

FUTURE DECODED

6-7 OTT '16 / MILANO

IN PARTNERSHIP WITH:



CommunityDays.it

www.futuredecoded.it



#FutureDecoded

Ottimizzare le prestazioni di SQL Server con la configurazione

Sergio Govoni

Microsoft Data Platform MVP

@segovoni

http://sqlblog.com/blogs/sergio_govoni

www.futuredecoded.it

 #FutureDecoded



Sergio Govoni



Microsoft Data Platform MVP




sqlblog.com/blogs/sergio_govoni



@segovoni



Vice Presidente



www.manning.com/books/sql-server-mvp-deep-dives-volume-2



segovoni@gmail.com

Agenda

- Common misconfigured options
 - Hypervisor level
 - Balloon Driver
 - OS & Storage level
 - Power Profile
 - Network Card Optimizations
 - Block Alignment
 - Instant File Initialization
 - Instance level
 - Memory settings
 - Tempdb data files
 - Cost threshold of parallelism



Hypervisor

Memory Ballooning

www.futuredecoded.it

 #FutureDecoded


Ballon Driver

- Balloon driver is loaded into the guest operating system as a “pseudo-device” driver
- No external interfaces to the guest operating system
- It communicates with the hypervisor through a private channel

Memory Ballooning

- Hypervisor is not aware of the guest OS memory management (used or free)
- When memory pressure occurs on the hypervisor, it requires reclaiming some memory from VMs, it will utilize the "Balloon Driver"
- Memory is paged to disk in guest OS

Memory Ballooning



▼ *Memory			
RAM (*)	16	GB	
Reservation (*)	16	GB	
	<input checked="" type="checkbox"/> Reserve all guest memory (All locked)		
Limit	Unlimited	MB	
Shares	Normal	163840	
Memory Hot Plug	<input type="checkbox"/> Enable		

Memory Ballooning

- If your SQL Server is designed for performance, we have to eliminate any chance of paging memory to disk!
- Setting Memory Reservation
 - Hypervisor won't be able to reclaim memory from the guest OS
 - [SQL Server on VMware: Best Practices Guide](#)
 - [SQL Server on Hyper-V: Best Practices Guide](#)
- Leave the Balloon Driver installed to manage memory pressure and prevent loss of service

A nighttime cityscape featuring several tall skyscrapers with illuminated windows. In the foreground, a multi-level highway interchange is visible, with long-exposure light trails from cars creating streaks of white and red light. The overall scene is dark, with the city lights providing the primary illumination.

OS & Storage



Power Profile, Network
Optimizations, Block Alignment, IFI

www.futuredecoded.it

 #FutureDecoded


Power Profile

- “Balanced” or “Power Saved” by default
- “High Performance” is better for SQL Server 😊



Network Card Optimizations

- Speed & Duplex
 - Don't use "Auto Negotiation" (default), set speed explicitly
 - Whole network must supports the speed you have chosen
- Jumbo packets
 - They are extra-size network packets, reduce network protocol overhead
 - Normal: 4096 bytes
 - Jumbo frames: ~9000 bytes
- Receive Side Scaling (RSS)
 - Disabled by default
 - It allows to better distribute workload over CPUs



A nighttime cityscape featuring several tall skyscrapers with illuminated windows. In the foreground, a multi-level highway interchange is visible, with long-exposure light trails from cars creating streaks of white and red. The overall scene is bathed in a cool blue light, with some warmer tones from the city lights and light trails.

DEMO

Network Card Optimizations


www.futuredecoded.it

 #FutureDecoded

Block Alignment

- On particular workload, can improve I/O performance up to 100%
 - RAID Stripe Unit is misaligned with NTFS Cluster Units
 - Accessing an NTFS Cluster Unit needs access to more than one RAID Stripe Unit
- Before Windows Server 2008
 - 63 hidden Sectors (31,5 KB)
- Windows Server 2008 and higher
 - Default offset: 1MB
 - Alignment by default


Misalignment – 4 KB



Block Alignment

- Must be done
 - When the disk is formatted
- OS
 - Windows Server 2003: diskpart.exe
 - Windows Server 2008 (R2) and later: Automatic
- Best practices
 - NTFS cluster size = RAID stripe size
 - RAID stripe unit size = 64KB (1 extent, 8 data pages)
 - $\text{Partition_Offset} / \text{Stripe_Unit_Size}$ must be a whole integer (8, 16, ..)

Disk Alignment Case Study



[Disk Partition Alignment Best Practices for SQL Server](#)


Database File Initialization

- By default, for security reasons, Data and Log files are first initialized by filling the space with zeros when SQL Server performs a
 - CREATE or ALTER DATABASE
 - RESTORE
 - AUTOGROW
- Zero-initialization takes lots of time, especially for large space allocation

Database Instant File Initialization

- If you don't need SQL Server is doing zero-initialization on data file, you can configure the "Performance Volume Maintenance Task" privileges at OS level
- It works only for Data file allocation
- Log file are always zero-initialized, no way around that
- Service restart is required

Database Instant File Initialization



Database Instant File Initialization

> Performance considerations

Performance Test with Zero Initialization

Hardware: Dell Precision 670 Dual Proc (x64) with Dual Core, 4 GB Memory, RAID 1+0 array w/4-142 GB, 15000rpm disks

CREATE DATABASE with 20 GB Data file = **14:02 minutes**
ALTER DATABASE BY 10 GB = **7:01 minutes**
RESTORE 30 GB DATABASE (EMPTY Backup) = **21:07 minutes**
RESTORE 30 GB DATABASE (11GB Backup) = **38:28 minutes**

