



Best Drive Configuration Practices for PTC Windchill on Microsoft SQL Server

Production-proven content and process management software

*Microsoft Corporation
Published: May 2010*

Author: Ken Lassesen

Reviewers: Tim Atwood (PTC), Victor Gerdes (PTC), David Bernbeck (PTC), Richard Waymire (Solid Quality Mentors)

Abstract

PTC Windchill is the only Product Lifecycle Management (PLM) solution designed from the ground up to work in an Internet-based, distributed-design environment. Whether you need core product data-management capabilities, optimization of processes to meet industry-specific requirements, or support for global product development, Windchill has the capabilities you need for effective management of global product development teams. Microsoft® SQL Server® provides an ideal database platform for Windchill. With SQL Server as a foundation, Windchill can further reduce the time and costs related to managing product development.

This white paper provides best practices for configuring and running Windchill on the SQL Server database platform. The information in this paper complements the detailed support documentation provided on the PTC support Web site. Implementing these best practices can help you avoid or minimize common problems and optimize the performance of Windchill on SQL Server so that you can effectively manage your resources, reduce operating expenses, increase productivity, and improve employee satisfaction.

Microsoft®

Copyright Information

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication. This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2010 Microsoft Corporation. All rights reserved.

Microsoft, SQL Server, Hyper-V, MSDN, and Windows are trademarks of the Microsoft group of companies. All other trademarks are property of their respective owners.

Table of Contents

OVERVIEW	5
INTRODUCING PTC WINDCHILL: A COMPLETE SOLUTION FOR PLM.....	5
SQL SERVER: AN ENTERPRISE-READY DATABASE PLATFORM FOR WINDCHILL.....	6
PTC AND MICROSOFT: A HISTORY OF INNOVATION.....	6
BETTER TOGETHER: WINDCHILL ON THE SQL SERVER DATABASE PLATFORM	6
BEST DRIVE PRACTICES FOR INSTALLING WINDCHILL ON SQL SERVER	7
RAID FOR DATA PROTECTION	7
SQL SERVER REQUIRES AT LEAST 2 PHYSICAL DRIVES.....	8
MIRROR, MIRROR ON THE WALL, WHO IS THE SMARTEST OF THEM ALL?	9
DRIVE CONFIGURATIONS FOR PERFORMANCE.....	10
PHYSICAL HARD DRIVES.....	12
<i>Getting Some Hard Numbers about Your Drives.....</i>	<i>12</i>
DRIVE ALLOCATION HEURISTIC TO ELIMINATE BOTTLENECKS	15
STORAGE AREA NETWORKS (SANS)	16
INTERNET SMALL COMPUTER SYSTEM INTERFACE (ISCSI)	16
TEMPDB ON SOLID STATE DRIVES.....	17
SATA vs. SCSI.....	18
<i>Bus Speeds</i>	<i>18</i>
ECONOMICAL SQL SERVER BOX FOR WINDCHILL	19
HIGH-END CONFIGURATION	21
SUMMARY	21
KEY RECOMMENDATIONS	21
LINKS FOR FURTHER INFORMATION	23
BIBLIOGRAPHY.....	24
APPENDIX: DRIVE ASSIGNMENTS FOR LIMITED LOGICAL DRIVES.....	26
ADEQUATE LOGICAL DRIVES – 8 OR MORE	26
RECOMMENDATIONS FOR LIMITED LOGICAL DRIVES.....	26

FIGURES

Figure 1 - The Policies Tab 13
Figure 2 - Disk Management Table..... 13
Figure 3 - Example of SQLIO Commands to Test a Drive..... 13
Figure 4 - Intel® Matrix Storage Console - RAID Management 20

TABLES

Table 1 - Standard RAID Types 8
Table 2 - Windchill Logical Files for Medium Size Installation..... 11
Table 3 - Recommended SQLIO Tests..... 12
Table 4 - Sample Average Latency Times (Milliseconds)..... 14
Table 5 - Sample MBs/sec Throughput 14
Table 6 - Operational Size of Windchill 15
Table 7 - SAN Configuration Option for HBA..... 16
Table 8 - Performance of Different Drive Types..... 17
Table 9 - Performance of Disk Rotation Speed..... 18
Table 10 - Data Transfer Speed by Bus..... 18
Table 11 - Economical Server Drive Assignments 20

Table A - Suggested Configuration for 7 Logical Drives..... 27
Table B - Suggested Configuration for 6 Logical Drives..... 27
Table C- Suggested Configuration for 5 Logical Drives..... 27
Table D - Suggested Configuration for 4 Logical Drives 28
Table E - Suggested Configuration for 3 Logical Drives 28

Overview

For manufacturers, there is constant pressure today to compete in rapidly changing markets. Companies must be innovative in keeping costs low while maintaining high-quality results. The stakeholders involved in the product development process have also changed. Teams today extend far beyond the central engineering department to include globally dispersed cross-functional groups working on hundreds of products with thousands of parts. The result is an enormous amount of data.

To stay ahead of the global competition, companies find that they must create a collaborative environment that brings together engineering, manufacturing, marketing, and sales teams. Product Lifecycle Management (PLM) solutions can help manufacturers achieve this collaboration, while also streamlining operations and keeping costs down.

PLM is the process of managing all phases of product development—from initial concept through end of life. Effective PLM combines information, methodology, and available resources for each phase of a product's lifecycle, improving a manufacturer's ability to respond swiftly and effectively to changes, new markets, and competitors.

Introducing PTC Windchill: A Complete Solution for PLM

PTC Windchill provides a complete family of solutions for content and process management, helping manufacturers efficiently control all information assets related to product development while optimizing associated business processes. Windchill is the only PLM solution that was designed from the ground up to work in an Internet-based, distributed-design environment.

Windchill technology forms a solid foundation for a variety of packages that PTC offers to address data, change, configuration, and process management; product development collaboration; project management and execution; and the release of product information to manufacturing management systems. For example,

- **Windchill PDMLink** consolidates scattered islands of product content into a single information source, which can help bring order to chaotic product development processes such as change management and speed the development of new product configurations.
- **Windchill ProjectLink** creates a virtual workspace that becomes the central access point for a project, enabling team members to collaborate with access to the same information. By automating project management activities, Windchill ProjectLink helps customers better manage all of their programs, project schedules, information, and processes.

