

# **Tuning Microsoft SQL Server 2005 Performance**

*Best practices and performance services offered  
by Hitachi Data Systems Global Solution  
Services (GSS) for Microsoft SQL Server  
applications*

<i>Introduction</i> .....	3
<i>Hitachi HiCommand® Tuning Manager Software</i> .....	3
<b>Custom Reports</b> .....	4
<i>Hitachi Data Systems Global Solution Services</i> .....	5
<b>Tuning Manager Software Implementation Services</b> .....	5
<b>Custom Reports Module for Microsoft SQL Server</b> .....	5
<i>Microsoft SQL Server Custom Reporting and Analysis Methodology</i> .....	5
<b>Mapping LUNs to Array Groups and Ports</b> .....	6
<b>Drill-down Performance Analysis Using Tuning Manager Software</b> .....	7
<b>Performance Analysis Flowchart</b> .....	8
<b>SQL Server Rules of Thumb and Best Practices</b> .....	9
<i>Summary</i> .....	11
<i>Appendix A</i> .....	11
<b>Microsoft SQL Server Performance Monitor Counters</b> .....	11
<i>Appendix B</i> .....	11
<b>Storage Top 10 Best Practices</b> .....	11

## Introduction

Microsoft SQL Server has become a leading database development platform for businesses. While there are some similarities, each application developed using SQL Server has its own I/O profile and therefore its own performance characteristics.

Gaining a complete view of the I/O path—software, servers, host bus adapters (HBAs), switches, switch and storage subsystem ports, storage subsystem, and disks that collectively make up the SQL Server infrastructure—is essential to understanding and managing performance. But accessing this information is no easy task. Few tools offer a sufficiently detailed picture for administrators to state with certainty the cause of poor response times.

Hitachi Tuning Manager software offer full application-to-spindle performance monitoring, including an agent specifically for monitoring SQL Server databases. It tracks over 150 performance metrics within SQL Server, Microsoft Window Server hosts, SAN switches and within Hitachi storage systems. To make the product effective in managing SQL Server databases, however, requires some best practices, or “Rules-of-Thumb”, to know what to look for. Hitachi Data Systems Global Solution Services (GSS) offers a custom reporting service based on Tuning Manager, applying the Rules of Thumb, that allows customers to monitor and manage their SQL Server applications, Combined these provide administrators the tools and knowledge to ensure storage systems used by the enterprise for database applications remain fully optimized.

## Hitachi HiCommand® Tuning Manager Software

Hitachi HiCommand® Tuning Manager software provides insight into the enterprise IT environment along the entire path taken by application I/O. An automated, intelligent, and path-aware storage resource management tool, Tuning Manager software maps, monitors, analyzes, and reviews storage network resources in real-time, giving administrators instant access to the current state of hosts, file systems, databases, SANs, and storage systems.

Using agent software, installed on platforms in the path of application I/O, Tuning Manager software is able to provide a comprehensive assessment of the performance and capacity of the entire I/O infrastructure. In addition to real-time reports, Tuning Manager software compiles metrics in a historical database allowing administrators to compare current and past behavior to quickly identify anomalies.

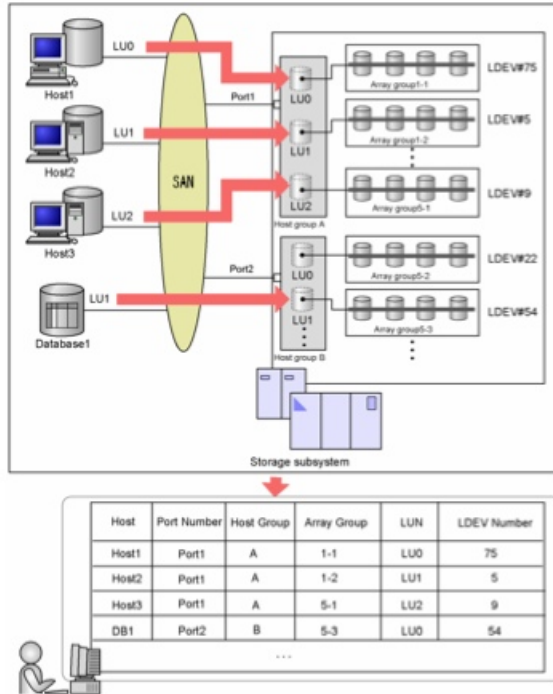


Figure 2. Hitachi HiCommand® Tuning Manager software integrates performance and capacity data from the entire I/O path.

