Jeff Mlakar

SQL Saturday #721 – Raleigh 2018

Environmental SQL Server Troubleshooting

Who Am I?

DBA/Developer about 15 years
MS Information Science – University of Pittsburgh
Databases – Security – Privacy
Escalation Engineer – Tech Lead at Varonis 

Twitter: @jmlakar
LinkedIn: https://www.linkedin.com/in/jeffmlakar
Blog: www.MlakarTechTalk.com

Agenda – Troubleshooting

- Environmental?
- SQL Server
- Windows Server
- VMware

SQL Server Configurations

SQL Server Configurations

- SQL Server Configurations
  - Resource Allocation
    - CPU
    - Memory
    - Disk
  - Instance level configurations
  - Database level configurations

SQL Server - CPU

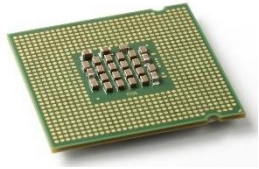

- MAXDOP
  - Max # of processors used for the execution of a query in a parallel plan
  - Determines computing and thread resources

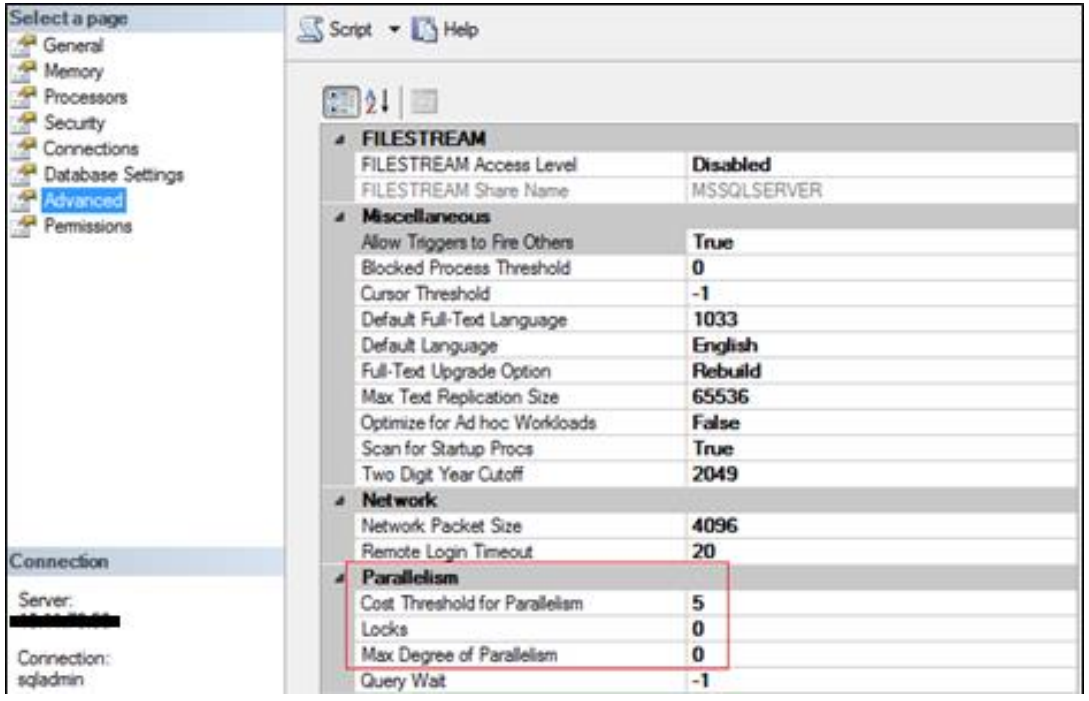

There are exceptions but this is a good starting place:
https://support.microsoft.com/en-us/kb/2806535

SQL Server - CPU

- The default is 0 – i.e. use all cores
- For < 8 logical cores, assign the value to be the number of logical cores
- For >= 8 logical cores, assign the value to be 8
- Exceptions exist – good starting place
  - SharePoint, OLTP vs OLAP
  - Diminishing returns
- How do you know if there are MAXDOP problems?

SQL Server - CPU

- Cost threshold for parallelism
  - Default is 5
  - Based on estimated query cost
  - Suggest 50
  - Legend of CTP
    - https://sqlstudies.com/2017/04/17/what-is-the-cost-in-cost-threshold-for-parallelism/
- Processor and I/O Affinity
  - Controls CPU processors that SQL Server will use
  - We typically want SQL to use all processors

https://blogs.msdn.microsoft.com/psssql/2010/11/19/how-it-works-io-affinity-mask-should-i-use-it/

SQL Server - Memory

- The default memory allocated to a SQL Instance is 0 for min server memory and 2147483647 for max server memory (i.e. all).
- Set the minimum server memory to 0

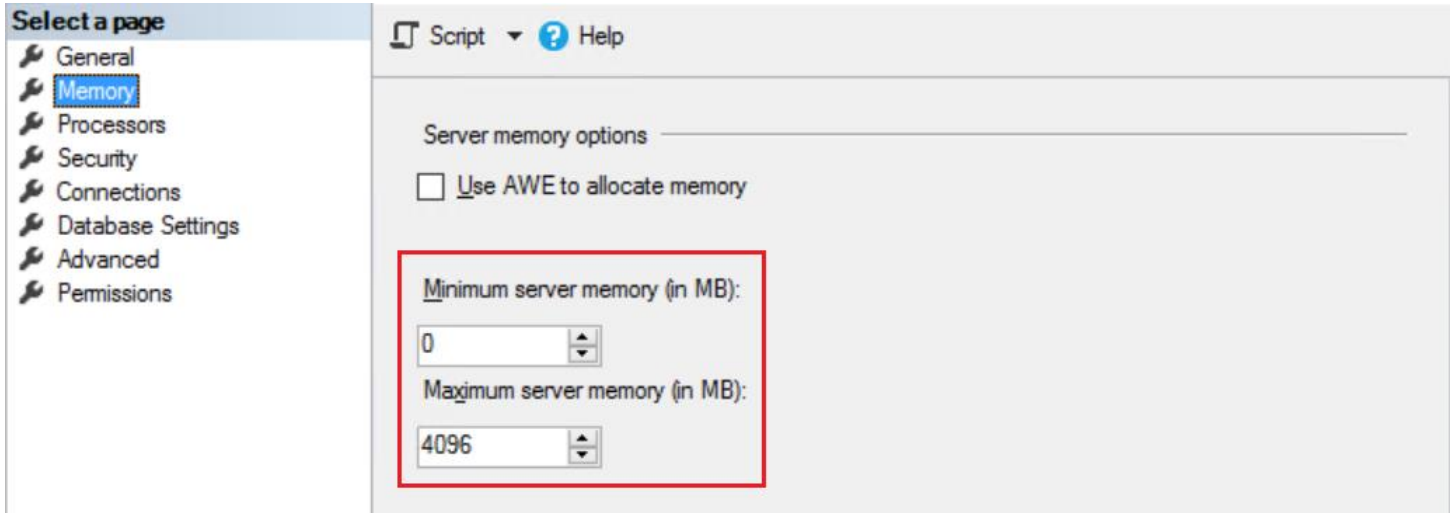

SQL Server - Memory

- Set the maximum server memory depending on what else is running on the machine
- Typically leave Windows Server >= 4GB
- SSRS, SSIS, SSAS should have !< 4GB each
- If other apps on machine add more memory
- Assign the rest to SQL Server

SQL Server - Memory

- Examples:

Total RAM = 16GB	
Component	Memory Allocated (GB)
OS	4
SSRS	4
App	4
SQL	4

Total RAM = 64GB	
Component	Memory Allocated (GB)
OS	4
SSRS	8
App	4
SQL	48

Total RAM = 128GB	
Component	Memory Allocated (GB)
OS	8
SSRS	8
App	4
SQL	108

https://www.sqlskills.com/blogs/jonathan/how-much-memory-does-my-sql-server-actually-need/

SQL Server - Memory

- CLR
  - Starting with SQL Server 2012, CLR allocations are also included in memory limits that are controlled by max server memory (MB) and min server memory (MB).
- Prior it was addressed in the OS memory space
  - https://support.microsoft.com/en-us/kb/2663912

SQL Server - Disk

- Separate DB files on different disks
  - Database files (system / user) – MDF, NDF
  - Transaction logs – LDF
  - Tempdb
  - Backups - BAK
  - Trace files – TRN
  - OS / SQL binaries / page file
- At a minimum - data and xact log files separate
- What do you put on the fastest disk?