As an integral component of PTC's Product Development System (PDS), Windchill manages all product content and business processes throughout a product's lifecycle. Windchill connects seamlessly to Pro/ENGINEER for three-dimensional (3-D) computer-aided design (CAD) models, ProductView® for advanced mock-up and interactive visualization, Mathcad® for engineering calculations, and Arbortext® for dynamic publishing.

SQL Server: An Enterprise-Ready Database Platform for Windchill

Microsoft SQL Server provides an ideal database platform for Windchill. SQL Server is a high-performance, integrated database and business intelligence (BI) solution for data management and analysis. This easy-to-implement, easy-to-support foundation provides a multifunctional solution for large-scale online transaction processing (OLTP), data warehousing, and e-commerce applications and a solution for data integration, analysis, and reporting.

SQL Server can help companies manage large volumes of mission-critical data and run software applications—such as PTC Windchill—to optimize their business performance. SQL Server can extract and transform data from a variety of sources, including XML data files, flat files, and relational data sources, and then load it into one or more destinations. In addition to rapid data mining, analysis, processing, and reporting capabilities, SQL Server has built-in features that give you a secure, reliable, and productive data management environment that truly protects your data. With its scalable infrastructure, SQL Server has the capability to grow with your business and keep up with your toughest data challenges.

PTC and Microsoft: A History of Innovation

PTC and Microsoft deliver complementary product development solutions that organizations can use broadly across their infrastructure. With a Microsoft IT infrastructure, you get an open, extensible platform and a simplified user experience. You can take advantage of your existing IT investments for a lower total cost of ownership (TCO). PTC provides full integration with your Microsoft infrastructure, including Microsoft® Office SharePoint® Server, Microsoft® Office Project Server, Microsoft® Office Communications Server, Windows Server®, and SQL Server. The PDS architecture gives you end-to-end solutions for product development and a single source for product and process knowledge.

Together, PTC and Microsoft deliver powerful product development solutions that provide tremendous customer value.

Better Together: Windchill on the SQL Server Database Platform

Running Windchill on SQL Server delivers measureable value by channeling data into manageable, automated processes. This decreases administrative time, improves productivity, reduces costs, and generates greater employee satisfaction. Benchmarking tests confirm that SQL Server scales to meet the performance needs of even the largest enterprise customers, while providing lower initial costs and licensing fees.¹

Benchmark test results showed that Pro/ENGINEER and Windchill 9.1 performed up to 50 percent faster on SQL Server 2005 than on a competitor's database, with an average performance advantage of approximately 10 percent. These results confirm SQL Server 2005 as a superior database choice for Windchill 9.1.

¹ "Microsoft SQL Server: Improved Performance for PTC Windchill," http://www.cio.com/white-paper/588485/Microsoft_SQL_Server_Improved_Performance_for_PTC_Windchill

Best Drive Practices for Installing Windchill on SQL Server

Best drive practices covers the recommended ways to configure and use combinations of physical drives (Seagate, Western Digital, etc.), virtual drives (SAN, RAID), and logical drives (C:, D:) for Windchill on SQL Server. This paper focuses on the following critical goals:

- Protecting data
- Optimizing performance
- Minimizing downtime if Windchill is unavailable due to a catastrophic failure of hardware, software, and/or human-ware.

There may be trade-offs between these aspects, especially if there are limited financial or skilled human resources. Most of this paper covers recommendations across all scenarios; however, we also summarize best practices specifically for the following installation categories:

- The small installation (typically fewer than 100 users, but could apply for up to 800 users) where money is tight—also known as the pedal-yourself bicycle scenario.
- The large installation (more than 5000 users with a strong IT department) where there are a lot of resources available for making the use of Windchill a success—also known as the chauffeur-driven limousine scenario.

RAID for Data Protection

At one time, Redundant Array of Independent Disks (RAID) was exclusive to serious enterprise-class storage solutions. Today, you can go to a Big Box or high street Web site and buy a “prosumer” computer that has RAID as part of the motherboard. If you are running old machines but a modern operating system (OS) such as Windows Server 2008 or Windows 7, you will find that RAID is built into the OS—a mainstream feature available for no extra cost.

RAID provides data protection or better performance, or both. There are three concepts involved with RAID drives:

- **Mirror:** Data is written to two or more drives, so each has a mirror copy.
 - 10GB of data requires at least 20GB of storage.
 - Data access is the speed of one drive (or slightly slower).
 - If a drive fails, the other drive retains and provides the data.
- **Striped:** Data is written by splitting it between two or more drives, so less data is written to or read from each drive.
 - 10GB of data requires 10GB of storage.
 - Data access is at half the speed (or better) of one drive, depending on the number of drives.
 - If a drive fails, data is lost.
- **Parity:** Data is written in a way that lets it be effectively mirrored without consuming as much space.

- 10GB of data may require just 12GB of storage across six or more drives or 16GB of storage with three drives.
- Data access may be half the speed (or better) of one drive depending on the number of drives. Write time is much higher because parity must be calculated.
- If a drive fails, data is not lost.
- With double parity (RAID 6), two drives failing will not result in data loss.

You can find a detailed explanation of RAID on Wikipedia.org² or in most SQL Server administration books. As a quick review, Table 1 defines the four common types of RAID available. RAID requires an array of two or more physical drives. A non-RAID drive is typical for a system with just one physical drive.

Table 1 - Standard RAID Types

RAID Type	Description
0	Used to improve performance. Data is split across multiple drives. If any drive fails, all data is lost. Also called <i>striped drives</i> .
1	Used to protect data. Data is written to different drives at the same time so if one drive fails, the data still exists on the other drive. Also called <i>mirrored drives</i> .
5	Used to protect data. Uses three drives to protect the data while consuming less space than mirroring (RAID 1). Data still exists if one drive fails. Space savings is achieved by extra computational load.
6	Used to protect data. This is an extension of RAID 5 that requires more drives but allows recovery if two or more drives fail.