Tuning Manager software reports are based on real-time snapshots, past point in time, or historical trending over a specified time period. And, reports are available from the perspective of the host, storage system, storage area network (SAN) fabric, application, or network attached storage (NAS) device. The application perspective is provided via integration with common enterprise applications, such as Oracle, IBM® DB2®, and Microsoft SQL Server, giving administrators a complete picture of I/O path activity.

Tuning Manager software reports give administrators access to critical information necessary for addressing performance problems, maintaining service level agreements (SLAs), predicting long-term trends, and forecasting future application storage needs. Integration with HiCommand Device Manager software allows Tuning Manager software to analyze performance and capacity utilization on all Hitachi storage systems.

### Custom Reports

In addition to standard solution set reporting options, provided with each Tuning Manager software agent, custom reports are available to tailor the display of information for specific needs. Custom reports help administrators quickly access a comprehensive view of a particular object being monitored in the environment. For example, during a period of degraded performance, custom reports can show real-time IOPS viewed from both the host server and the storage subsystem, giving SQL Server administrators an instant picture of the I/O environment. Contrasting this information with historical data can help to quickly highlight problem areas needing further investigation.

Tuning Manager software custom reports are generated through a wizard-driven interface that prompts the user for specific information to tailor the data view. However, developing effective custom reports requires a thorough understanding of the different Tuning Manager software agent data models, and sufficient knowledge of the storage subsystem to be able to interpret reported results. To bridge this gap between the raw performance data and effective problem analysis, Hitachi Data Systems GSS consultants have developed a solution designed specifically for the needs of SQL Server support staff.

## **Hitachi Data Systems Global Solution Services**

Even at the most effective organizations, the day-to-day challenge of managing complex storage infrastructures can stretch IT staff to the limit. New solutions run the risk of becoming shelf-ware if administrators do not have the time to effectively deploy and get to know the technology.

Hitachi Data Systems Global Solution Services (GSS) group understands the tremendous challenge of managing today's storage infrastructures and has developed specific services to help IT staff install, configure, customize, and come up to speed on new technologies and solutions, in a fraction of the time it would take a person to perform these tasks on their own. GSS specialists will help deploy new solutions to meet specific business requirements, upgrade existing technologies to accommodate changing workloads, solve performance problems, and perform capacity planning to ensure future growth is accommodated without infrastructure disruption.

### **Tuning Manager Software Implementation Services**

To help organizations gain the maximum benefit from Tuning Manager software, in the shortest possible time, GSS has developed a service offering specifically designed to help administrators rapidly come up to speed with the software's capabilities. During the short-term engagement GSS engineers ensure Tuning manager software is installed and configured correctly and perform a knowledge transfer enabling IT staff to benefit immediately from access to the data in Tuning Manager software reports.

The Implementation Service for Hitachi HiCommand® Tuning Manager software provides installation and configuration of Tuning Manager software, and validation for up to five RAID agents and up to ten host agents for data collection and reporting by the management server for up to five storage systems and up to ten hosts.

As part of the Tuning Manager implementation service, GSS consultants conduct a three-day workshop with customer technical staff. The workshop provides in-depth training on successful Tuning Manager software installation and configuration processes, in addition to a comprehensive overview of the software's functionality and its architecture. After completing the service, IT staff are able to use Tuning Manager software to tune the storage environment and more efficiently manage growth by monitoring, reporting, and forecasting capacity and performance.