SQL Server - Disk

- Disk Setup - example

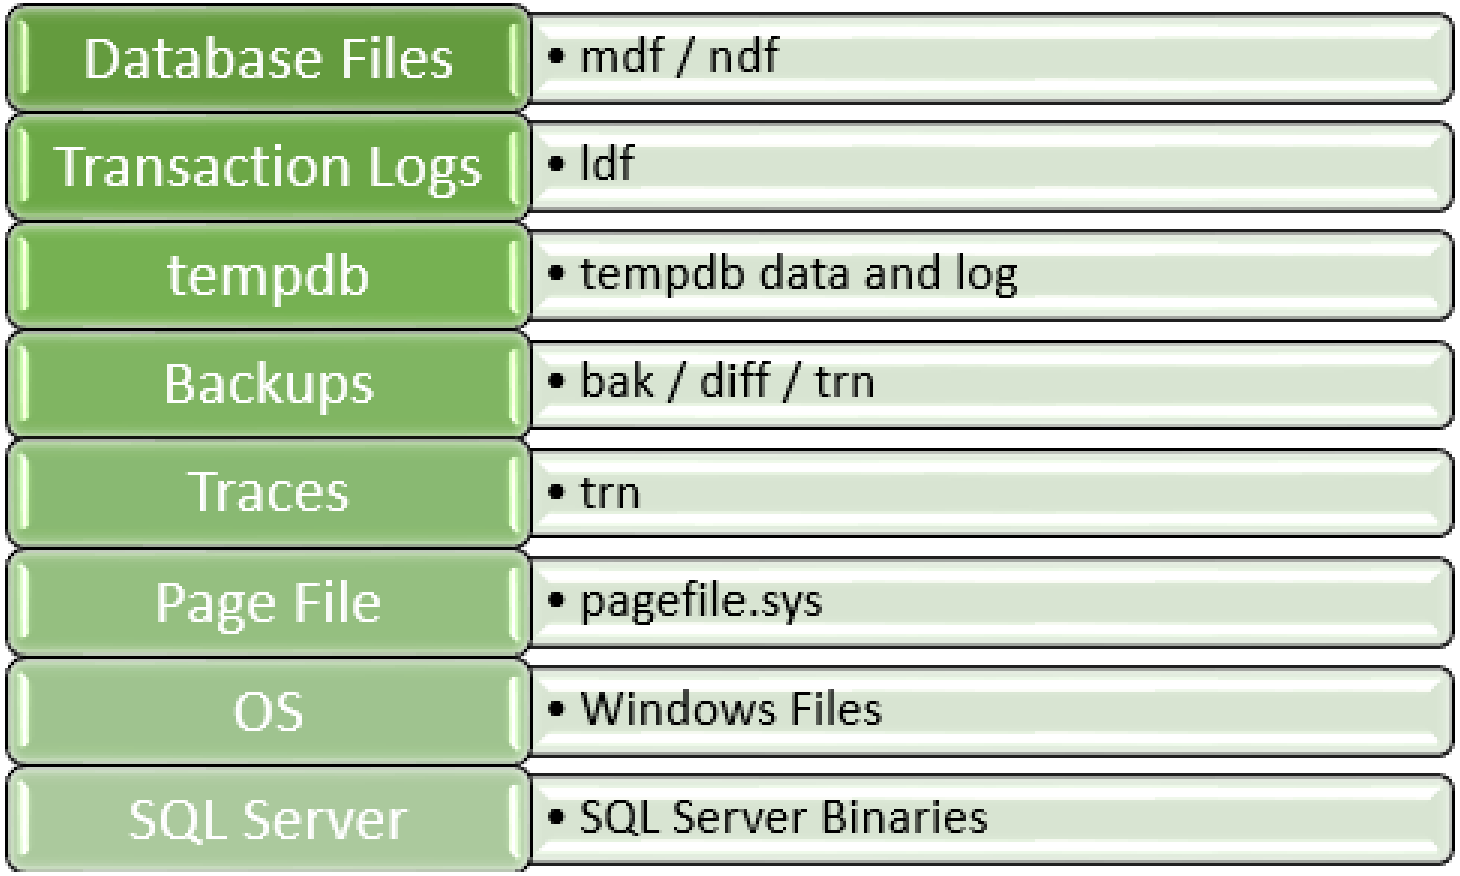

SQL Server – Instance Configuration

- Maximum worker threads
  - Default is 0 – leave it alone
  - https://msdn.microsoft.com/en-us/library/ms190219(v=sql.110).aspx
  - I have seen both too limited and too much – rare scenarios

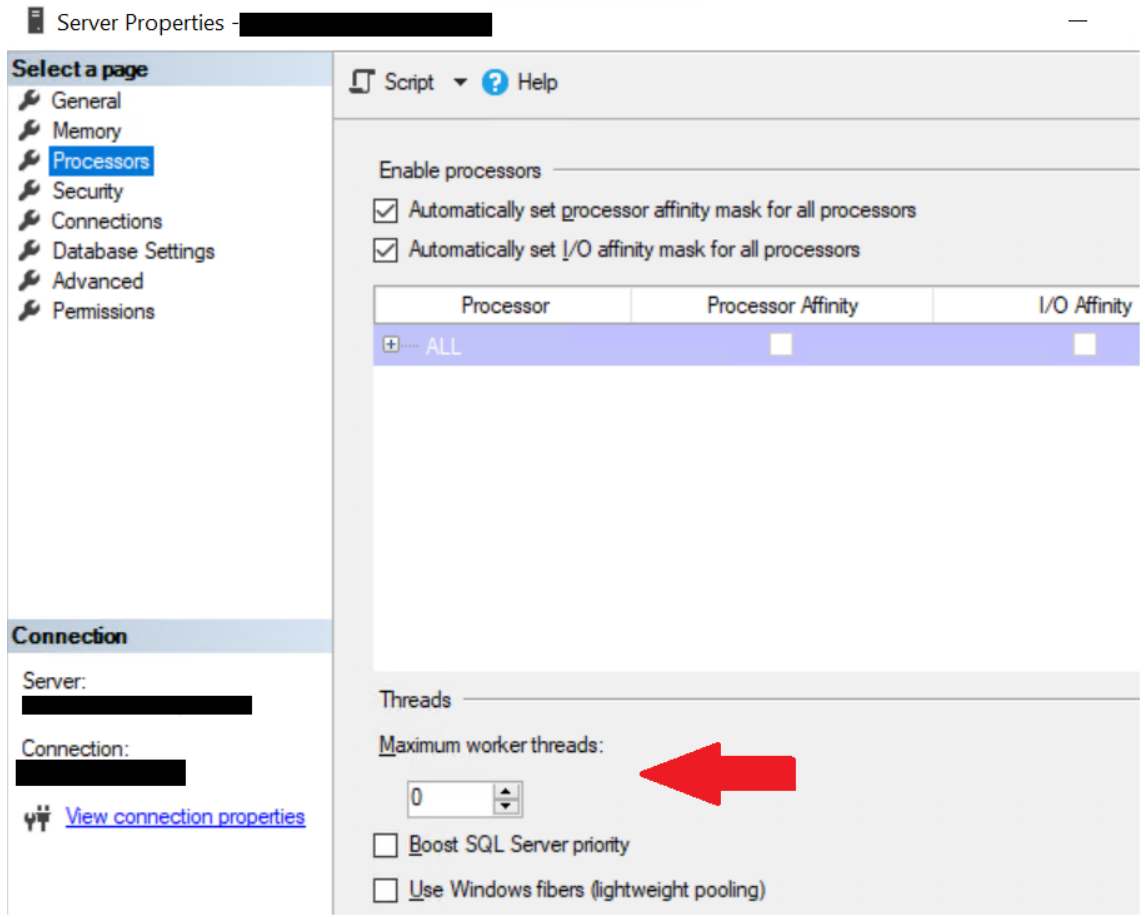

https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/configure-the-max-worker-threads-server-configuration-option

SQL Server – Instance Configuration

- Boost SQL priority
  - Default is unchecked – leave alone

FWIW marked for deprecation
https://technet.microsoft.com/en-us/library/ms180943(v=sql.105).aspx

SQL Server – Instance Configuration

- Remote server connections
  - Default is checked – leave it alone
  - This is an obscure SQL Server to SQL Server communication feature.
  - It has been **deprecated** and should not be used.
  - It is often confused with remote access to the instance. However, that is not the case.

