


# Consolidating Microsoft SQL Server on the IBM System x3950 M2

Suggests strategies and techniques to consolidate SQL Server systems

Explains when consolidation is and isn't practical

Helps you prepare for a consolidation activity



David Watts  
Daniel Soares de Barros  
Simon Champion





International Technical Support Organization

**Consolidating Microsoft SQL Server on the IBM  
System x3950 M2**

February 2008

**Note:** Before using this information and the product it supports, read the information in “Notices” on page v.

**First Edition (February 2008)**

This edition applies to SQL Server 2005 running on the IBM System x3950 M2, machine type 7141.

**© Copyright International Business Machines Corporation 2008. All rights reserved.**

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Notices</b> .....	v
Trademarks .....	vi
<b>Preface</b> .....	vii
The team that wrote this paper .....	vii
Become a published author .....	ix
Comments welcome .....	ix
<b>Chapter 1. Overview of SQL Server</b> .....	1
1.1 Common issues with SQL Server environments .....	2
1.2 SQL Server 2005 and server consolidation .....	3
1.2.1 Support for 64-bit computing (x64) .....	4
1.2.2 Other features .....	4
1.3 64-bit computing and SQL Server 2005 .....	8
1.3.1 Windows, SQL Server, and 32-bit versus 64-bit .....	8
1.4 Looking forward to SQL Server 2008 .....	10
1.4.1 New features in SQL Server 2008 .....	10
1.4.2 Migrate to SQL Server 2005 or SQL Server 2008? .....	12
<b>Chapter 2. Consolidation</b> .....	15
2.1 Concepts of consolidation .....	16
2.2 Forms of consolidation .....	16
2.3 Consolidation strategies .....	17
2.3.1 Vertical consolidation: Scale up .....	17
2.3.2 Horizontal consolidation: Scale out .....	18
2.3.3 SQL Server instances: Single or multiple? .....	18
2.4 Considering service availability .....	19
2.4.1 Assessing critical systems .....	19
2.4.2 Ensuring system availability .....	20
2.5 Virtualization technologies .....	24
<b>Chapter 3. IBM System x3950 M2 for consolidation</b> .....	27
3.1 Introduction to IBM System x3950 M2 .....	28
3.1.1 Comparing the x3950 M2 with the x3950 .....	29
3.1.2 IBM XA-64e fourth-generation chipset .....	30
3.1.3 Processors .....	31
3.1.4 System memory .....	33
3.1.5 NUMA .....	35
3.1.6 Broadcom dual Gigabit Ethernet controller .....	37
3.1.7 SAS disk subsystem .....	37
3.1.8 PCI subsystem .....	37
3.1.9 Redundancy .....	38
3.2 Scalability with Windows and SQL Server .....	38
3.3 IBM sizing guide tool .....	39
<b>Chapter 4. Consolidation strategy and implementation</b> .....	43
4.1 The five phases of consolidation .....	44
4.2 Assessment and analysis .....	44
4.2.1 Assessment guidelines .....	45

4.2.2 Sample script to gather information . . . . .	49
4.3 Consolidation project . . . . .	56
4.3.1 Critical success factors . . . . .	57
4.3.2 SQL Server 2005 features for consolidation . . . . .	57
4.3.3 Storage . . . . .	59
4.3.4 Backup . . . . .	61
4.3.5 High availability . . . . .	62
4.3.6 Monitoring . . . . .	67
4.3.7 Online analytical processing and online transaction processing . . . . .	67
4.3.8 SQL Server upgrade . . . . .	68
4.3.9 NUMA Considerations . . . . .	68
4.4 Testing . . . . .	71
4.5 Implementation and stabilization . . . . .	72
4.6 Best practices . . . . .	72
4.7 Capacity management . . . . .	73
4.8 Services offerings . . . . .	75
<b>Abbreviations and acronyms . . . . .</b>	<b>77</b>
<b>Related publications . . . . .</b>	<b>79</b>
IBM Redbooks . . . . .	79
Other publications . . . . .	79
Online resources . . . . .	80
Help from IBM . . . . .	81

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

# Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Redbooks (logo) ®  
eServer™  
xSeries®  
Active Memory™  
BladeCenter®

Chipkill™  
Collation®  
IBM®  
Predictive Failure Analysis®  
Redbooks®

System x™  
System Storage™  
Tivoli®  
X-Architecture®

The following terms are trademarks of other companies:

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

Snapshot, and the Network Appliance logo are trademarks or registered trademarks of Network Appliance, Inc. in the U.S. and other countries.