### **Custom Reports Module for Microsoft SQL Server**

For immediate identification of host device performance problems with Microsoft SQL Server applications, the GSS implementation service offers an optional Custom Reports Module for Microsoft SQL Server. With this module, GSS installs and configures Tuning Manager software Custom Performance Reporter user reports specifically designed to diagnose SQL Server performance problems.

The Custom Reports Module engagement begins by tracking device allocations at the SQL Server host into the storage system. During this process, all storage resources used by the host are identified and profiled. This information is then fed into custom Tuning Manager report templates and the results are analyzed to determine the cause of the performance problem. In addition to resolving specific issues, GSS specialists teach administrators how to interpret data in the custom Tuning Manager software reports and methodically diagnose future performance problems.

## **Microsoft SQL Server Custom Reporting and Analysis Methodology**

The GSS Custom Reports Module for Microsoft SQL Server is designed to help administrators identify and measure key host and storage performance metrics. Uniquely identifying the location of SQL databases and logs on the storage system, GSS consultants give SQL administrators the information they need to track ongoing compliance with best practice thresholds.

The Tuning Manager custom reports offer a single source for all performance data, supporting a drill down methodology for quickly analyzing and resolving performance issues. By aligning host and storage performance with best practice thresholds defined by Microsoft and Hitachi Data Systems, administrators can ensure SQL Server is always functioning within defined guidelines.

### Mapping LUNs to Array Groups and Ports

The first step in analyzing performance is to identify all storage resources used by the SQL Server database and log files. GSS consultants combine information from an Hitachi Data Systems proprietary SQL Server Storage Analyzer Tool with array group and port data from Tuning Manager software RAID Map reports.

SQL Server TPCC Database					
Name	Type	File	Allocated (MB)	Used (MB)	Source
MSSQL_tpc_log	Log File	s:\MSSQL_log.ldf	50000	-	S:\
MSSQL_misc1	Data File	f:\MSSQL_misc1.ndf	5000	3804	F:\
MSSQL_misc2	Data File	g:\MSSQL_misc2.ndf	5000	3803	G:\
MSSQL_misc3	Data File	h:\MSSQL_misc3.ndf	5000	3803	H:\
MSSQL_misc4	Data File	i:\MSSQL_misc4.ndf	5000	3804	I:\
MSSQL_misc5	Data File	j:\MSSQL_misc5.ndf	5000	3803	J:\
MSSQL_misc6	Data File	k:\MSSQL_misc6.ndf	5000	3805	K:\
MSSQL_cs1	Data File	l:\MSSQL_cs1.ndf	7300	7194	L:\
MSSQL_cs2	Data File	m:\MSSQL_cs2.ndf	7300	7195	M:\
MSSQL_cs3	Data File	n:\MSSQL_cs3.ndf	7300	7195	N:\
MSSQL_cs4	Data File	o:\MSSQL_cs4.ndf	7300	7194	O:\
MSSQL_cs5	Data File	p:\MSSQL_cs5.ndf	7300	7194	P:\
MSSQL_cs6	Data File	r:\MSSQL_cs6.ndf	7300	7194	R:\

Figure 3. Hitachi Data Systems SQL Server Storage Analyzer Tool.

The Tuning Manager software RAID Map details allocations on the storage subsystem showing file system identifier, port name, channel adapter, logical unit number (LUN), logical device (LDEV), RAID group, RAID level, and disk size and speed. This information is correlated with data from the Storage Analyzer tool and placed in an Excel spreadsheet to drive the performance analysis.

RAID MAP								
File Sys	RAID Lvl	Prod ID	Prod Name	Ser Num	LUN	Target ID	CHA MP	
S:	5	OPEN-V	USP	10075	0	0	CHA-2Q 0	
F:	1	OPEN-V	USP	10075	1	0	CHA-2Q 0	
G:	1	OPEN-V	USP	10075	2	0	CHA-2Q 0	
H:	1	OPEN-V	USP	10075	3	0	CHA-2Q 0	
I:	1	OPEN-V	USP	10075	4	0	CHA-2Q 0	
J:	1	OPEN-V	USP	10075	5	0	CHA-2Q 0	
K:	1	OPEN-V	USP	10075	6	0	CHA-2Q 0	
L:	1	OPEN-V	USP	10075	7	0	CHA-2Q 0	
M:	1	OPEN-V	USP	10075	8	0	CHA-2Q 0	
N:	1	OPEN-V	USP	10075	9	0	CHA-2Q 0	