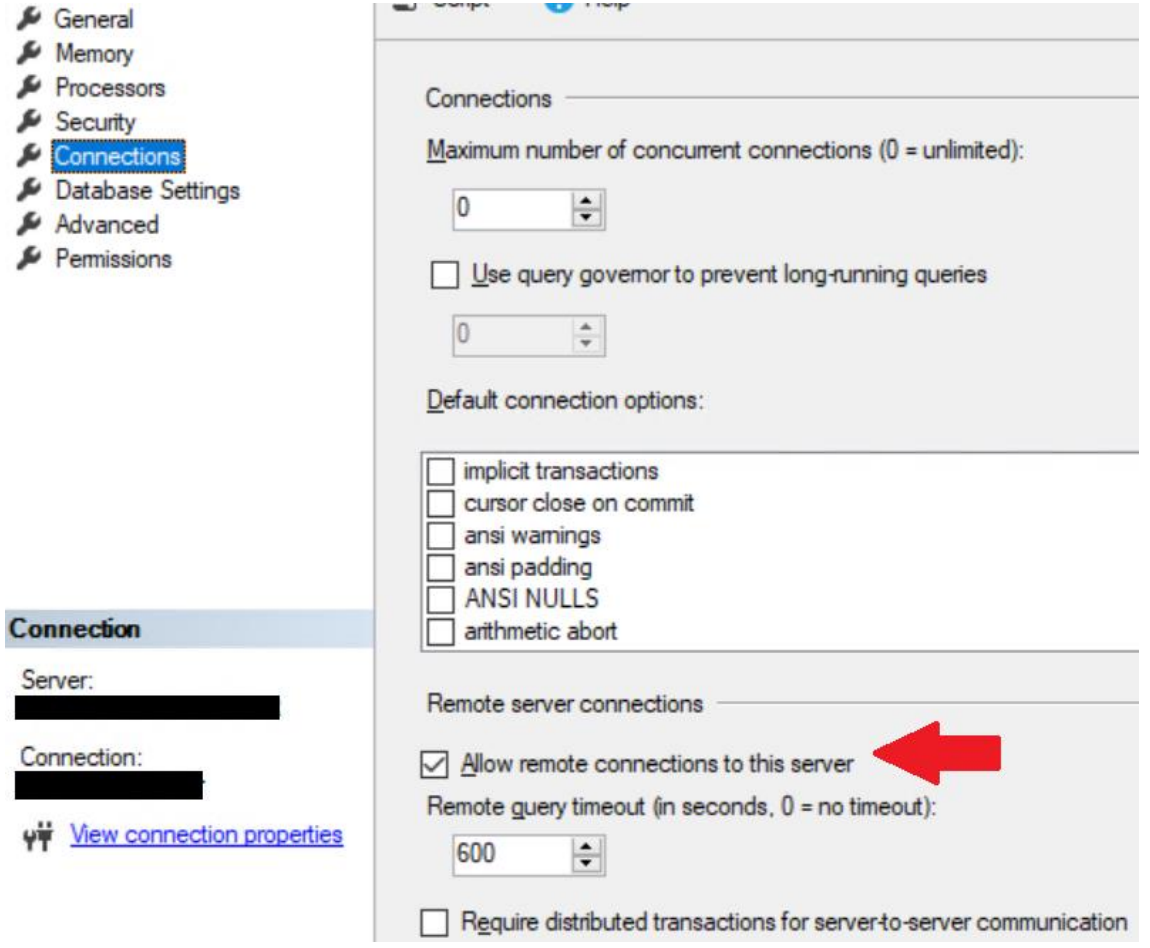

SQL Server – Instance Configuration

- Maximum # of concurrent connections
  - Default is 0 – leave it alone
- Remote query timeout
  - Default is 0 (no timeout) – leave it
- Query wait
  - Time in seconds that a query waits for resources before timing out
  - Default is -1 – leave it alone

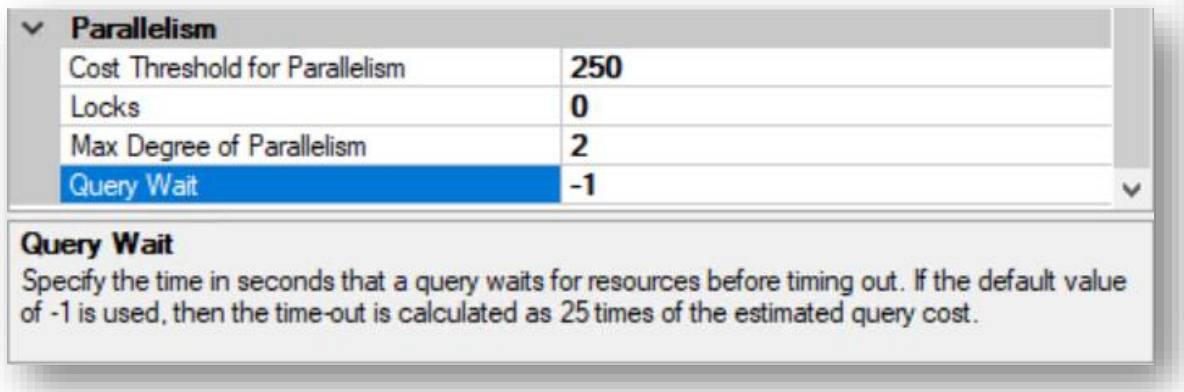

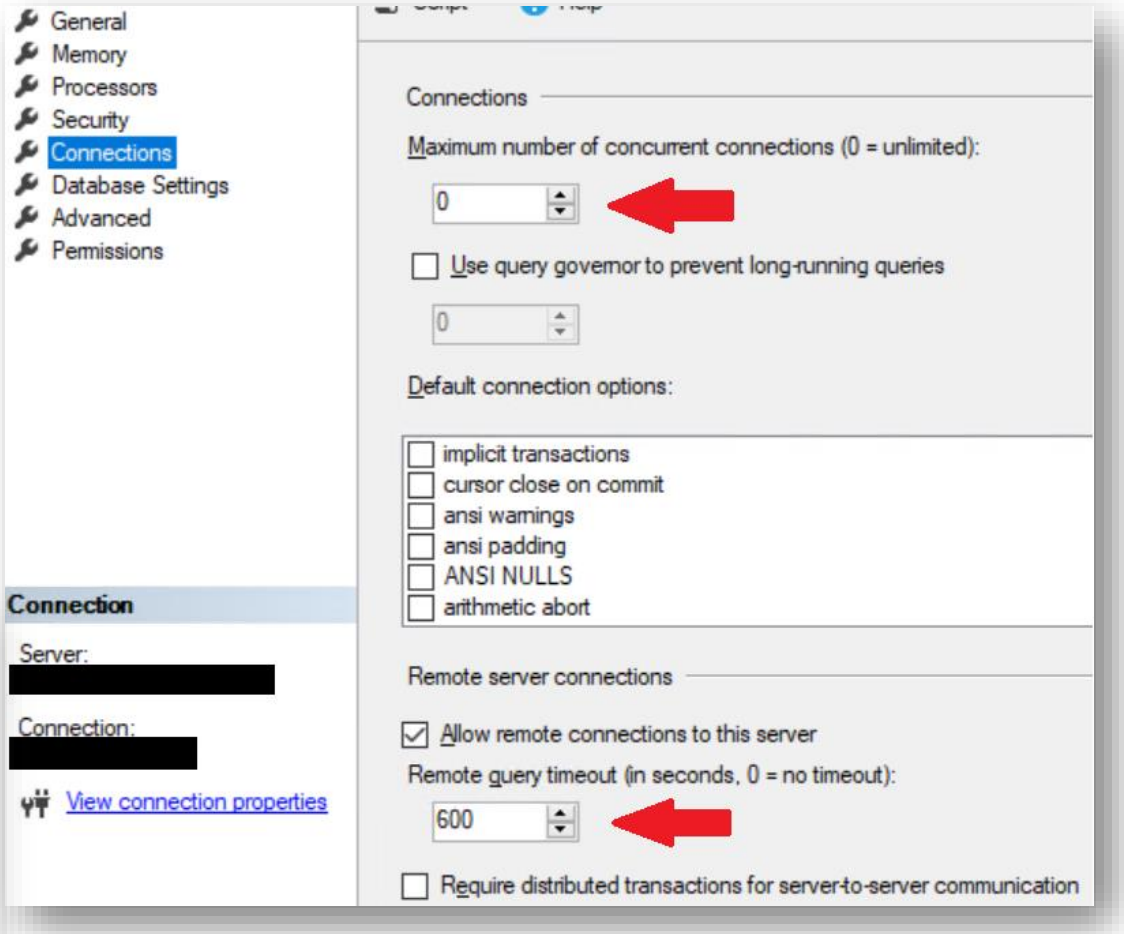

SQL Server – Instance Configuration

- TempDB
  - 1 file per logical core up to 8 then evaluate
  - Increase in sets of 4
  - Initial size should be the same for all
  - Autogrow in MB not %
  - Trying to reduce allocation contention
  - Trace Flags – default in SQL Server 2016
    - 1117
    - 1118

https://support.microsoft.com/en-us/help/2154845/recommendations-to-reduce-allocation-contention-in-sql-server-tempdb-d