Advanced Micro Devices, AMD, ATI, ES1000, Radeon, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

Access, Microsoft, SQL Server, Visual Basic, Windows NT, Windows Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Xeon, Itanium, Itanium 2, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.



# Preface

Database servers have become a vital part of the IT infrastructure for organizations of all sizes. This has been fueled by the requirement for those organizations to capture, analyze, and report on an ever increasing volume of data of all types and sizes.

Today we see a number of factors combining to provide cost-effective solutions for every budget. These factors include the comparatively low cost of disk storage and high-capacity memory chips, the availability of mainframe-class processing power on the Microsoft® Windows® Server 2003 platform, and the Microsoft SQL Server™ 2005 relational database management system. However, the proliferation of systems running SQL Server in many organizations, while tactically solving many business problems, has brought with it a number of management issues and costs. This has been accentuated by the growth in importance of the regulatory and compliance controls that have imposed new demands on most businesses around the globe.

The concept of server consolidation is not new and the cost reduction benefits that a physical server consolidation project can bring to an organization are well understood. If implemented appropriately, most organizations can realize real server consolidation, which can deliver actual cost savings while improving business capability.

This paper explores how the server consolidation concept is best applied to a SQL Server environment by suggesting strategies and techniques that apply specifically to SQL Server. We include a script to help you gather some of the data you will need as part of your planning. In addition, we explore the IBM® System x3950 M2 server and explain how it can be an ideal platform for a SQL Server consolidation exercise because of its advanced hardware features and, in particular, its ability to scale from 4-processors/16-cores right up to a 16-processors/64-cores.

The paper is written for technical decision makers, infrastructure architects, solution designers, and technical managers and is positioned between a business overview and a technical user guide.

## The team that wrote this paper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Raleigh Center.

**David Watts** is a Consulting IT Specialist at the IBM ITSO Center in Raleigh. He manages residencies and produces IBM Redbooks® publications on hardware and software topics related to IBM System x™ and BladeCenter® servers and associated client platforms. He has authored over 80 books, papers, and technotes. He holds a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM both in the U.S. and Australia since 1989. He is an IBM Certified IT Specialist.

**Daniel Soares de Barros** is an Advisory IT Specialist for Database Administration in IBM Global Services in Brazil, supporting IBM domestic and international accounts. He has over 9 years of IT experience and his areas of expertise include Application Development, Database Administration and Windows Server® administration. He has passed 25 technical exams and holds 12 certifications for Microsoft, IBM, and Oracle®.

**Simon Champion** is an IT Specialist in UKISA ITDelivery based in the United Kingdom. He has over 19 years of IT experience. He has worked at IBM for 5 years and his areas of expertise include Windows Server consolidation, VMware ESX Server, Windows Server administration, SQL Server implementation, and Windows clustering. Simon has a Higher National Certificate in Software Engineering from Brighton College of Technology and is a Microsoft Certified Professional in Managing and Supporting Windows Server 2003.



*The team (l-r): David, Daniel, and Simon*

Thanks to the following people for their contributions to this project:

From the ITSO:

- ▶ Tamikia Barrow
- ▶ Byron Braswell
- ▶ Carolyn Briscoe
- ▶ Linda Robinson
- ▶ Margaret Ticknor

From IBM Marketing:

- ▶ Craig Alexander
- ▶ Paul Branch
- ▶ Jay Bretzmann
- ▶ Brian Sanders

## Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

[ibm.com/redbooks/residencies.html](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

[ibm.com/redbooks](http://ibm.com/redbooks)

- ▶ Send your comments in an e-mail to:

[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400





# Overview of SQL Server

In this chapter, we discuss some of the problems that organizations face with their older and existing SQL Server environments and introduce some strategies that can be used to address them. We also preview the release of the next edition of the Microsoft relational database management system, SQL Server 2008.

The specific topics covered in this chapter are:

- ▶ 1.1, “Common issues with SQL Server environments” on page 2
- ▶ 1.2, “SQL Server 2005 and server consolidation” on page 3
- ▶ 1.3, “64-bit computing and SQL Server 2005” on page 8
- ▶ 1.4, “Looking forward to SQL Server 2008” on page 10

## 1.1 Common issues with SQL Server environments

Because it is relatively easy to deploy and the initial investment costs are relatively low, the installed base of SQL Server has grown considerably over the years. This growth, while solving many business problems along the way, has created new challenges for organizations of all sizes. Some of these challenges are:

- ▶ **Server sprawl**

The number of SQL Servers has grown organically, often without any central coordination, leading to a vast array of different combinations of hardware, operating systems, and SQL Server versions. Many organizations do not know exactly how many SQL Servers they have, making the management of these systems inconsistent.