O:	1	OPEN-V	USP	10075	10	0	CHA-2Q 0
P:	1	OPEN-V	USP	10075	11	0	CHA-2Q 0
R:	1	OPEN-V	USP	10075	12	0	CHA-2Q 0

Figure 4. Tuning Manager software custom RAID map report for SQL Server platform.

### Drill-down Performance Analysis Using Tuning Manager Software

With a complete picture of the storage groups, message store, and logs used by SQL Server, GSS consultants customize report templates so that only relevant performance metrics are displayed by Tuning Manager software. The reports show configuration and performance data for the SQL Server host and the storage subsystem.

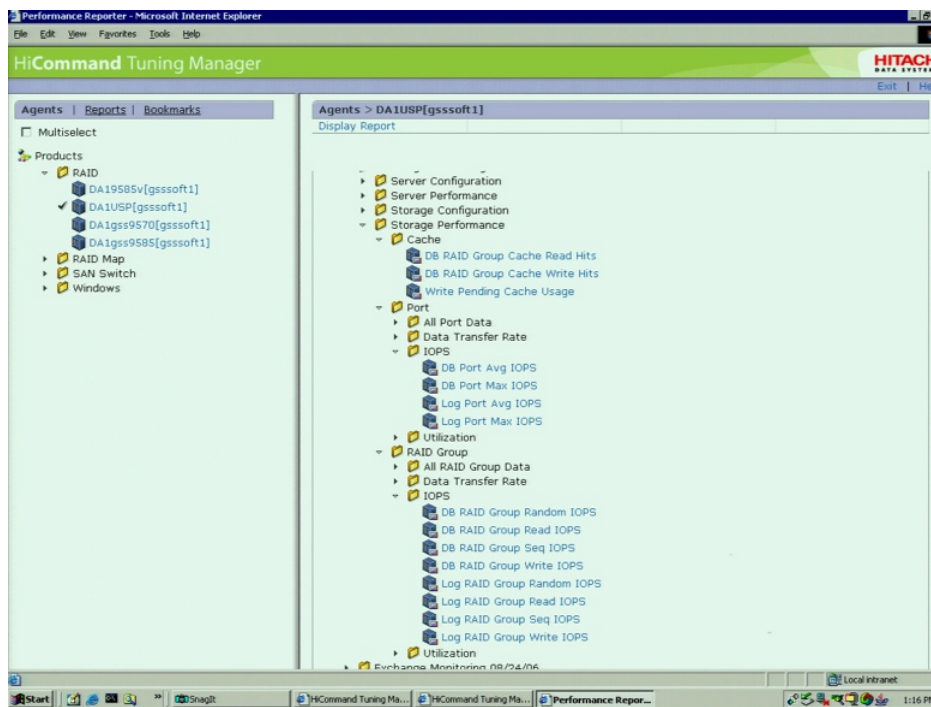


Figure 5. Hitachi Tuning Manager software performance reports.

A browser-based interface to Tuning Manager software supports fast drill-down through the performance data to highlight specific problem areas. GSS consultants use a flowchart-based methodology when performing the analysis to further speed identification of problems.