SQL Server – Database Level Configurations

- Auto Close
- Auto Shrink

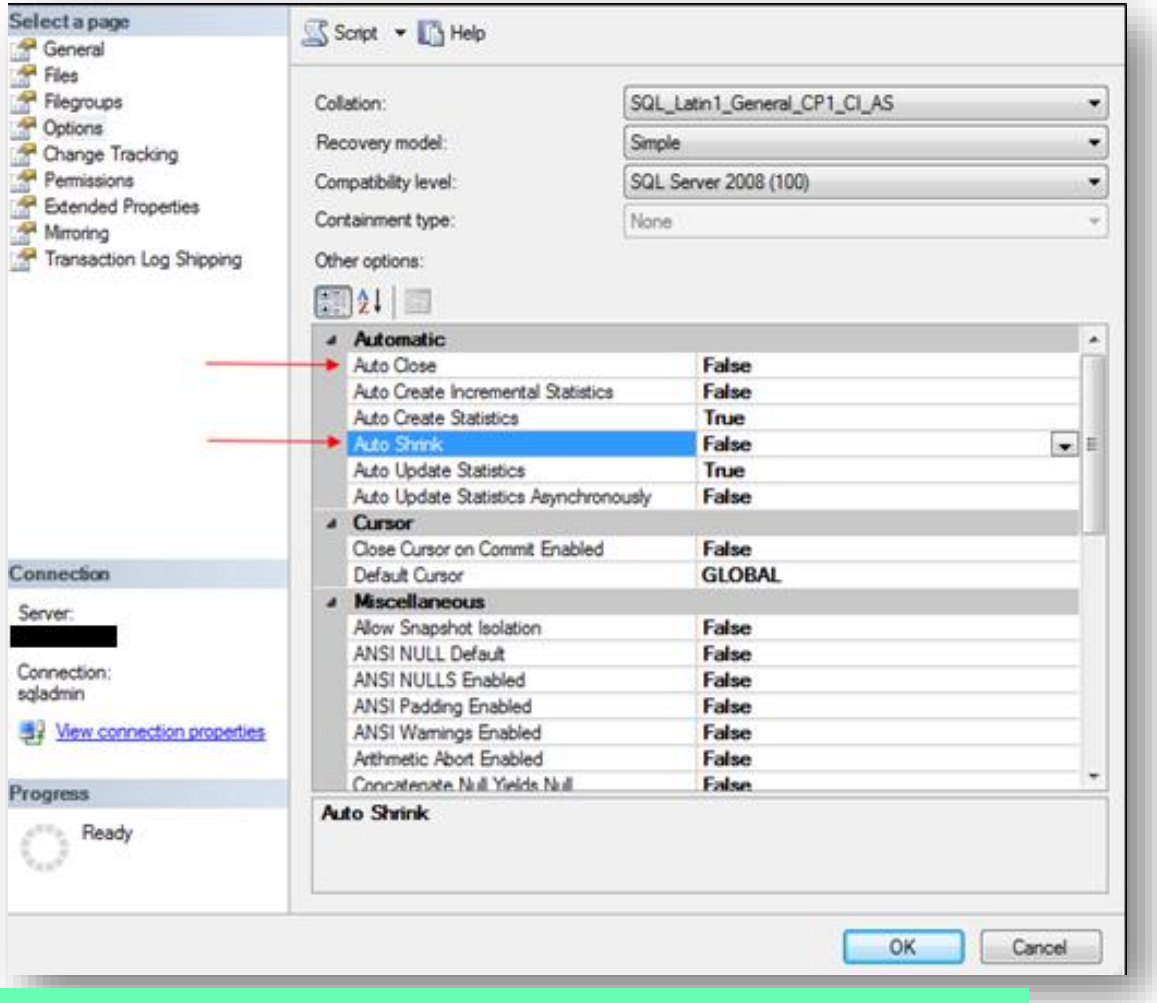

MS Best Practice: Considerations for the "autogrow" and "autoshrink" settings in SQL Server
Read this: http://support.microsoft.com/kb/315512

SQL Server – Database Configuration

- Database file defaults
  - Initial size
  - Autogrowth / Maxsize

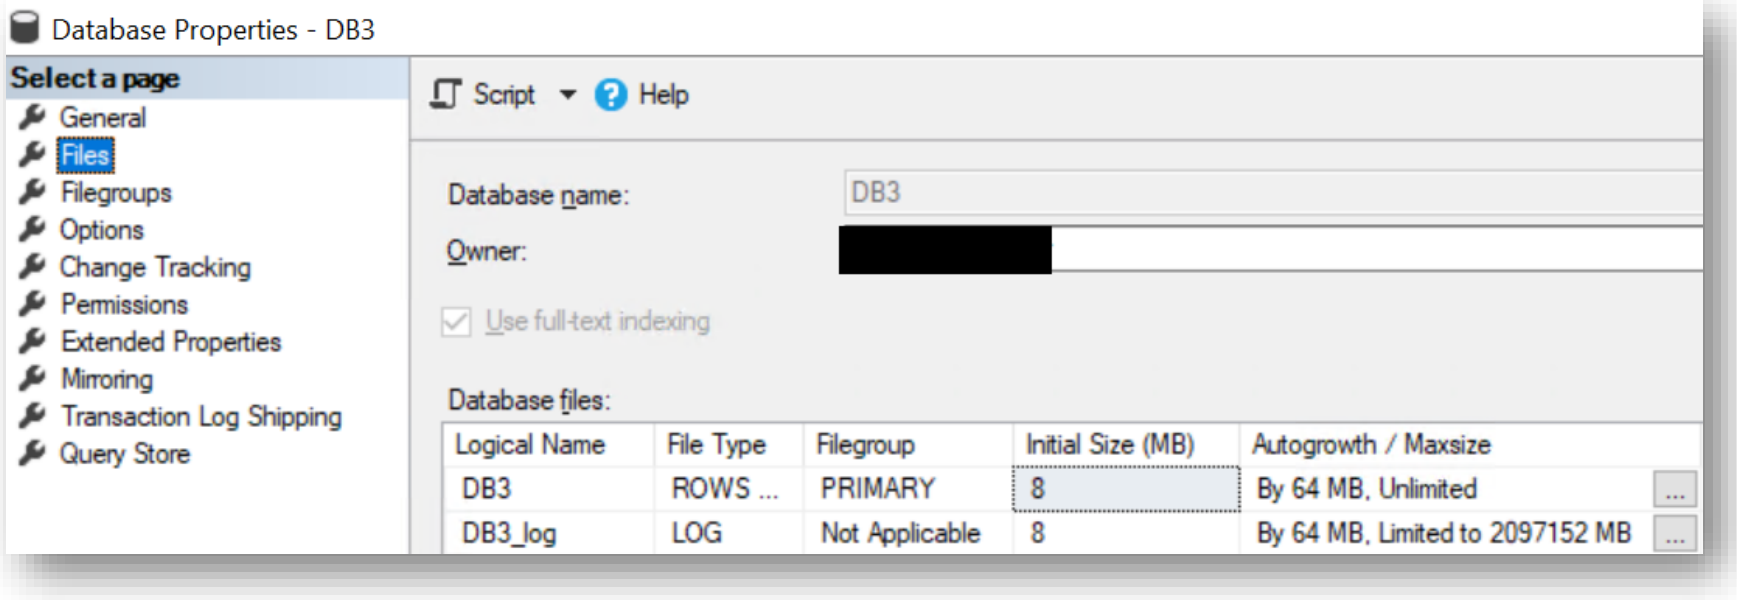

Storage Top 10 Best Practices: https://technet.microsoft.com/en-us/library/cc966534.aspx

Operating System Troubleshooting

Operating System

- Windows Server
  - General Configurations
  - Power Settings
  - AV
  - IFI
  - Page file
  - WSFC
  - Firewall
  - Scheduled Tasks
  - Disk Partition Alignment

Windows Server – General Configurations

- Windows Update
  - Disable automatic updates
  - Otherwise there will be unmanaged downtime to the applications running on this machine
- Windows OS Roles
  - Only install and activate necessary roles
- Windows OS Features
  - Only install and activate necessary features
- What non-default software is running on a machine running a SQL Instance?

Windows Server - Power

- Default power setting is “Balanced”.
  - This is **not** acceptable
  - Will throttle system resources globally to all applications and significantly hinder SQL performance
  - **ALWAYS** set to **HIGH** performance!

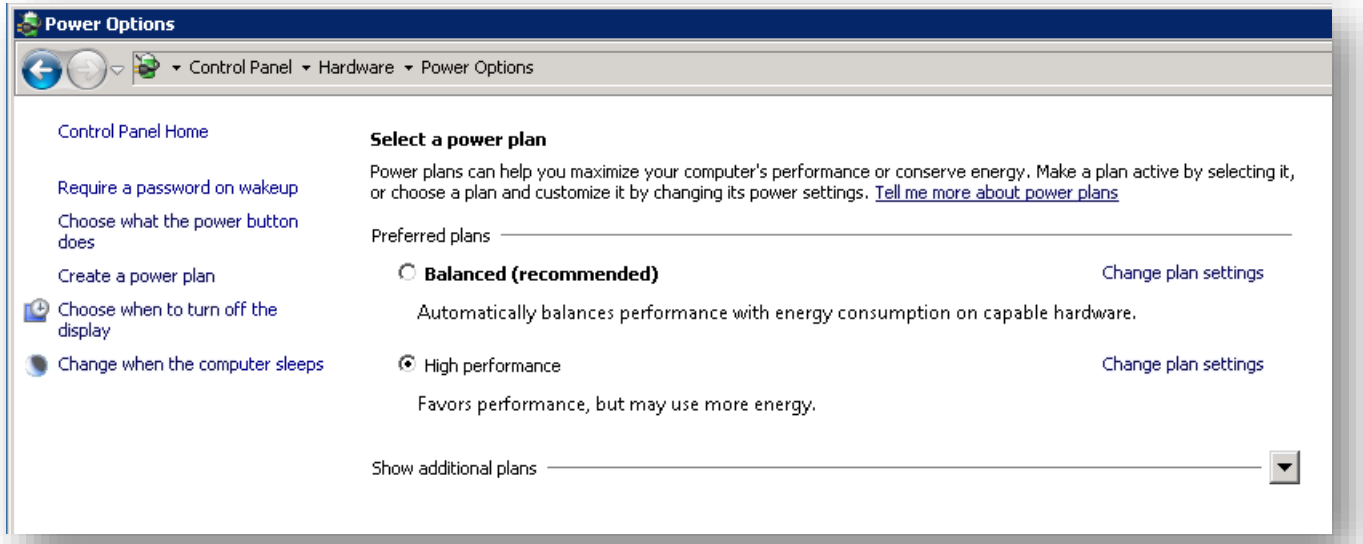

Windows Server - Power

- Power Setting Trickery
  - BIOS level power setting
    - HP ProLiant
    - Dell PowerEdge
  - Group Policy level power setting