- ▶ **Scalability limitations**

Since the release of the Microsoft SQL Server relational database management system (RDBMS), it has gained considerable market share in the small business/departmental database space. However, because of architectural limitations that are present in all 32-bit editions of SQL Server, it cannot scale high enough to support enterprise-level applications with high transaction throughput. One limitation relates to the amount of physical memory SQL Server can use directly for all its functions. Microsoft did release a 64-bit edition of SQL Server 2000 for the Itanium® 64-bit processor, which resolved this limitation, but its deployment was not widespread because of high implementation costs.

- ▶ **Low server utilization**

Average server utilization rates of 5% are commonplace, although many organizations find their servers are running at utilization rates of about 10-15%. Therefore, there are clear opportunities in many environments to consolidate SQL Server workloads onto fewer servers, which can improve the return on investment (ROI) for companies.

- ▶ **Lack of system fault tolerance for critical systems**

Many SQL Server deployments start small and grow in size and importance over time. However, a server environment that might have been suitable for a small, unproven system when it was first installed is not likely to be appropriate after the system has grown into a critical business tool. Many times this is overlooked; important business processes are not being supported by server infrastructures that are capable of recovering from a hardware component failure or server outage. Apart from hardware considerations, SQL Server can be easily configured so that a database can be recovered in the event of server loss or data corruption. Too often, though, we see that the link between business criticality and server systems is not understood and simple measures are not put in place.

- ▶ **Standardization of configuration**

The proliferation and lack of central control has led to variations in server hardware, operating system, and SQL Server configuration. This complicates the already complex task of ensuring that all systems meet the business, operational, and regulatory requirements now expected from server environments.

- ▶ **Meeting initial regulatory compliance**

Often dictated by national government legislation, like the U.S. Government's Sarbanes-Oxley Act of 2002, audit compliance is a top priority in today's computing environments and failure of an audit can have far reaching consequences for organizations. The areas open for audit are generally security patch deployment, user ID management, and server health checking. Maintaining the procedures to support these audit requirements is a continual process in itself and new server deployments must ensure that they meet the necessary audit requirements at the point when they are implemented.

► **Maintaining ongoing regulatory compliance**

To ensure that server systems continue to comply with regulatory audit requirements, there must be strictly adhered to processes for performing periodic reviews of all systems, user IDs, and security patch distribution. Many systems were installed before the current audit compliance specifications were a requirement and these systems also need to be able to pass an audit.

► **System administration**

Enterprise Manager in SQL Server 2000 can administer SQL Server 7.0 and SQL Server 2000 databases; however, versions before SQL Server 7.0 must use their own administration tools. This requires the maintenance of multiple management tools which increases administrative costs.

► **Inconsistencies in administration**

In larger SQL Server environments, there will be more than one, and possibly a team of, database administrators (DBAs). Given the proliferation of tools and human nature, DBAs are not always consistent in how they configure and maintain a database. This can lead to gaps in operational support that can then affect service availability.

► **Backup and recovery**

Because multiple tools are required to manage existing and older SQL Server environments, there is often a lack of consistency in approach and implementation. This can lead to ineffective backups that are of no use in a data restoration scenario.

► **Server maintenance, service packs, and hot-fixes**

The larger the server estate, the greater the effort required to ensure that all systems are kept up-to-date and not exposed to security threats, even with automated deployment tools. If “service window” documentation for each server does not exist, then server maintenance is difficult to perform without unexpected system outages.

► **Hardware configuration and maintenance**

An aging server estate contains a greater risk of hardware failure. If the original hardware was not designed to be fault tolerant (that is, they do not contain redundant hardware components that can keep the server online when a component fails), or if server clustering has not been deployed, there is little or no protection from hardware failure. Spare parts for old servers might not be readily available, and that might lead to an extended outage of the server and the database.

Many of these issues are addressed by a well-planned and well-executed server consolidation exercise. With the environment consolidated, the remainder of the issues can be resolved with proper procedures and controls for maintaining the environment in a managed state.