## Performance Analysis Flowchart

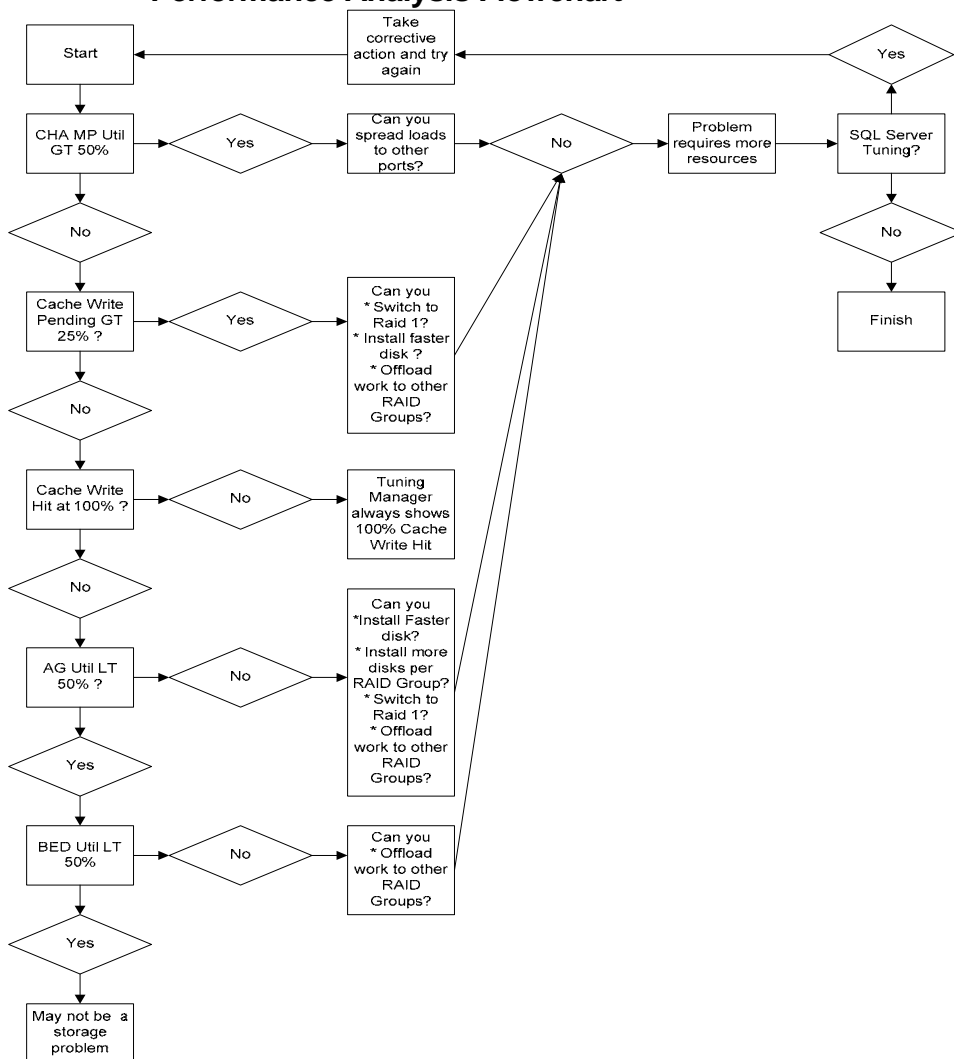


Figure 6. Hitachi Data Systems Global Solution Services SQL Server Performance Analysis Flowchart.

The performance analysis flowchart offers a guide to diagnosing problems in the SQL Server environment. The Tuning Manager user interface supports a step-through approach that enables the GSS specialist and administrators to quickly drill down into the performance metrics to determine the cause of a problem.

The first step of the analysis is to examine channel processor (CHP) microprocessor (MP) utilization. The CHP MP Busy metric must be below 50 percent to accommodate failover in an active-active port balanced configuration. Note that, each port is connected to a CHP and each CHP is connected to multiple ports. SQL Server configurations will almost certainly use multiple paths to the storage subsystem spread across multiple ports and multiple CHPs. This allows for load balancing and failover in case of path outage. Spreading the I/O workload across more ports can alleviate high CHP utilization.

The Cache Write Pending (CWP) metric must be viewed in conjunction with the Side File (SF) usage. The processes measured by these two metrics use the same cache pool and together, they must never be allowed to approach 70 percent. Although CWP can get high during load operations, in general, it must be kept below 25 percent. SF usage must not exceed 10 percent. If CWP utilization is high it is necessary to consider moving to RAID-1+0 array groups and to spread the I/O load across more array groups.