Our discussion will focus on RAID 1, RAID 0, and their composite (RAID 1+0 or RAID 0+1). Using RAID 5 and RAID 6 is complex unless performed by hardware. If using the OS to provide redundancy, use RAID 1; RAID 5 or 6 will take cycles away from the CPU, reducing performance but saving disk space.

SQL Server Requires at Least 2 Physical Drives

Regardless of whether or not you use RAID, the minimum recommended number of physical drives for SQL Server in production is two. A physical drive may fail, so you need to have another drive available after a failure in case of the following scenarios:

- If the database remains available but the backup and transaction logs are lost, you can recover the backup and transaction logs by:
 1. Replacing the drive
 2. Creating a new full backup
 3. Creating a new transaction log

² http://en.wikipedia.org/wiki/Standard_RAID_levels, <http://en.wikipedia.org/wiki/RAID>, http://en.wikipedia.org/wiki/Nested_RAID_levels

- If you have the backup and transaction logs, you can recreate the database at the point of failure by:
 1. Replacing the drive
 2. Recreating the database from the last full backup
 3. Applying the differential backups to bring it almost up-to-date
 4. Applying the transaction log(s) to bring it up-to-date
- If you are running RAID 1, with each drive in the array having a complete image of everything, you can:
 1. Replace the drive
 2. Rebuild the new drive with the image of the mirror drive

For example, consider the scenario for a single physical drive failure. After a failure, you will need to:

- Replace the drive and either
 - Restore the last disk image from external backup, or
 - Reinstall and rebuild by
 - Reinstalling the OS and SQL Server
 - Reinstalling the application
 - Applying all of the updates and security patches
- Restore the database from external media copies of
 - Database backups
 - Transaction logs
- Persuade all users to recreate or resubmit all changes since the last backup or transaction log
 - In some cases, the users because of cumulative changes must perform the changes in the same sequence.

This last type of failure can be a career-ending event.

Mirror, Mirror on the Wall, Who Is the Smartest of Them All?

Running RAID 1 on two drives means that if one drive fails, the system will still be up with no loss of data. Overnight, you can replace the drive. The new drive will rebuild (i.e., copy an image from the good drive), and the system may be back in mirror operation when work starts the next morning. The system will be offline only for as long as it takes to replace the drive—and you get to choose when this will happen.

A question that is sometimes asked is, *“If I have mirrored drives, can’t I just forget about doing backups?”* The answer is no, and here is why:

- If there is a power failure (let’s say someone trips over the cord that connects the hardware to an uninterruptable power supply), you might have perfect mirror images of two corrupted databases. You would need to restore from backup and apply the transaction log to recover.
- If the physical machine is destroyed, you need an off-machine backup and perhaps an off-site backup to restore operations.

Recommendation: All physical drives should be redundant (mirrored or parity—RAID 1, RAID 5, or RAID 6).

With today's systems, including consumer machines, it is possible to have hot-swappable drives. *Hot swappable* means drives may be replaced without having to turn off the machine. The cost to convert a consumer system to hot-swappable Serial Advanced Technology Attachment (SATA) drives is less than \$120³ for three drives. The main advantage of using hot-swappable drives with redundant drives is the elimination of down time. You can immediately replace a bad drive when it fails without taking the system down.

Suggestion: All drives should be hot-swappable.

Drive Configurations for Performance

You need to consider multiple layers of drive configuration as you seek top performance. A few of the key layers are:

- Physical performance of the drive
- Number of drives in a striped or parallel array (RAID-0)
- Locations of various database files

There is a balance required between the number of physical drives in parallel creating a logical drive and the number of logical drives you have. Parallel drives increase throughput. A computer having fewer than eight (8) logical drives will have increased drive-head contention and latency.

Caveat: We are assuming that data redundancy or mirroring is in place.

To illustrate this scenario, consider four (4) physical drives on a system:

- If all of the drives are placed into one logical drive, data to and from the drives will flow extremely fast. But because data is written to the database, the transaction log, tempdb, and the tempdb log, the drive heads may be jumping all over the surfaces of the drives, resulting in high data transfer rates and high latency.
- If one drive is allocated for each of the following—databases, transaction log, tempdb, and tempdb log—the drive heads will require less movement (especially for the log files, which are sequentially written so there are no random movements). The data transfer rate may be only 25% of that in the above case, but because of the low latency, the actual performance may be better.

Finding the best combination is part art and part science. This issue is made a little more complex because just a one logical drive may consist of multiple physical drives, a database may have many filegroups, and filegroups may be a set of database files on multiple logical drives.

³ <http://www.startech.com/category/parts/serial-ata/sata-storage/list.aspx>

This means that one record could be stored on one logical drive and the next record may be on a different logical drive.

Filegroups are useful because they can reduce latency by distributing data across different logical drives. The key word is *can*—this benefit comes only when logical drives are appropriately set up.

Windchill has four filegroups: WCAUDIT, PRIMARY, BLOBS, and INDX. The name of the filegroup indicates its contents:

- **WCAUDIT** – Contains records of who modified what. These are *sequential writes* with few reads.
- **PRIMARY** – Contains records holding general information but not documents; these records are indexes. These are random reads and writes.
- **BLOBS** – Contains records that are large, consisting of documents, and not indexed. These are sequential with more reads than writes.
- **INDX** – Containing the index records, searched to find records. These are random with a high read-to-write ratio.

In addition to these application filegroups, there are three more filegroups used at the server or database level:

- **Tempdb** – A system database serving as the work area for SQL Server, characterized by random reads and writes.
- **Windchill Transaction Log** – The log of changes made to all of the Windchill filegroups, characterized by sequential writes.
- **Tempdb Transaction Log** – The log of changes made in tempdb, characterized by sequential writes.

The Windchill installation program will install these filegroups on the same drive by default. **Error! Reference source not found.** shows a summary of the default installation for medium sized installations. In this section, we will look at distributing these files to different drives to improve performance.