Performance Test with Instant Initialization

Hardware: Dell Precision 670 Dual Proc (x64) with Dual Core, 4 GB Memory, RAID 1+0 array w/4-142 GB, 15000rpm disks

CREATE DATABASE with 20 GB Data file = **1.3 seconds**
ALTER DATABASE BY 10 GB = **.4 seconds**
RESTORE 30 GB DATABASE (EMPTY Backup) = **5 seconds**
RESTORE 30 GB DATABASE (11GB Backup) = **19:42 minutes**

[Instant Initialization – What, Why and How?](#) (Kimberly Tripp)

Database Instant File Initialization

> Security Considerations from BOL <http://bit.ly/1dK32Dc>

- Deleted disk content is overwritten only as new data is written to the files
 - Deleted content might be accessed by an unauthorized principal
- While the database file is attached to the instance, this information disclosure threat is reduced by the discretionary access control list (DACL)
- When the file is detached or backed up, it may be accessed by a user or service that does not have SE_MANAGE_VOLUME_NAME



Demo

Instant File Initialization

www.futuredecoded.it

 #FutureDecoded



Instance

Memory settings, Tempdb data files, Cost threshold of parallelism

www.futuredecoded.it

 #FutureDecoded

Memory settings

- SQL Server hasn't a memory configuration by default (Min and Max server memory)
- Setting the Maximum Server Memory is a good thing for the OS
 - Max Memory = (Max Physical Memory – 3 GB)
 - Based on the services you have on the server
 - Take care to other services such as SSIS, SSRS, SSAS, etc..

Tempdb Overview

- It stores
 - User objects
 - Work objects (worktable for Sort and Spool, ...)
 - Version Store (Row Versioning)
- It's always recreated after SQL Server restart
- It uses Simple recovery model
- One tempdb for the entire instance = It's a **bottleneck by design!!**

User Objects

- Local temporary tables
 - Prefix "#", Scope limited to the local session
 - Auto dropped after the session is closed
- Global temporary tables
 - Prefix "##", Visible in all sessions
 - Auto dropped after the session is closed
- Table variables
- Tables returned from the "Table Valued Functions"

Latch Contention

- Latch
 - It's a synchronization object used to protect physical pages
 - Every time a page is accessed, a latch is needed
 - Other threads can't access the page in the mean time
- Occurs when many objects (table variables, temp table, ...) are created/destroyed
- Latch Contention on
 - PFS page
 - GAM/SGAM page

Creating a temporary table means..

- Reading the SGAM page (2:1:3) to find an extent with free space
 - An exclusive latch is active during the update
- Reading the PFS page (2:1:1) to find a free page within the extent
 - An exclusive latch is active during the update
- A PAGELATCH wait type occurs
 - 2:1:3 (SGAM page)
 - 2:1:1 (PFS page)

How to reduce Latch Contention

- Multiple tempdb data files
 - From $\frac{1}{4}$ to $\frac{1}{2}$ data files of the CPU cores, including HT cores
 - New objects will be allocated using round-robin algorithm
 - All data files must have the same size!
- Don't use the Microsoft recommendation: one-to-one mapping between data files and CPU cores
 - <http://bit.ly/1f9sQYZ>

A nighttime cityscape featuring several tall skyscrapers with illuminated windows. In the foreground, there is a multi-level highway with long-exposure light trails from cars, showing red and white streaks. A pedestrian bridge with a railing and some greenery is visible on the right side. The overall scene is lit with a mix of blue and white light from the buildings and streetlights.

DEMO

Reduce Latch Contention on tempdb

www.futuredecoded.it

 #FutureDecoded

Cost threshold of parallelism

- Specify the threshold at which SQL Server Query Optimizer creates and runs parallel plans for queries
- Determines which queries are considered "short", so they should be run using "serial" plans
- The default value is 5, queries above that cost are executed with parallelism
 - Is it a good thing or not?

Parallelism

- Parallelism in an OLTP system is not so good!
 - CXPACKET: Parallelism issue through Parallel Execution Plan, threads are not given equal amount of work to do
- Cost threshold for parallelism
 - Set it to a much higher value than the default (5)
 - Usually, values from 10 to 40 are good values
- MAXDOP
 - Set it to 1 **is not the “wizard solution”**
 - Limit it to the max core in the NUMA node [KB 2806535](#)



Demo

Cost threshold of parallelism

www.futuredecoded.it

 #FutureDecoded

Resources

- SQL Server in a virtual environments
 - [SQL Server on VMware: Best Practices Guide](#)
 - [SQL Server on Hyper-V: Best Practices Guide](#)
- Disk Alignment
 - [Disk Partition Alignment Best Practices for SQL Server](#)
- Database Instant File Initialization
 - [What, Why and How?](#)
 - [Misconceptions around IFI](#)



Domande?

Materiale su

<http://www.communitydays.it/>

www.futuredecoded.it

 #FutureDecoded

A nighttime cityscape featuring several tall skyscrapers with illuminated windows. In the foreground, there is a multi-level highway with light trails from cars, and a pedestrian bridge with a railing and some greenery. The overall scene is lit with a mix of blue and white light from the buildings and warmer yellow and red light from the traffic.

Thanks!

Thanks for attending this session and
the entire event!

Have a nice day!

@segovoni

www.futuredecoded.it

 #FutureDecoded