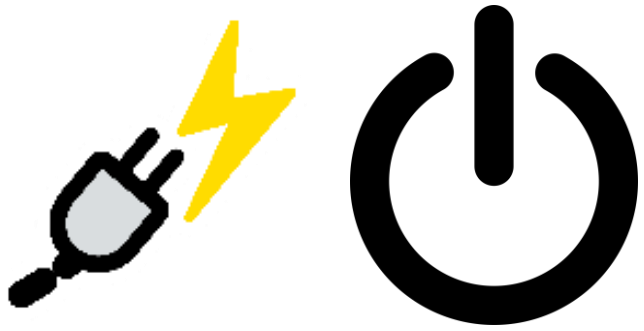

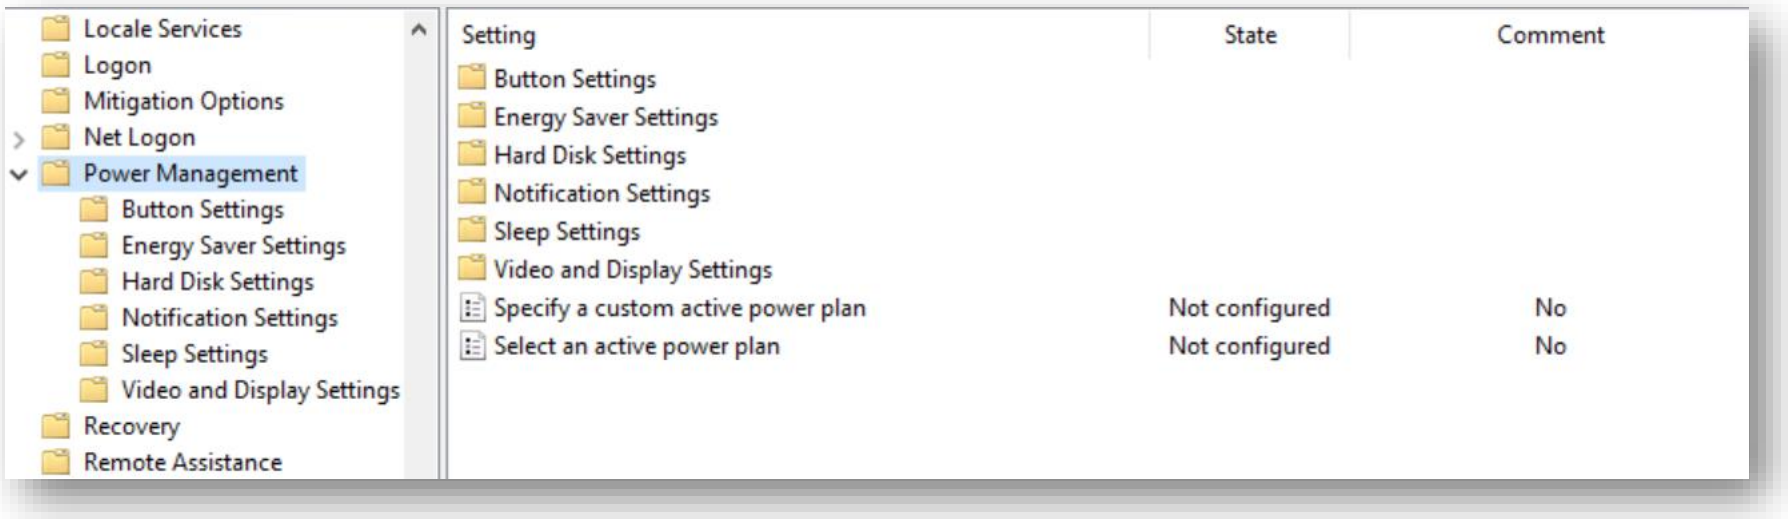

Windows Server - Power

- Power Setting Trickery
  - vSphere power setting

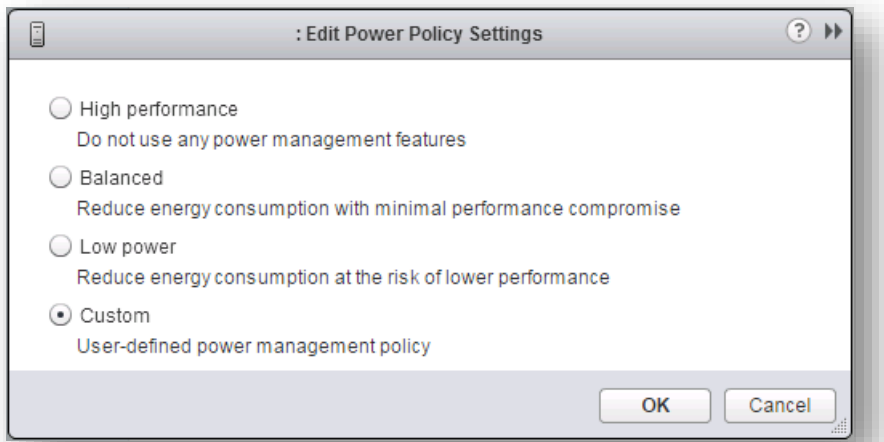

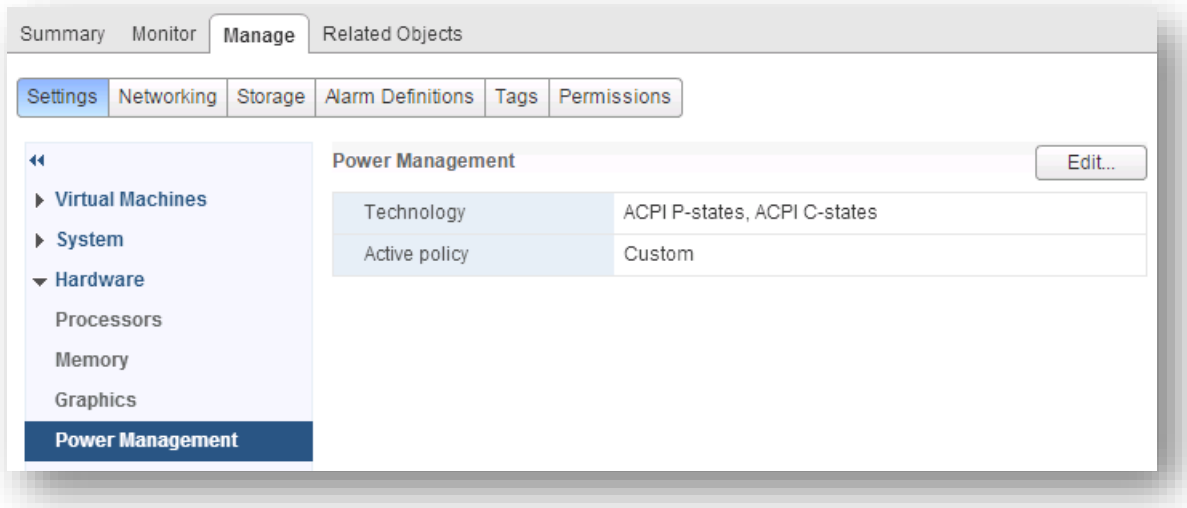

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/hpm-performance-vsphere55-white-paper.pdf

Windows Server - AV

- Anti Virus – Exclusions
  - If AV is running on SQL host then whitelist DB files
    - MDF –file extensions associated with SQL Server database files
    - LDF – file extensions associated with SQL Server transaction log files
    - BAK – file extensions associated with SQL Server backup files
    - TRN – file extensions associated with SQL Server trace files
  - Directories and file name extensions to exclude from scanning
    - https://docs.microsoft.com/en-us/sql/sql-server/install/file-locations-for-default-and-named-instances-of-sql-server

https://support.microsoft.com/en-us/help/309422/how-to-choose-antivirus-software-to-run-on-computers-that-are-running

Windows Server - Instant File Initialization

- File initialization
  - NTFS feature added to SQL Server 2005 for IFI using Windows Server 2003
  - Overwrite any existing data – fill with zeros
    - Create a database
    - Add data or log files to an existing database
    - Increase the size of an existing file (includes autogrowth)
    - Restore a database or filegroup
  - File won't be usable until it finishes zeroing out
  - SQL Error log after creating a DB post enabling IFI

```
2018-04-08 20:17:30.25 spid56      Zeroing C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA\IFI.mdf from page 0 to 1024 (0x0 to 0x800000)
2018-04-08 20:17:30.28 spid56      Zeroing completed on C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA\IFI.mdf (elapsed = 29 ms)
2018-04-08 20:17:30.30 spid56      Zeroing C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA\IFI_log.ldf from page 0 to 1024 (0x0 to 0x800000)
2018-04-08 20:17:30.34 spid56      Zeroing completed on C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA\IFI_log.ldf (elapsed = 38 ms)
2018-04-08 20:17:30.48 spid56      Starting up database 'IFI'.
2018-04-08 20:17:30.52 spid56      FixupLogTail(progress) zeroing C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA\IFI_log.ldf from 0x5000 to 0x6000.
2018-04-08 20:17:30.52 spid56      Zeroing C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA\IFI_log.ldf from page 3 to 249 (0x6000 to 0x1f2000)
2018-04-08 20:17:30.53 spid56      Zeroing completed on C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA\IFI_log.ldf (elapsed = 1 ms)
```