Table 2 - Windchill Logical Files for Medium Size Installation

Database Files Logical Name	Initial Size	Activity Type	Volume	Primary Activity	Filegroup
WcAdmin ⁴	1GB	Random	High	R/W	PRIMARY
WcAdmin_blobs	1GB	Random	Low	R/low W	BLOBS
WcAdmin_indx	1GB	Random	Medium	R/W	INDEX
WcAdmin_wcaudit	240MB	Sequential	Low	W	WCAUDIT
WcAdmin_log	1GB	Sequential	Low	W	

Physical Hard Drives

In general, obtain the best performing drives that you can afford. Performance is typically a function of rotations per minute (RPM), cache size, interface cards, connection, and drivers. The performance for a RAID array is often more dependent on the RAID card than on the hard drives. If not all of the drives are identical, evaluate each drive's performance before allocating them to logical drives.

Getting Some Hard Numbers about Your Drives

If you have five (5) or fewer logical drives, using the drives for different database files usually results in the best performance, rather than striping the drives into arrays, because latency is usually the biggest performance killer. The recommended analysis is to verify the performance of drives and then assign the fastest drives to tempdb, the tempdb log, and the database transaction log (in this sequence).

The recommended tool for determining drive performance before installing Windchill is Microsoft's [SQLIO](#). You should execute the tests in Table 3 on each drive.

Table 3 - Recommended SQLIO Tests

Purpose	Description
Test Sequential Write	This is the behavior that logging and auditing exhibits.
Test Random Read	This is the behavior typically seen for retrieving data.
Test Random Write	This is the behavior when data changes.

Before you run the tests, make sure you:

- Turn off write caching on the drives⁵ through the Properties dialog box of the OS's drive. Clicking the Properties button on the Hardware tab opens a dialog box like that shown in Figure 1, where you turn off write caching.
- Turn off write caching on the controllers.

⁴ WcAdmin is used to indicate the name of the Windchill database, which may vary across installations.

⁵ <http://support.microsoft.com/default.aspx/kb/234656?p=1>; some drives may not honor this.

- Verify that all drives are formatted as NTFS. PTC recommends the use of the NTFS format only. The disk-management table in Figure 2 - Disk Management Table shows the drives' format types.

Figure 1 - The Policies Tab

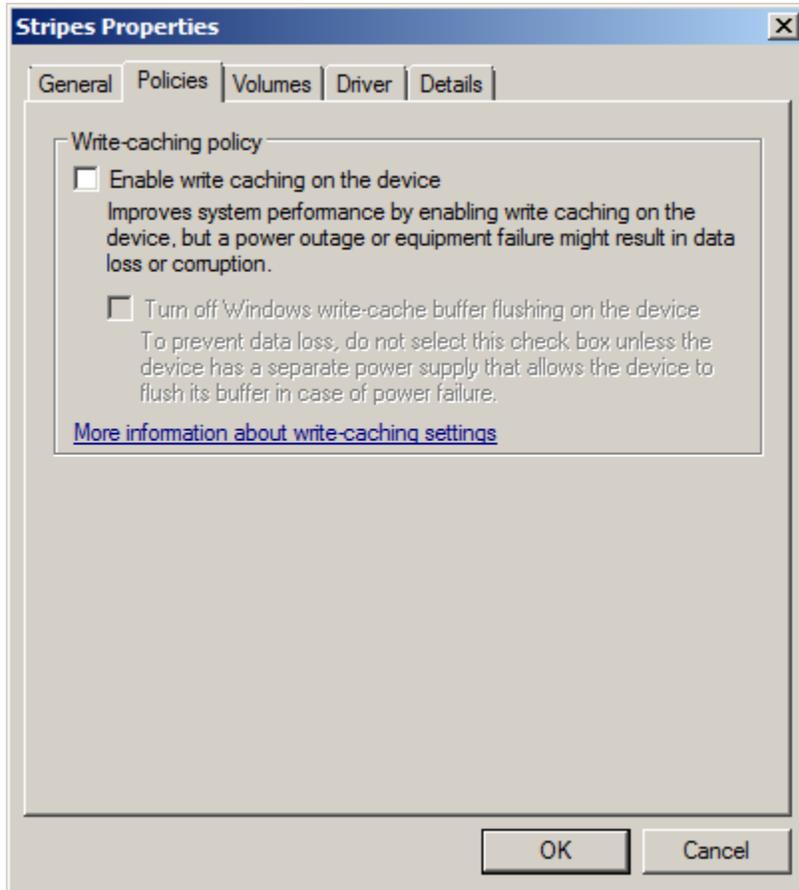


Figure 2 - Disk Management Table

Volume	Layout	Type	File System	Status	Capacity	Free Space
(C:)	Mirror	Dynamic	NTFS	Healthy (Boot, Page File)	465.76 GB	255.89 GB
eSATA1 (H:)	Simple	Dynamic	NTFS	Healthy	784.22 GB	386.13 GB
Local Disk (D:)	Simple	Basic	NTFS	Healthy (System, Active, Primary Partition)	455.49 GB	356.56 GB
New Volume (E:)	Simple	Basic	NTFS	Healthy (Primary Partition)	298.09 GB	249.63 GB
New Volume (O:)	Simple	Basic	NTFS	Healthy (Primary Partition)	465.76 GB	50.44 GB
Recovery (F:)	Simple	Basic	NTFS	Healthy (Primary Partition)	10.27 GB	6.02 GB

Figure 3 shows an example of a command file to test performance of drive M:, with the output redirected to a text file on the E: drive. SQLIO documentation explains the command options.

Figure 3 - Example of SQLIO Commands to Test a Drive

```

sqlio -kR -s360 -frandom -o32 -b128 -LS timeout /T 60 M:\testfile.dat >E:\M_randomRead.txt
sqlio -kW -s360 -frandom -o32 -b128 -LS timeout /T 60 M:\testfile.dat >E:\M_randomWrite.txt
sqlio -kW -s360 -fsequential -o32 -b128 -LS timeout /T 60 M:\testfile.dat >E:\M_SeqWrite.txt

```

We used SQLIO to test some physical drives, whose labeling suggested that similar performance was to be expected, and received the results⁶ shown in Table 4 and Table 5.