## 1.2 SQL Server 2005 and server consolidation

Microsoft SQL Server 2005 has many new features, not previously available in SQL Server, that add considerably to its abilities to scale from a small single-user database to a huge, enterprise-wide, mission-critical, multi-user database. This section highlights these features and discusses how they are relevant to server consolidation.

## 1.2.1 Support for 64-bit computing (x64)

The combination of Windows Server 2003 for x64 and SQL Server 2005 for x64 offer directly addressable physical memory up to the memory limit of the operating system: 32 GB for Windows Server 2003 Standard Edition and 1024 GB (1 TB) for Windows Server 2003 Enterprise Edition. This effectively resolves the memory constraint that exists with the 32-bit versions of Windows Server and SQL Server. We describe the advantages that x64 computing brings to server consolidation more fully in 1.3, “64-bit computing and SQL Server 2005” on page 8.

There are several editions of SQL Server 2005 with varying support for x64; however, only these versions are suitable for creating a consolidated SQL Server environment:

- ▶ SQL Server 2005 Enterprise Edition (32-bit and 64-bit)
- ▶ SQL Server 2005 Standard Edition (32-bit and 64-bit)

For medium and large-scale SQL Server consolidation projects, the Standard Edition and Enterprise Edition versions both have native x64 versions; however, many of the advanced scalability features are only found in the Enterprise Edition. Developer Edition has all the features of Enterprise Edition, but is licensed only for development and testing, not for production use.

It is important to note that a database created using SQL Server 2005 Express Edition can be moved to an installation of SQL Server 2005 Enterprise Edition without any modifications. This provides a clear growth strategy for all new databases created with SQL Server 2005 and demonstrates the ease with which databases can be scaled-up on this platform.

In section 4.3.2, “SQL Server 2005 features for consolidation” on page 57, we go into more detail about which scalability features are available in each SQL Server edition.

## 1.2.2 Other features

Most of the features that are mentioned in the following sections are only found in the Enterprise Edition. For a detailed analysis of what is supported by the Standard and Enterprise Editions, see the Microsoft Web document “Comparison Between SQL Server 2005 Standard and Enterprise Editions” at the following location:

<http://www.microsoft.com/sql/editions/enterprise/comparison.mspx>

### Database mirroring

Database mirroring provides online protection for your critical databases by maintaining a secondary copy of the database. SQL Server 2005 detects when the database is off-line and automatically switches over to the mirrored copy in only a few seconds. Failover time is typically less than if a clustered SQL Server instance is used. With database mirroring, you can protect only those databases that have particular value to your business because it works at the database level rather than at the SQL instance level like Microsoft Clustering.

### Data partitioning

Partitioning a large table or index enables partitions, or sub-sets of the data, to be updated or managed independently while maintaining the integrity of the overall table or index. Certain operations that might take many minutes on an unpartitioned database take only seconds to complete on a correctly partitioned database. Querying large tables or indexes that have been partitioned is likely to be faster and more efficient on multiple CPU computers because partitions can be worked on concurrently.



## **Hot-add memory**

Additional physical memory can be installed in a running server, and SQL Server 2005 will recognize and use the additional memory immediately. This could prove useful if you must increase available memory to service new business requirements without affecting database availability. This feature also requires hot-add memory support as provided in servers such as the IBM System x3950 M2.

## **Online restore**

With SQL Server 2005, database administrators can perform a restore operation while the rest of the database remains online and available. Online restoration improves the availability of SQL Server because only the data being restored is unavailable. In a consolidated environment with very large tables, a large system might continue to be available to its users while a small portion of corrupted data is recovered.

## **Online indexing operations**

With the online index option, concurrent modifications (updates, deletes, and inserts) to the table or any associated indexes can occur during index maintenance operations. For example, while a clustered index is being rebuilt, users can continue to make updates to the underlying data and perform queries.

## **Fast recovery**

A new faster recovery option improves the availability of SQL Server databases. Users can reconnect to a recovering database after the transaction log has been rolled forward. It is no longer necessary to wait for the rollback phase to complete.

## **Security enhancements**

SQL Server 2005 includes security enhancements such as database encryption, secure default settings, password policy enforcement, fine-grained permissions control, and an enhanced security model. In a consolidated SQL Server environment where many users must have access to different databases and their tables, security controls are an important part in protecting data integrity.

## **Dedicated administrator connection**

SQL Server 2005 introduces a dedicated administrator connection (DAC) that administrators can use to access a running server even if the server is locked or otherwise unavailable. Administrators can troubleshoot problems on a server by executing diagnostic functions or Transact-SQL (TSQL) statements.