Array Group (AG) utilization can indicate problems when I/O requests are experiencing excessive waits. The AG Busy metric must be kept below 50 percent to avoid waiting. Excessive AG Busy is addressed by one of several options: installing faster disk drives; configuring more disk per RAID group to spread individual disk I/O; adding RAID groups; and switching to alternative RAID configurations, for example from RAID-5 to RAID-1+0.

Back End Disk (BED) utilization can indicate problems when I/O requests are experiencing excessive waits because multiple high activity RAID groups are allocated on the same BED. The BED Busy metric must be kept below 50 percent to avoid waiting. Excessive BED Busy is addressed by one of several options: spreading RAID groups to other BEDs, installing faster disk drives; switching to alternative RAID configurations, for example from RAID-5 to RAID-1+0.

### SQL Server Rules of Thumb and Best Practices

In addition to customized Tuning Manager software reports, GSS consultants have developed specific rules of thumb and best practices to help administrators configure the SQL Server storage environment and measure performance. Although differences between environments will always produce individual results, the key metrics identified by GSS consultants provide a starting point when analyzing SQL Server performance problems.

Note that these metrics are for an Online Transaction Processing (OLTP) type of SQL Server query. OLTP workloads consist mainly of processing application transactions. Examples include: order entry systems, online banking, airline reservations, and eBusiness. A profile of an OLTP transaction is

Read/Write ratio: from 1:1 to 4:1

I/O type: Random / Sequential Access

Operation	Random / Sequential	Read / Write	Size Range
OLTP – Log	Sequential	Write	Sector Aligned Up to 60K
OLTP – Log	Sequential	Read	Sector Aligned Up to 120K
OLTP – Data (Index Seeks)	Random	Read	8K
OLTP - Lazy Writer	Random	Write	Any multiple of 8K up to 256K
OLTP - Checkpoint	Random	Write	Any multiple of 8K up to 256K
Read Ahead (DSS, Index/Table Scans)	Sequential	Read	Any multiple of 8KB up to 256K
Bulk Insert	Sequential	Write	Any multiple of 8K up to 128K
CREATE DATABASE	Sequential	Write	512KB (SQL 2000) , Up to 4MB (SQL2005) (Only log file is initialized in SQL Server 2005)
BACKUP	Sequential	Read/Write	Multiple of 64K (up to 4MB)
RESTORE	Sequential	Read/Write	Multiple of 64K (up to 4MB)
DBCC – CHECKDB	Sequential	Read	8K – 64K
ALTER INDEX REBUILD - replaces DBREINDEX	Sequential	Read	Any multiple of 8KB up to 256K

(Read Phase)				
ALTER INDEX REBUILD - replaces DBREINDEX (Write Phase)	Sequential	Write	Any multiple of 8K up to 128K	
DBCC – SHOWCONTIG (deprecated, use sys.dm_db_index_physical_stats)	Sequential	Read	8K – 64K	

The metrics described below are measured by Tuning Manager Custom Reports for SQL Server. They can be used to determine whether SQL Server Databases and Logs are meeting best practice rules of thumb. The Database and Log metrics are separated since an OLTP transaction typically is composed of 5 to 10 percent Sequential IO processing for Log versus Random IO processing for Database. Performance characteristics will be different for Sequential versus Random IO processing.