Table 4 - Sample Average Latency Times (Milliseconds)

Drive	Sequential Write	Random Read	Random Write
C (SSD)	90	52	134
D (SATA)	143	297	576
E (SATA)	211	337	873
F (SATA)	256	376	923
K (SATA)	49	137	319
L (SATA)	60	160	319
M (SATA)	263	490	1152

Table 5 - Sample MBs/sec Throughput

Drive	Sequential Write	Random Read	Random Write
C (SSD)	176.7	302.1	116.4
D (SATA)	123.4	142.4	23.4
E (SATA)	75.4	47.2	18.1
F (SATA)	100.6	80.9	25.6
K (SATA)	320.8	115.5	49.4
L (SATA)	261.4	99.2	48.9
M (SATA)	37.7	10.9	8.5

As you can see, drives with the same RPMs and cache size may exhibit very different behavior. The hardware manufacturer's specifications are technically correct but do not reflect expectations or actual behavior when being used by SQL Server. If you want the best performance from your Windchill system, you need to test each drive by using a tool such as SQLIO. This applies to bare drives, RAID arrays, iSCSI, and SANs.

For more information about using SQLIO, see:

- [Examples of using SQLIO to perform drive analysis](#)
- [SAN Performance Tuning with SQLIO](#)
- [Disk Performance Hands On , Part 1: Test Harness and Results](#)
- [Webcast: Storage Performance Testing with SQLIO](#)

⁶ These are actual results from a consumer-grade box with consumer-grade drives.

Drive Allocation Heuristic to Eliminate Bottlenecks

We can use the information from SQLIO to eliminate poor assignments that could cause bottlenecks. The allocation of drives should follow this heuristic pattern:

- Because transactions cannot happen faster than they can be logged, the lowest latency *sequential write* drive(s) should be assigned to **tempdb log**, with preference being given to the highest throughput metrics.
- The lowest latency remaining drive(s) for *random read and write* should be assigned to **tempdb**, with preference being given to the highest throughput metrics.
- The lowest latency remaining *sequential write* drive(s) should be assigned to the transaction log (**WcAdmin_log**), with preference being given to the highest throughput metrics.
- The lowest latency remaining drive for *sequential write* should be assigned to **WcAdmin_wcaudit**, with preference being given to the highest throughput metrics; audit trails are similar to logs in behavior.
- The lowest latency remaining drive(s) for *random read and write* should be assigned to **WcAdmin_indx**, with preference being given to the highest throughput metrics.
- The lowest latency remaining drive(s) for *random read and write* should be assigned to **WcAdmin**, with preference being given to the highest throughput metrics.
- The remaining drive(s) should be assigned to **WcAdmin_blobs**, with preference being given to the highest throughput metrics.
- Backups should be assigned to whatever is left over (or to network locations).
- There is no need for a Windows swap file (assuming adequate RAM).

The goal of this heuristic is to minimize head movement and thus reduce latency. The number of drives assigned should be sized to provide the storage space anticipated (see Table 6 - Operational Size of Windchill). The WcAdmin_blobs filegroup will likely require the greatest amount of storage over time. If you have fewer than eight (8) logical drives available, see Appendix: Drive Assignment for suggested assignments with fewer drives.

Table 6 - Operational Size of Windchill

Installation Size	Data Volume
Very Large	> 500GB
Large	300GB-500GB
Medium	100GB-300GB
Small	4GB-100GB

Storage Area Networks (SANs)

A SAN is often the preferred storage medium because it is the easiest scale-out solution to implement⁷. A SAN Logical Unit Number (LUN) is equivalent to a RAID unit or a logical drive. A SAN may consist of hundreds of individual drives, and thus the issue of moving drive heads and latency may disappear. On the other hand, a SAN is a big black box of complex technology, which often means that what is expected may not be delivered.

The above SQLIO testing approach allows the performance of each LUN to be determined. The SAN administrator may argue that such testing is not needed (just as hardware manufacturers would argue that testing the physical hard drives is not needed). But LUNs performing differently have been documented⁸.

For more information about SANs, see

- [Best Practices for Optimizing SQL Server in SAN Environments Using Altiris Products](#)
- [SAN Best Practices](#)
- [SQL Server on SANs](#)

If you are using SAN drives, some special settings affect performance. In general, the Host Bus Adapter is defaulted to too low a value. Use [SQLIO.exe](#) to find the optimal setting for your hardware, as Table 7 summarizes. For more information, see Mike Ruthruff's [SQL Server Best Practices Article](#).

Table 7 - SAN Configuration Option for HBA

Database Option	Default Value	Recommended Value
Host Bus Adapter (HBA) Queue Depth	8-32	64 (Varies by hardware)

Internet Small Computer System Interface (iSCSI)

iSCSI drives are similar to SAN drives but require some care in environmental configuration. When you deploy SQL Server in an iSCSI environment, Microsoft recommends that you use appropriate caution because of latency and bandwidth issues⁹. The iSCSI drives should be:

- On at least a 1000 BaseT (1GB) network with all network interface cards (NICs) supporting at least 1GB
- On an isolated network
- For every three (3) cores used for Windchill, you should have an additional NIC card:
 - 1-3 cores: 1 NIC

⁷ <http://msdn.microsoft.com/en-us/library/aa479364.aspx>

⁸ http://sqlserverpedia.com/wiki/Using_SQLIOSIM_to_isolate_DB_corruption

⁹ <http://support.microsoft.com/kb/833770>

- 4-6 cores: 2 NICs
- 7-9 cores: 3 NICs and a 10GB network

For additional information, see:

- [Installing and Configuring Microsoft iSCSI Initiator](#)
- [Fiber Channel and iSCSI Performance Comparison for DSS Workloads Using SQL Server 2005](#)
- [Deployments and Tests in an iSCSI SAN](#)

Tempdb on Solid State Drives¹⁰

The performance measurements of different drive types from Table 8 below, along with the disk assignment heuristic, point to using a Solid State Drive (SSD) for tempdb if your budget permits. There are three types of SSDs available:

- Multi-Level Cell flash memory (MLC) – consumer grade
- Single-Level Cell flash memory (SLC) – enterprise grade
- Dynamic Random Access Memory (DRAM)-based drive

Only DRAM-based drives are suitable for tempdb. The flash-memory SSD suffers from a low expected number of lifetime writes and slow write speed compared to read speeds (as we saw above). A DRAM-based drive uses the same type of memory that computer systems use and does not maintain the information in the event of a power loss. Because tempdb is recreated whenever SQL Server starts, this loss of information is acceptable. Table 8 illustrates the reported differences between Hard Disk Drives (HDD), flash-memory SSD, and DRAM SSD.

