition level so only 20% of rows changed per partition are required

Resources

- [SQL Server Sessions @ TechEd North America 2014](#)
- [Free eBook: Introducing Microsoft SQL Server 2014](#)
- [Updating your Database Management Skills to SQL Server 2014 - MVA Course](#)
- [Microsoft SQL Server 2014 Product Guide](#)

ITCAMP

Premium community conference on Microsoft technologies

@ itcampro

itcamp14

THANK YOU!

Q & A

silviu@niculita.ro

ITCAMP Premium community conference on Microsoft technologies

@ itcampro

itcamp14

Huge thanks to our sponsors & partners!