Metric Name	Type	Description	Normal Value	Bad Value	View
I/O Rate	IOPS	I/Os per second	Varies		Storage
Read Rate	Read IOPS	I/Os per second			Storage
Write Rate	Write IOPS	I/Os per second			Storage
Read Block Size	Avg Disk Bytes/Read	Bytes transferred per I/O	multiple of 8192	N/A	Host
Write Block Size	Avg Disk Bytes/Write	Bytes transferred per I/O	multiple of 8192	N/A	Host
Database Read Response Time	Avg Disk Sec/Read	Read I/O response (Millisecond)	4 to 8	> 20	Host
Database Write Response Time	Avg Disk Sec/Write	Write I/O response (Millisecond)	4 to 8	> 20	Host, Port
Log Read Response Time	Avg Disk Sec/Read	Read I/O response (Millisecond)	1 to 5	> 5	Host, Port
Log Write Response Time	Avg Disk Sec/Write	Write I/O response (Millisecond)	1 to 5	> 5	Host, Port
Average Queue Length	Avg Disk Queue Length	Average Number of disk requests queued for execution on one specific LUN	< 2 per disk in array	> 2 per disk in array	Host (Max Qdepth may vary with different hardware settings)
Read Hit Ratio	Read Hit %	% of Read I/Os satisfied from Cache	25 to 50	< 25	RAID Group
Write Hit Ratio	Write Hit %	% of Write I/Os satisfied from Cache	100% (0% on USP)	< 100%	RAID Group
Average RAID Group Utilization	RAID Group Utilization Rate	Percentage of the RAID Group Utilization	1% to 50%	> 50%	RAID Group
Average Port Processor Utilization	Port Processor Utilization Rate	Percentage of the Port Processor Utilization	1% to 50%	> 50%	Port
Average Write Pending	Write Pending Rate	Percentage of the Cache used for Write Pending	1% to 25%	> 25%	Cache

Figure 7. Hitachi Data Systems GSS Rules of Thumb for Microsoft SQL Server.

## Summary

With a software tool, such as Hitachi Tuning Manager, and customized reports based on best practices, businesses can do an effective job managing their SQL Server applications. Using the drill-down methodology to analyze reported data will significantly ease the identification of trouble areas when diagnosing performance problems. The reporting service ensures that customers gain the maximum benefit from Tuning Manager software in the shortest amount of time possible.

To learn more about how Hitachi Data Systems can help with SQL Server applications performance analysis and to read more about Tuning Manager software and the custom reporting service, please visit [www.hds.com](http://www.hds.com) or call Hitachi Data Systems to explore an engagement that will result in the optimal solution for your SQL Server storage needs. (Hitachi Data Systems Channel Partners should contact their Channel Managers for information.)

## Appendix A

### Microsoft SQL Server Performance Monitor Counters

In addition to the storage system performance metrics, the Microsoft SQL Server 2005 performance metrics are directly integrated into Microsoft Windows Perfmon when SQL Server is installed on the physical server. There are many object and counters for Microsoft SQL Server to choose from and it can be confusing when trying to decide what to look at and what not to for metrics. This whitepaper contains some of the more important counters to keep an eye on when characterizing performance of a Microsoft SQL Server environment.

[http://www.microsoft.com/technet/prodtechnol/sql/bestpractice/performance\\_tuning\\_waits\\_queues.msp](http://www.microsoft.com/technet/prodtechnol/sql/bestpractice/performance_tuning_waits_queues.msp)

## Appendix B

### Storage Top 10 Best Practices

Published: October 17, 2006. <http://www.microsoft.com/technet/prodtechnol/sql/bestpractice/storage-top-10.msp>

Proper configuration of IO subsystems is critical to the optimal performance and operation of SQL Server systems. Below are some of the most common best practices that the SQL Server team recommends with respect to storage configuration for SQL Server.

1. Understand the IO characteristics of SQL Server and the specific IO requirements/characteristics of your application.

In order to be successful in designing and deploying storage for your SQL Server application, you need to have an understanding of your application's IO characteristics and a basic understanding of SQL Server IO patterns. Performance monitor is the best place to capture this information for an existing application. Some of the questions you should ask yourself here are:

What is the read versus write ratio of the application?

What are the typical IO rates (IO per second, MB/s and size of the IOs)? Monitor the Perfmon counters:

1. Average read bytes/sec, average write bytes/sec
2. Reads/sec, writes/sec

3. Disk read bytes/sec, disk write bytes/sec
4. Average disk sec/read, average disk sec/write
5. Average disk queue length

How much IO is sequential in nature, and how much IO is random in nature? Is this primarily an OLTP application or a Relational Data Warehouse application?

To understand the core characteristics of SQL Server IO, refer to [SQL Server 2000 I/O Basics](#).