Table 8 - Performance of Different Drive Types¹¹

Technology	Capacity (GB)	Latency (milliseconds)
Flash SSD (SLC)	700	0.200
7200 HDD	700	7
DRAM SSD	250	0.003

Recommendation: The equivalent performance of an SSD may be obtained from suitably large striped RAID 0 arrays. Flash SSD is not recommended as being cost- or performance-effective. Only DRAM SSD should be considered.

¹⁰ SSD drives should be mirrored; a failure of the tempdb drive would immediately stop SQL Server.

¹¹ <http://www.storagesearch.com/ssd-ram-v-flash.html> retrieved 2010-03-11 and <http://www.storagesearch.com/soliddata-art2-comparisons.pdf>

SATA vs. SCSI

SATA drives are cheap and may be purchased almost anywhere. SCSI drives are expensive. Both are being used in SANs, which are strictly enterprise-class components. For example, for around \$30,000¹² HP is offering today either a:

- 12TB SATA starter SAN or
- 4.8TB SAS [SCSI] starter SAN

The common reasons that a SCSI drive array may be required include:

- Performance cannot be obtained from SATA drives (see Table 9 - Performance of Disk Rotation Speed).
 - SATA drives typically rotate at 5200-7200RPM.
 - SCSI drives typically rotate at 10000-15000RPM.
- The cost of any transactions being lost is high.
- The forecasted throughput volume is greater than what SATA drives can handle.
- Heavy loads are expected.
 - SATA tends to be light-duty engineered (consumer grade).
 - SCSI tends to be heavy-duty engineered (enterprise grade).
- Lower drive failure rates are desired.

The bottom line is that SCSI can cut latency by up to 75% and double the data transfer speed over SATA.

Table 9 - Performance of Disk Rotation Speed¹³

Speed (RPM)	Rotational Latency (milliseconds)	Track-to-Track Latency	Seek Time	Data Transfer Rate (MB/sec)
5400	5.5	6.5	12	90
7200	4.1	6.5	10.7	120
10000	3	1.5	4.5	166
15000	2	1.5	3.5	250

Bus Speeds

Table 10 illustrates the main differences between SCSI and SATA. Most modern SCSI drives support fiber optic interfaces.

Table 10 - Data Transfer Speed by Bus

¹² <http://www.hp.com/hpinfo/newsroom/press/2009/090211b.html>

¹³ *Professional Microsoft SQL Server 2008 Administration* (Wrox, 2009), Brian Knight et al. p 438

Bus (not in Array)	Maximum Transfer Speed (MB/sec)
SATA 300 – consumer/prosumer	300
SATA 600 – On market in 2010	600
SATA 150 – last generation	150
USB 1	1.5
USB 2	60
SCSI-320	320
Fiber Optic (available on SCSI)	2000
SAN on fiber optics networks	10000

Recommendation: Always use SCSI drives in RAID arrays or SANs.

In reality, for small organizations, the cost¹⁴ of using SCSI in RAID arrays or SANs may be prohibitive. If you must use SATA drives, consider using SATA-600 drives with RAID controllers that support SATA-600. In the next section, we look at a suggested reasonable solution when you are unable to follow the best recommendations.

For more information, see

- [Storage Top 10 Best Practices](#)
- Mike Ruthruff's [Predeployment I/O Best Practices](#)

Economical SQL Server Box for Windchill

This section describes a pragmatic configuration for Windchill when it is not financially feasible to follow the recommended best practices. It is an alternative and not a recommendation.

In today's mid-2010 consumer market, you can purchase the following hardware configuration for around \$1,000 at COSTCO.com or other sites:

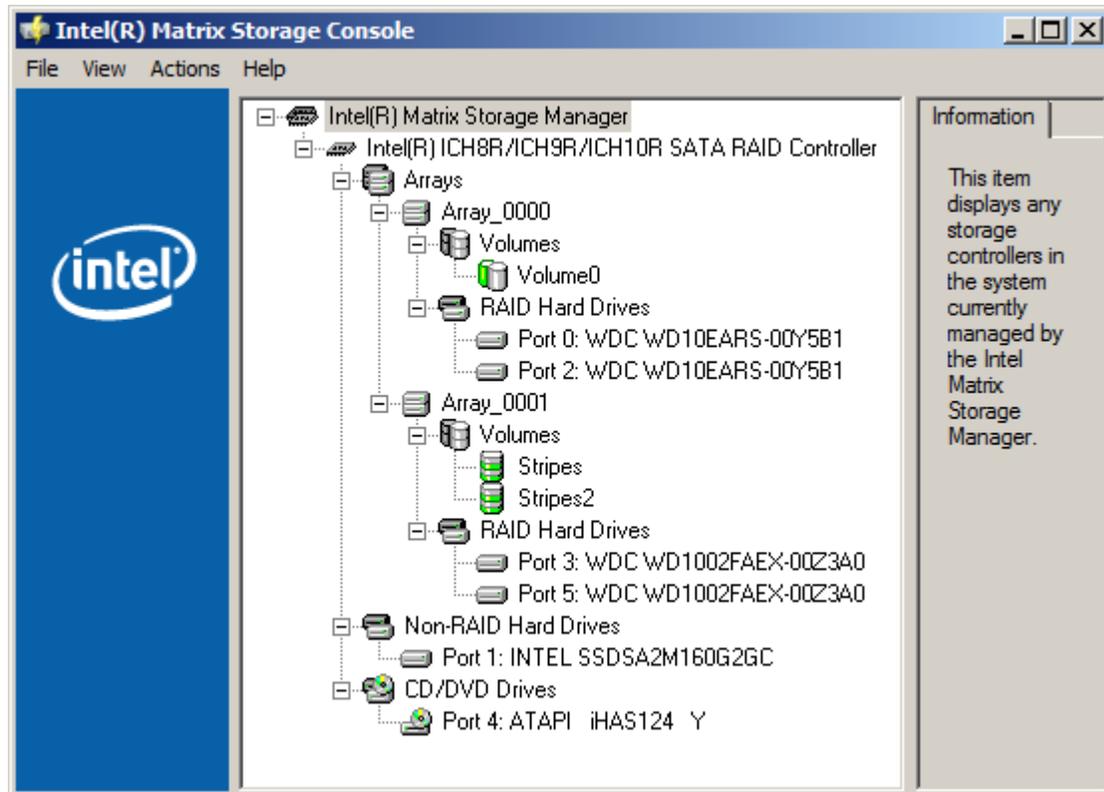
- 4-core processor such as the Intel I-7 920 with hyper-threading
 - 8 CPUs appear in Task Manager
- Motherboard with a RAID controller on it, such as Intel's Matrix Storage
- 8 separate SATA connectors that may be used by RAID
- eSATA connector
- Additional IDE connectors for DVDs, tape backups, etc.
- 12GB of DDR3 memory