Windows Server – Instant File Initialization

- Local security policy
  - Local policies →
  - User rights assignment →
  - Perform volume maintenance tasks
- Add SQL Service account to enable IFI
- If changing the SQL Server service account then must restart service

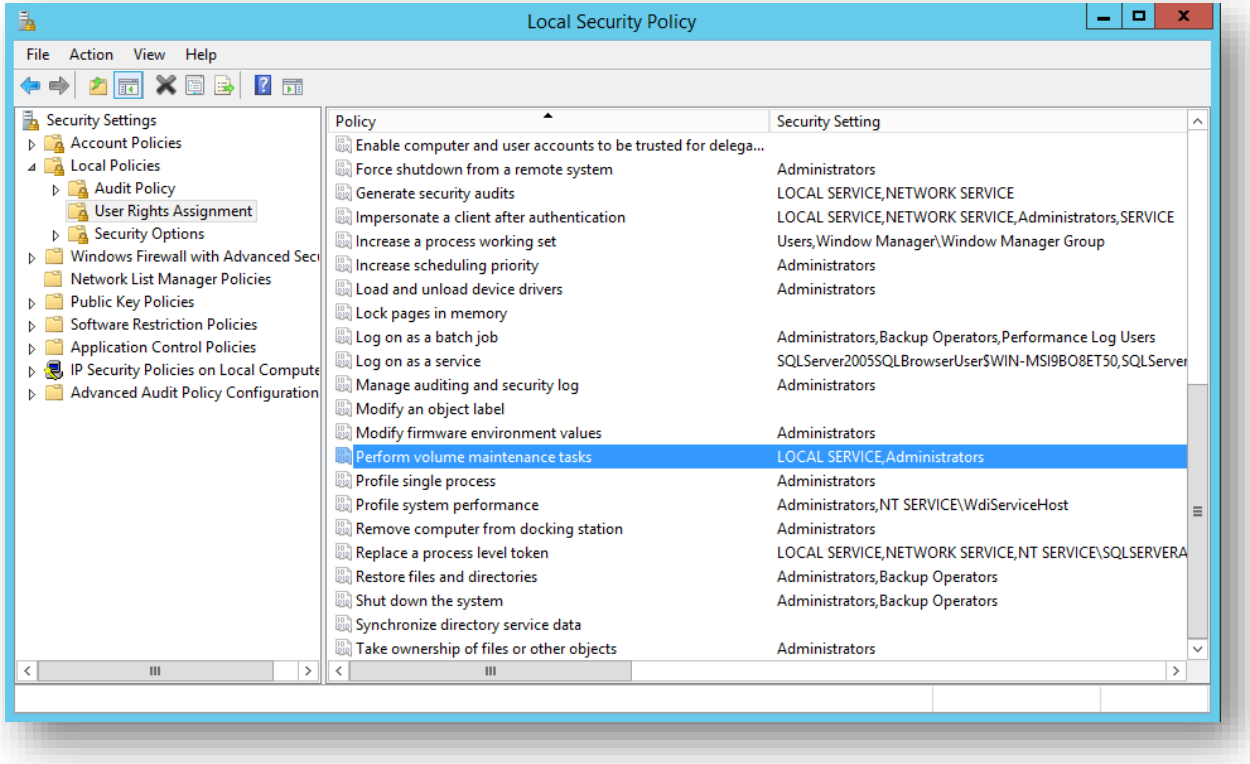

Windows Server – Instant File Initialization

- SQL Server 2016 added IFI option during install

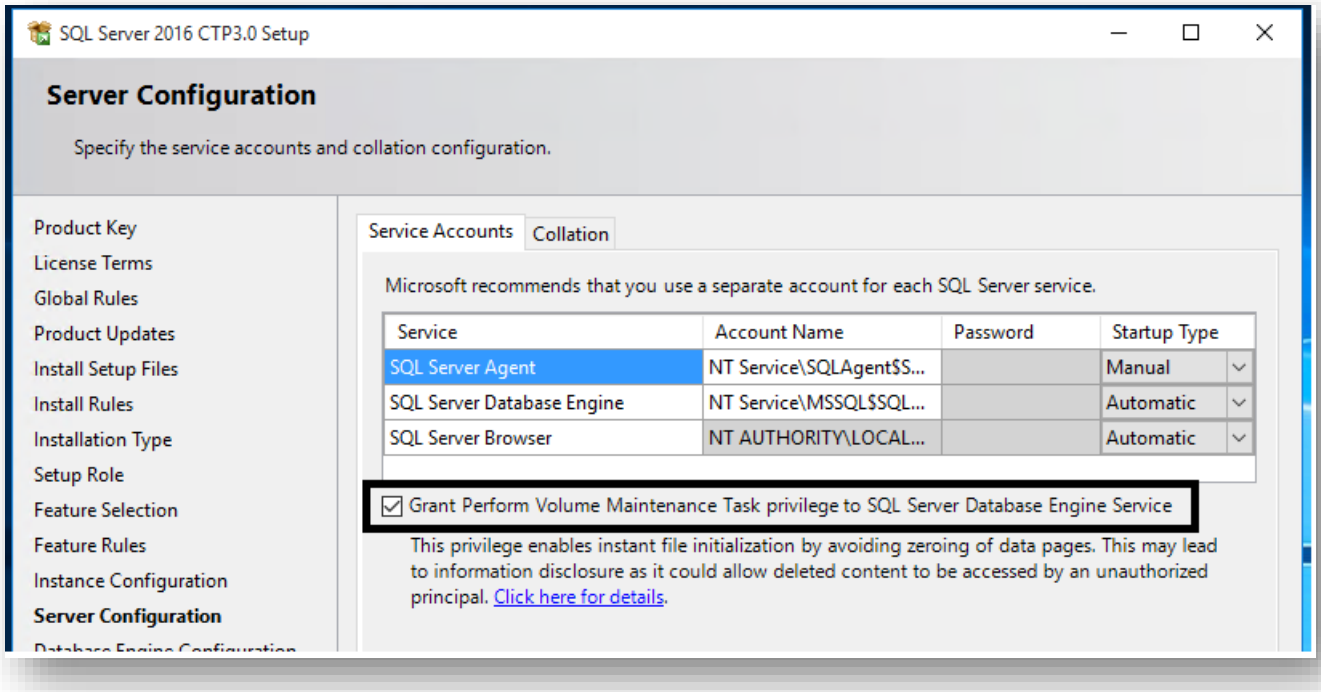

- Command line install option
  - setup.exe /Q /ACTION="INSTALL" /IACCEPTSQLSERVERLICENSETERMS /FEATURES="SQL" /INSTANCENAME="SQL2016" .. /**SQLSVCINSTANTFILEINIT**="True"

Windows Server – Instant File Initialization

PROS

1. Improved performance during file initialization…significant

CONS

1. Security Consideration
    a. Small possibility of disclosure of deleted content
    b. Granting “Perform Volumes Maintenance Tasks” to a SQL Instance means you are giving admins of the instance the ability to read the encrypted contents of a recently deleted file
    c. Mitigation
        a. Always make sure any detached DBs and backup files have restrictive DACLs
        b. Disable IFI for the instance
2. TDE

Overall – the PROS > CONS. Enable IFI as your default.

Windows Server – Page Files

- Page files are a special kind of file used as a temp workspace for storing modified pages from disk still in use by a process
- Holds data which is in the process of being swapped in and out of physical memory
- Allows a larger virtual memory set
- Large page files deserve their own disk (like data, xact log, tempdb, etc.)
- What does lots of page file usage mean?

Windows Server – Page Files

- PerfMon counters
  - **Memory: Committed Bytes** – number of bytes of virtual memory that has been committed
  - **Memory: Commit Limit** – number of bytes of virtual memory which can be committed without having to extend the paging files
  - **Paging File: % Usage** - % of the paging file committed
  - **Paging File: % Usage Peak** – highest % of the paging file committed

Windows Server – Page Files

- What is the Page File for anyway
  - https://blogs.technet.microsoft.com/askperf/2007/12/14/what-is-the-page-file-for-anyway/
- How to Determine the Appropriate Page File Size for 64-bit Versions of Windows
  - https://support.microsoft.com/en-us/help/2860880/how-to-determine-the-appropriate-page-file-size-for-64-bit-versions-of
- Page File – The Definitive Guide
  - https://blogs.technet.microsoft.com/motiba/2015/10/15/page-file-the-definitive-guide/

Windows Server - WSFC

- Windows Server Failover Cluster
  - If the WSFC feature is installed and running then make sure the best practices are being employed
  - Microsoft Windows Multi-Site Failover Cluster Best Practices (2012)
    - https://blogs.technet.microsoft.com/meamcs/2013/11/09/microsoft-windows-multi-site-failover-cluster-best-practices/
  - Windows Server 2008 R2 Failover Clustering – Best Practices Guide (2008 R2)
    - https://blogs.technet.microsoft.com/aevalshah/2012/05/15/windows-server-2008-r2-failover-clustering-best-practice-guide/

Windows Server - Firewall

- Windows Server Firewall
  - If running then make sure there are port exclusions for necessary ports for application to communicate