## 2. More/faster spindles are better for performance

Ensure that you have an adequate number of spindles to support your IO requirements with an acceptable latency.

Use filegroups for administration requirements such as backup/restore, partial database availability, etc.

Use data files to “stripe” the database across your specific IO configuration (physical disks, LUNs, etc.).

3. Try not to “over optimize” the design of the storage; simpler designs generally offer good performance and more flexibility.

Unless you understand the application very well avoid trying to over optimize the IO by selectively placing objects on separate spindles.

Make sure to give thought to the growth strategy up front. As your data size grows, how will you manage growth of data files/LUNs/RAID groups? It is much better to design for this up front than to rebalance data files or LUN(s) later in a production deployment.

## 4. Validate configurations prior to deployment

Do basic throughput testing of the IO subsystem prior to deploying SQL Server. Make sure these tests are able to achieve your IO requirements with an acceptable latency. SQLIO is one such tool which can be used for this. A document is included with the tool with basics of testing an IO subsystem. Download the [SQLIO Disk Subsystem Benchmark Tool](#).

Understand that the of purpose running the SQLIO tests is not to simulate SQL Server’s exact IO characteristics but rather to test maximum throughput achievable by the IO subsystem for common SQL Server IO types.

IOMETER can be used as an alternative to SQLIO.

## 5. Place log files on RAID-1+0 (or RAID-1) disks. This provides:

better protection from hardware failure, and

better write performance.

Note: In general RAID-1+0 will provide better throughput for write-intensive applications. The amount of performance gained will vary based on the hardware vendor’s RAID implementations. Most common alternative to RAID-1+0 is RAID-5. Generally, RAID-1+0 provides better write performance than any other RAID level providing data protection, including RAID-5, but with Hitachi Universal Storage Platform, Adaptable Modular Storage, and Workgroup Modular Storage, RAID-5 should be considered because of the effectiveness of Write Caching.

## 6. Isolate log from data at the physical disk level

When this is not possible (for example, consolidated SQL environments) consider I/O characteristics and group similar I/O characteristics (for example, all logs) on common spindles.

Combining heterogeneous workloads (workloads with very different IO and latency characteristics) can have negative effects on overall performance (for example, placing Exchange and SQL data on the same physical spindles).

#### 7. Consider configuration of TEMPDB database

Make sure to move TEMPDB to adequate storage and pre-size after installing SQL Server.

Performance may benefit if TEMPDB is placed on RAID-1+0 (dependent on TEMPDB usage).

For the TEMPDB database, create one (1) data file per central processing unit (CPU), as described in #8 below.

#### 8. Lining up the number of data files with CPUs has scalability advantages for allocation intensive workloads.

It is recommended to have .25 to 1 data files (per filegroup) for each CPU on the host server.

This is especially true for TEMPDB where the recommendation is 1 data file per CPU.

Dual core counts as two (2) CPUs; logical processors (hyperthreading) do not.

#### 9. Don't overlook some of SQL Server basics

Data files should be of equal size—SQL Server uses a proportional fill algorithm that favors allocations in files with more free space.

Pre-size data and log files.

Do not rely on AUTOGROW, instead manage the growth of these files manually. You may leave AUTOGROW ON for safety reasons, but you should proactively manage the growth of the data files.

#### 10. Don't overlook storage configuration bases

Use up-to-date host bus adapter (HBA) drivers recommended by the storage vendor

Utilize storage vendor specific drivers from the HBA manufactures website

Tune HBA driver settings as needed for your IO volumes. In general driver specific settings should come from the storage vendor. However we have found that Queue Depth defaults are usually not deep enough to support SQL Server IO volumes.

Ensure that the storage array firmware is up to the latest recommended level.

Use multipath software to achieve balancing across HBAs and LUNs and ensure this is functioning properly

Simplifies configuration and offers advantages for availability

Microsoft Multipath I/O (MPIO): Vendors build Device Specific Modules (DSM) on top of Driver Development Kit provided by Microsoft.

Hitachi Data Systems offering is HiCommand Dynamic Link Manager software