This hardware configuration places itself at the 800-user level on the sizing table, according to "Windchill and Pro/INTRALINK 9.0 and 9.1 Server Hardware Sizing Guidelines – Microsoft Windows Platform."

¹⁴ Costs include capital cost of purchasing system, training support staff, and ongoing system monitoring and maintenance.

If you bring up the RAID interface in such a configuration, as **Error! Reference source not found.** shows, you'll see that it is running its own controller.

Figure 4 - Intel® Matrix Storage Console - RAID Management



This gives you the ability to install additional SATA drives (at \$100 each for 1TB) and create four (4) logical drives that consist of two (2) mirrored drives. If you are good with a screwdriver, you may want to make all eight (8) drives hot-pluggable by buying appropriate enclosures. For \$60, you can even toss in an extra 4-port RAID SATA card¹⁵. Go for two of these cards with mirroring, and you have eight (8) logical drives—the ideal minimum number. You can also mount these physical drives in aftermarket hot-swappable SATA holders.

With only eight (8) SATA connectors, it is not possible to do RAID 0 striping without giving up mirroring. Adding a 4-port RAID SATA-600 card exclusively for tempdb (configure the card for RAID 1+0) gives a striping performance boost to Windchill. The assignments of our four (4) logical drives would be as **Error! Reference source not found.** shows.

¹⁵ See <http://www.startech.com/item/PCISATA4R1-4-Port-Serial-ATA-RAID-01-PCI-Card.aspx> for purposes of illustration only. Not endorsed by PTC or Microsoft.

Table 11 - Economical Server Drive Assignments

Drive	Contents
1	Tempdb (database) Operating System
2	INDX (filegroup) BLOBS (filegroup) Tempdb Log (database log)
3	PRIMARY (filegroup)
4	WCAUDIT (filegroup) WcAdmin Log (database log)
ATA/IDE Drives	Backup Medium

Although this is not an ideal configuration, it minimizes data loss risk and maximizes performance on a tight budget. (And don't forget battery backup with automatic server shutdown.)

High-End Configuration

A high-end configuration would place all of the filegroups on a SAN, with each filegroup having its own LUN. The SAN will provide for both redundancy and striping (typically across many drives). Because SANs typically start at \$30,000, dedicated staff usually maintains it.

Summary

Running PTC on SQL Server provides a firm foundation for your mission-critical PLM activities. PTC Windchill supports essential product engineers who cannot do their work if SQL Server is not available. It is critical that contingency planning be at the forefront of your installation planning. At a minimum, it is strongly recommended that hard drives containing database files be redundant (mirrored or parity) on hot-swappable drives. If one drive fails, there is no loss of data nor is there a need to bring SQL Server down to replace the drive.

Hardware failures happen, and there are two common paths for recovery:

- Using an appropriate RAID configuration with hot-pluggable drives. You simply promptly replace the failed drive, with the new drive rebuilding dynamically from the other drives. The database is not offline.
- Selecting the SQL Server full recovery model. You replace the failed drive and recover the database. The database is offline until recovered, which may be hours or days depending on the size of the database and transaction logs.

The first path costs a little more in capital purchases, but the second path has a high lost-productivity cost across many users.

Key Recommendations

Here is a brief summary of the key drive recommendations from PTC for running Windchill on the SQL Server database platform:

- All drives should be redundant (either by mirroring or by parity).
- The use of SSD drives is not recommended.
- SCSI drives in RAID arrays or SANs should be used if possible.
- All drives used for SQL Server must have write caching turned off.
- Each filegroup should be on a separate logical drive or LUN.
- Use the drives with the lowest latency and highest throughput that you can afford.

By following these best practices for your drive configurations, you can minimize the risk of data loss and down time while optimizing the performance of your Windchill and SQL Server implementation.

Links for Further Information

PTC Windchill

- [Microsoft SQL Server: Improved Performance for PTC Windchill](#)
- [Windchill and Pro/INTRALINK 9.0 and 9.1 Server Hardware Sizing Guidelines – Microsoft Windows Platform](#)

SQL Server information can be found in Books Online:

- [SQL Server 2008 Books Online](#)
- [SQL Server 2005 Books Online](#)

SQL Server Books Online also includes best practice information in the following articles:

- [Best Practices for Replication Administration](#)
- [Replication Security Best Practices](#)
- [Best Practices for Recovering a Database to a Specific Recovery Point](#)

See the [SQL Server Best Practices](#) portal for technical white papers, the SQL Server Best Practices Toolbox, Top 10 Lists, and other resources.

Following is a list of technical white papers that were tested and validated by the SQL Server development team. These can help you learn more about specific SQL Server topics.