Port	Protocol	Usage
135	TCP	SSMS T-SQL Debugger
80	TCP	SSRS: http requests
443	TCP	SSRS: https requests SSL
1433	TCP	Default SQL Server port
1434	TCP	DAC
1434	UDP	SQL Server Browser

https://docs.microsoft.com/en-us/sql/sql-server/install/configure-the-windows-firewall-to-allow-sql-server-access

Windows Server – Scheduled Jobs

- Maintenance Jobs
  - SQL Server backups
  - SQL Server index maintenance
  - SQL Server dbcc checkdb
  - Disk space checks
- Make sure setup right and run off peak
- Monitor output and errors

Windows Server – Disk Partition Alignment

- Optimal disk configuration
  - Windows default is 1,024 kb cluster
  - Start at a more common sizing of 64 kb
  - Greater chance of playing nice with disks, controllers, and cache
  - Formatting disk to 64 kb cluster size can remediate suboptimal I/O performance

Disk Partition Alignment Best Practices for SQL Server
https://technet.microsoft.com/en-us/library/dd758814(v=sql.100).aspx

Virtualization

Virtualization

- VMware – accept no substitute
  - SQL Server on VMware best practices guide
- General Configurations
- CPU Ready
- Memory Ballooning
- Disk

VMware

- Troubleshooting Guidelines
  - Troubleshooting ESX/ESXi Virtual Machine Performance Issues
    - https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2001003
  - Tips for Configuring Microsoft SQL Server in a Virtual Environment
    - https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1002951

VMware

vSphere Key Performance Metrics

Table 9. Key Performance Metrics

Resource	Metric (resxtop)	Metric (vSphere Client)	Host/Virtual Machine	Description
CPU	%USED	Used	Both	CPU used over the collection interval (%).
	%RDY	Ready	Virtual Machine	CPU time spent in ready state.
	%SYS	System	Both	Percentage of time spent in the vSphere Server VMKernel.
Memory	Swapin, Swapout	Swapinrate, Swapoutrate	Both	Memory vSphere host swaps in/out from/to disk (per virtual machine, or cumulative over host).
	MCTLSZ (MB)	vmmemctl	Both	Amount of memory reclaimed from resource pool by way of ballooning.
Disk	READs/s, WRITEs/s	NumberRead, NumberWrite	Both	Reads and Writes issued in the collection interval.
	DAVG/cmd	deviceLatency	Both	Average latency (ms) of the device (LUN).
	KAVG/cmd	KernelLatency	Both	Average latency (ms) in the VMkernel, also known as queuing time.
	GAVG/cmd	TotalLatency	Both	Average latency (ms) in the guest. GAVG = DAVG + KAVG.
Network	MbRX/s, MbTX/s	Received, Transimitted	Both	Amount of data transmitted per second.
	PKTRX/s, PKTTX/s	PacketsRx, PacketsTx	Both	Packets transmitted per second.
	%DRPRX, %DRPTX	DroppedRx, DroppedTx	Both	Dropped packets per second.

VMware – CPU Ready

- Overcommitting the VM Host CPU to Guest VMs
  - Can cause more trouble than benefit.
  - Hypervisor must keep track of CPUs and context switch between them across all guest VMs.
  - Try to “Right-Size” the guest machines rather than over commit.
- ***Recommend a CPU Ready of under 5%.***
- The command “esxtop” can be run from the ESX host to get general statistics about the VM host.

VMware – CPU Ready

- Waits in CPU Ready below 10,000ms.
  - A range of 5000-8000ms should be as high as they get.
  - Reservations on CPU
- CPU Shares - High
- Converting Between CPU Summation and CPU % Ready Values
  - https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2002181
- Determining if Multiple Virtual CPUs are Causing Performance Issues
  - https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1005362

VMware – CPU Ready

- Examples of calculating CPU % Ready
  - Use the following formulas for the default chart update intervals
    - Realtime: CPU summation value / 200
    - Past Day: CPU summation value / 3000
    - Past Week: CPU summation value / 18000
    - Past Month: CPU summation value / 72000
    - Past Year: CPU summation value / 864000

VMware – CPU Ready

VMware – CPU Ready

- CPU Ready %
  - Min
    - 6645 / 18000 = 0.3692
  - Average
    - 16850 / 18000 = 0.9361
  - Max
    - 38067 / 18000 = 2.1148

VMware – CPU Ready

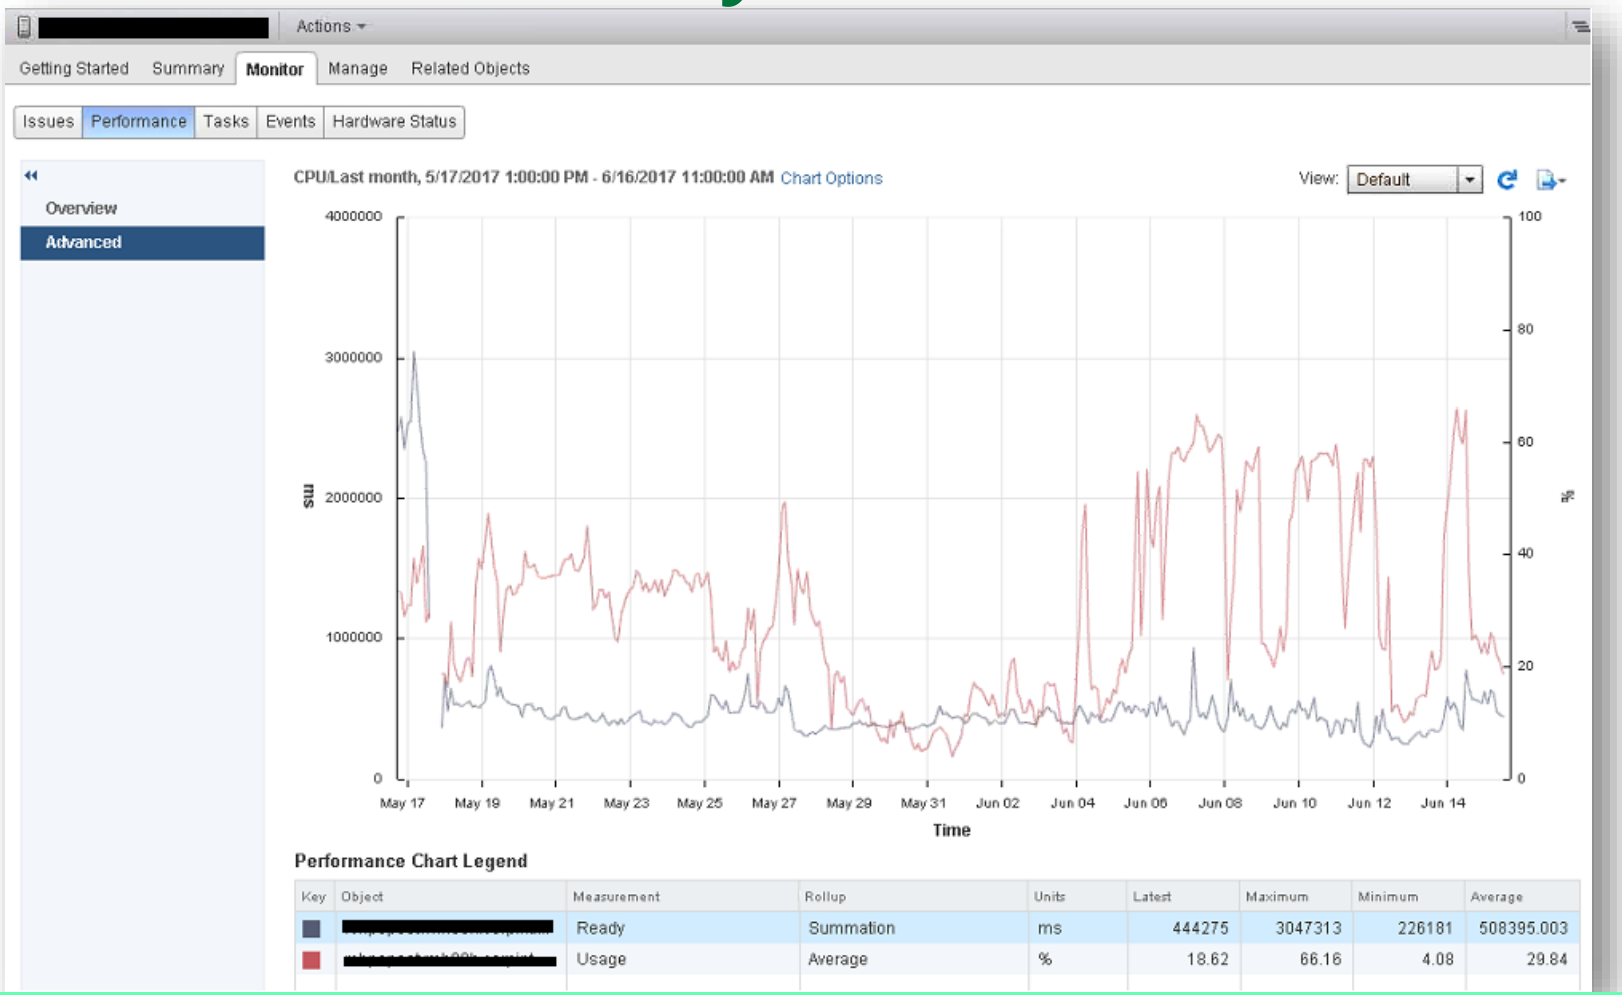

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2002181

VMware – Memory Ballooning

- Memory reservations
  - If in place make sure there is enough memory in the lower bound for the guest machine to perform without excessive paging.
  - Also make sure the VM host isn't stressed for memory and the hypervisor doesn't have to reclaim memory to service other guests.
- If memory is overcommitted then either increase memory to the host OR reduce memory to the guest VMs

VMware – Memory Ballooning

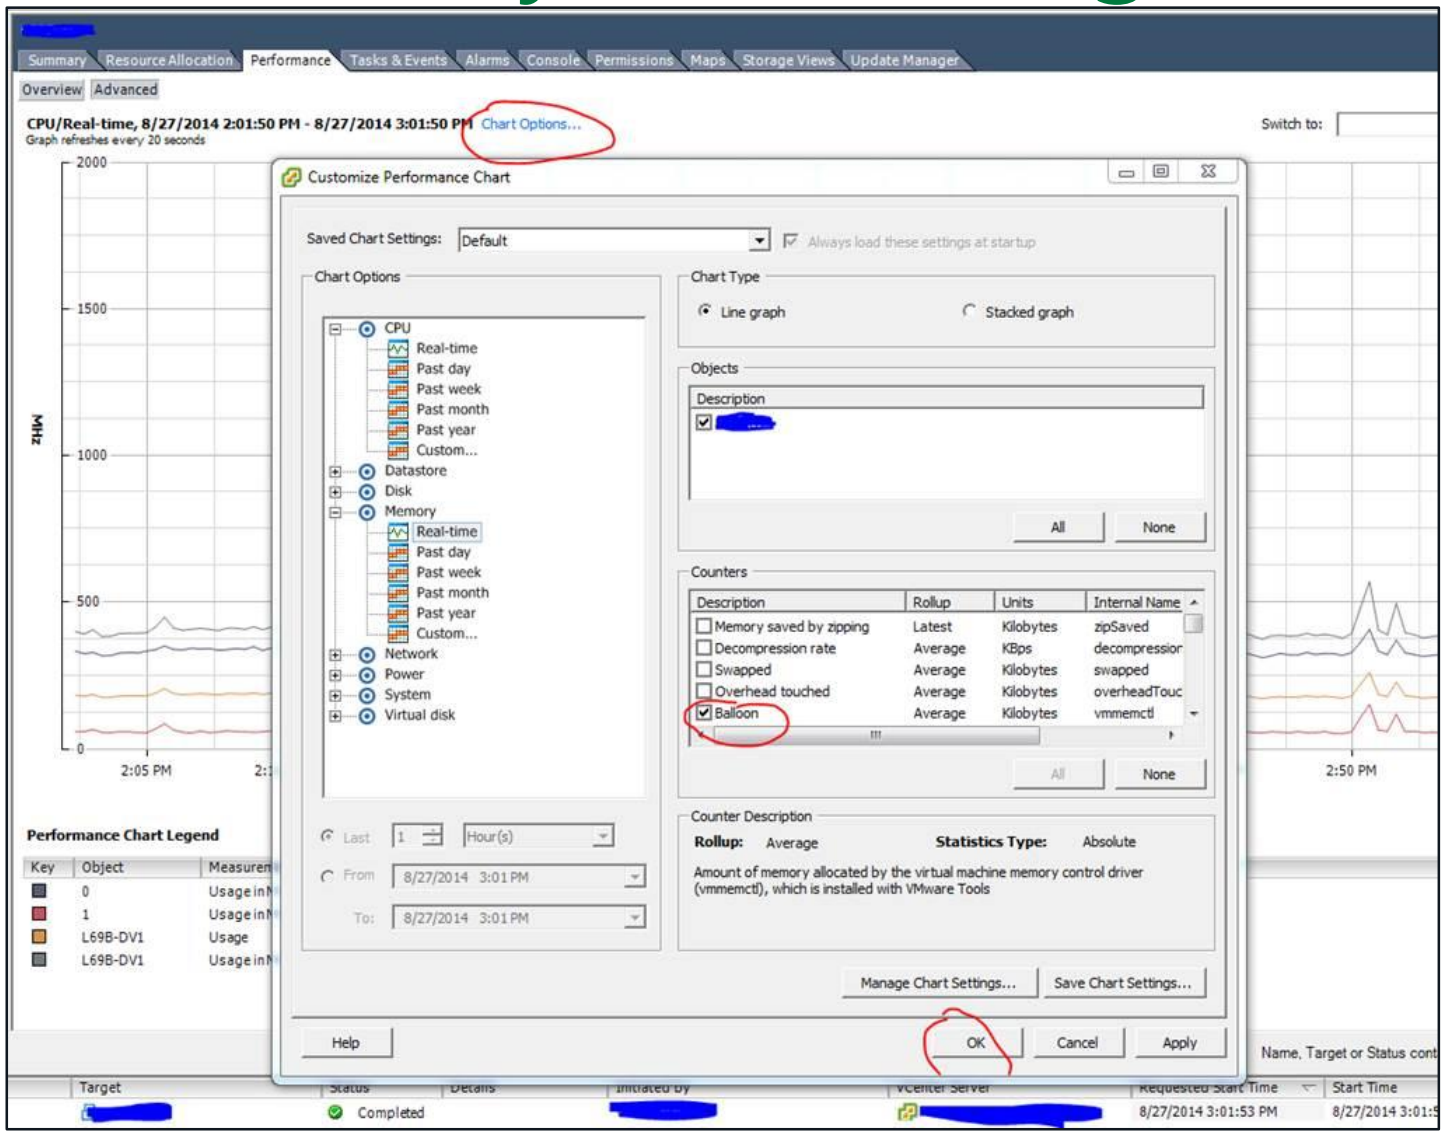

VMware – Memory Ballooning

VMware - Disk

- PVSCSI – Paravirtual SCSI adapter

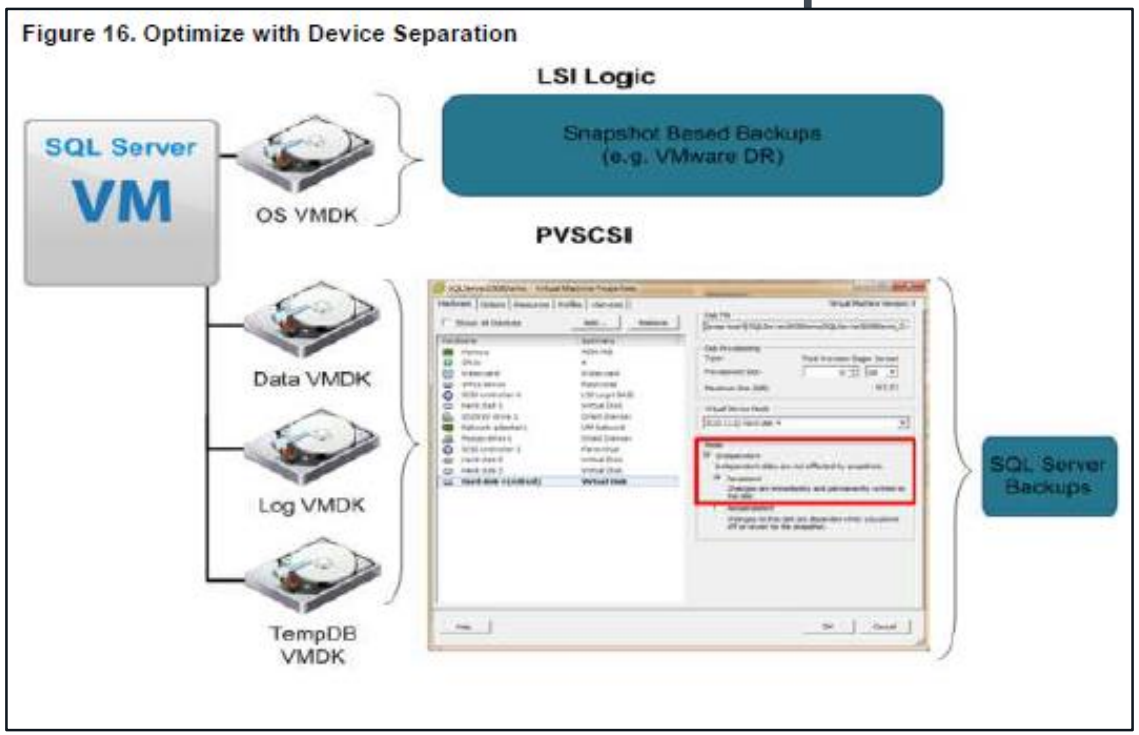

Configuring Disks to Use VMware Paravirtual SCSI (PVSCSI) Adapters
https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1010398

Summary

Conclusions

- Defaults are often poor!
- Use best practices to configure SQL Server to minimize performance issues
  - Test, test, test!
- Know when to step outside the guidelines
  - Learn the exceptions

Thank You!
Q&A

@jmlakar

https://www.linkedin.com/in/jeffmlakar/

www.MlakarTechTalk.com