- [A Quick Look at Serial ATA \(SATA\) Disk Performance](#)
- [Best practices for operational excellence](#)
- [Best Practices for Running SQL Server on Computers That Have More Than 64 CPUs](#)
- [Best Practices for Semantic Data Modeling for Performance and Scalability](#)
- [Checklist: SQL Server Performance](#)
- [Checksum Problems, Choosing the Correct Recovery Model and More](#)
- [Comparing Tables Organized with Clustered Indexes versus Heaps](#)
- [Database Mirroring and Log Shipping Working Together](#)
- [Database Mirroring Best Practices and Performance Considerations](#)
- [Database Mirroring in SQL Server 2005](#)
- [Database Snapshot Performance Considerations Under I/O-Intensive Workloads](#)
- [DBCC SHOWCONTIG Improvements and Comparison between SQL Server 2000 and SQL Server 2005](#)
- [Description of using disk drive caches with SQL Server that every database administrator should know](#)
- [Disk Partition Alignment Best Practices for SQL Server](#)
- [FLASH Disk Opportunity for Server-Applications](#)
- [How to mirror the system and boot partition \(RAID1\) in Windows Server 2003](#)

- [How to use the SQLIOSim utility to simulate SQL Server activity on a disk subsystem](#)
- [How To: Use SQL Profiler](#)
- [Identifying and Resolving MDX Query Performance Bottlenecks in SQL Server 2005 Analysis Services](#)
- [Implementing Application Failover with Database Mirroring](#)
- [Improving SQL Server Performance](#)
- [Microsoft SQL Server 2005 Tuning Tips for PeopleSoft 8.x](#)
- [Microsoft SQL Server Database Engine Input/Output Requirements](#)
- [Microsoft SQL Server I/O Basics \(2005\)](#)
- [Microsoft SQL Server I/O subsystem requirements for the temp db database](#)
- [Monitor and troubleshoot storage performance](#)
- [Partial Database Availability](#)
- [Performance of WD 250GB SATA Drives + 3ware Controller](#)
- [Physical Database Storage Design](#)
- [Precision Considerations for Analysis Services Users](#)
- [Pre-Configuration Database Optimizations](#)
- [Pre-deployment I/O Best Practices](#)
- [RML Utilities for SQL Server \(x64\)](#)
- [SAP with Microsoft SQL Server 2005: Best Practices for High Availability, Maximum Performance, and Scalability](#)
- [SQL Server 2005 Best Practices Analyzer \(August 2008\)](#)
- [SQL Server 2005 Deployment Guidance for Web Hosting Environments](#)
- [SQL Server 2005 Performance Dashboard Reports](#)
- [SQL Server 2005 Security Best Practices - Operational and Administrative Tasks](#)
- [SQL Server 2005 Waits and Queues](#)
- [SQL Server 7.0, SQL Server 2000, and SQL Server 2005 logging and data storage algorithms extend data reliability](#)
- [SQL Server Best Practices](#)
- [SQL Server Best Practices Article](#)
- [SQL Server Health and History Tool \(SQLH2\)](#)
- [SQL Server Replication: Providing High Availability Using Database Mirroring](#)

Bibliography

- ***Professional SQL Server 2005 Performance Tuning*** (Wiley Publishing, 2008), Steven Wort, Christian Bolton, et al.
- ***SQL Server MVP Deep Dives*** (Manning, 2010), Paul Nielsen, Kalen Delaney, et al.
- ***Inside Microsoft SQL Server™ 2005: The Storage Engine*** (Microsoft Press, 2006), Kalen Delaney

- ***Professional SQL Server 2005 Administration*** (Wiley, 2006), Brian Knight, Ketan Patel, et al.
 - ***SQL Server 2005 Bible*** (Wiley, 2006), Paul Nielsen
-

Appendix: Drive Assignments for Limited Logical Drives

The recommended minimum number of logical drives is eight (8). However, not all installations can follow that best practice. This Appendix shows how to make drive assignments if you have fewer than eight (8) logical drives. The tables below assume that the logical drives have been tested with SQLIO for random read/write, sequential write, and sequential reads.

Adequate Logical Drives – 8 or More

The location of the operating system (OS) is probably immaterial, and with a machine dedicated to SQL Server, a swap file will not be active.

1. Assign the fastest random R/W to **Tempdb**. Otherwise, it will bottleneck the reading of data to it from other drives (you want to avoid data arriving faster than it can be handled).
2. Assign the fastest sequential write to **Tempdb Log**. You cannot do things faster than you can log them!
3. Assign the fastest remaining sequential write to **WcAdmin Log**. You cannot change the data in the database faster than you can log the changes.
4. Assign the fastest remaining random R/W to **INDX**. A record cannot be changed until all the indexes associated with it are updated.
5. Assign the fastest remaining random R/W to **PRIMARY**. These are generally short records that are often changed.
6. Assign the fastest remaining random R/W to **BLOBS**. These are generally large files that are read/written sequentially from the disk. You may want to run a test for sequential read/write on the drives if time permits.
7. Assign the fastest remaining sequential write to **WCAUDIT**. The audit log is used less than either the WcAdmin Log or Tempdb Log.
8. Assign the fastest remaining sequential write to **Backup**. Backups are done by sequential writing to disk.

Recommendations for Limited Logical Drives

The following tables show the recommended drive configurations if you have fewer than eight (8) drives. The drive selection follows the pattern above for the first content in each set.

Table A - Suggested Configuration for 7 Logical Drives

Drive #	Contents
1	Tempdb (database)
2	Tempdb Log (database log) Backups (operating system)
3	WcAdmin Log (database log)
4	INDX (filegroup)
5	PRIMARY (filegroup)
6	BLOBS (filegroup)
7	WCAUDIT (filegroup)

Table B - Suggested Configuration for 6 Logical Drives

Drive #	Contents
1	Tempdb (database)
2	Tempdb Log (database log) Backups (operating system)
3	WcAdmin Log (database log)
4	INDX (filegroup) BLOBS (filegroup)
5	PRIMARY (filegroup)
6	WCAUDIT (filegroup)

Table C - Suggested Configuration for 5 Logical Drives

Drive #	Contents
1	Tempdb (database)
2	Tempdb Log (database log) Backups (operating system) WcAdmin Log (database log)
3	INDX (filegroup) BLOBS (filegroup)
4	PRIMARY (filegroup)
5	WCAUDIT (filegroup)

Table D - Suggested Configuration for 4 Logical Drives

Drive #	Contents
1	Tempdb (database) Backups (operating system) WcAdmin Log (database log)
2	INDX (filegroup) BLOBS (filegroup) Tempdb Log (database log)
3	PRIMARY (filegroup)
4	WCAUDIT (filegroup)

Table E - Suggested Configuration for 3 Logical Drives

Drive #	Contents
1	Tempdb (database) Backups (operating system) WcAdmin Log (database log)
2	INDX (filegroup) BLOBS (filegroup) Tempdb Log (database log)
3	WCAUDIT (filegroup) PRIMARY (filegroup)