## **Snapshot isolation**

SQL Server 2005 introduces a new snapshot isolation level that is intended to enhance concurrency for OLTP applications. In earlier versions of SQL Server, concurrency was based solely on locking, which can cause blocking and deadlocking problems for some applications. Snapshot isolation depends on enhancements to row versioning and is intended to improve performance by avoiding reader-writer blocking scenarios.

## **Compatibility mode for SQL Server 7.0 and 2000 databases**

With SQL Server 2005, you can attach a SQL Server 7.0 or 2000 database and run it unchanged in “compatibility mode.” SQL Server 2005 can emulate a SQL Server 7.0 or SQL Server 2000 server for the database so that it can be accessed as before. However, there are some limitations to the functionality that is available, but this could be used as a “staging post” during migration to SQL Server 2005. SQL Server 6.0 and 6.5 compatibility modes are supported by this in SQL Server 2005, but there are some issues with 6.0 support.

## Using different SQL Server versions side-by-side

It is possible to install SQL Server 2005, 2000, and 7.0 onto the same physical server and have them run alongside each other. This is, in fact, one of the migration strategies that you can use while moving existing SQL Server databases to SQL Server 2005. However, you can host your existing databases on the same consolidated server hardware without having to make substantive changes to related applications.

## Server Resource Management

Given the availability of server hardware that has large memory capacity, up to 32 processors, and multiple network cards, it is necessary to have control over how those considerable resources are allocated. This section introduces hardware and software features that can provide that control and ensure the most appropriate use of the available resources.

### ► Non-uniform memory addressing (NUMA)

NUMA is a scalability technology for splitting servers with numerous processors (CPUs) and large amounts of memory into resource groups, or NUMA nodes. The processors in a NUMA node work primarily with the local memory in that NUMA node while still having access to memory in other NUMA nodes (remote memory). Using local memory is quicker than remote memory because of the configuration of the NUMA node.

Figure 1-1 shows a generic NUMA node. This one has been configured with four CPUs and has its own memory and I/O interface. You can also see the system interconnect from the memory controller that will make requests to other NUMA nodes when access to remote memory is required.




Figure 1-1 Logical NUMA node (image from SQL Server 2005 Books Online)

Because SQL Server 2005 is NUMA aware, it tries to write data to physical memory that is associated with the requesting CPU to benefit from the better local memory access performance. If the requesting CPU does not have enough memory available, it is allocated from another NUMA node.

### ► Soft-NUMA, CPU affinity, and I/O affinity

Soft-NUMA is a SQL Server 2005 feature that you can use to group CPUs and network interfaces into soft-NUMA nodes. However, you cannot allocate memory to a soft-NUMA node and all memory requests are served from all memory available to SQL Server.

To group CPUs, you must edit the registry directly using a *node configuration affinity mask*. After the soft-NUMA nodes have been created, you can assign individual SQL Server instances to one or more soft-NUMA nodes.

You might create soft-NUMA nodes if your server hardware does not have hardware NUMA capabilities or to sub-divide a NUMA node further. Each soft-NUMA node gets its own I/O thread and lazy writer thread. If the SQL instance has a high I/O requirement, it could be assigned two soft-NUMA nodes. The SQL instance then has two I/O threads that can help it process I/O requests better.

Soft-NUMA provides the ability to fine-tune the use of the server resources to ensure that critical databases get the resources that they require within a consolidated environment.

CPU affinity and I/O affinity are SQL Server 2005 features for configuring each database instance to use specific CPUs for database processing and I/O requests. Assigning a set of CPUs just to handle I/O processing might provide performance benefits with a database that relies heavily on I/O operations.

Designating a certain number of CPUs to a critical database ensures that performance is not affected by other processes running on the same server because those processes are run on other CPUs in the server. CPU and I/O affinity are used for fine-tuning the allocation of server resources to where they are most required.

► **Windows Server Resource Manager (WSRM)**

WSRM comes with Windows Server 2003 Enterprise and Datacenter Editions and can be applied to any application running on the server. Using WSRM policies, it is possible to manage CPU and memory use by process or user. WSRM helps ensure that a process or user does not use more than its allotted quantity of CPU and memory resources meaning that multiple applications can run safely together.

SQL Server 2005 is also able to allocate individual CPUs to a SQL database instance using soft-NUMA and CPU affinity, so care must be taken during the configuration of WSRM to ensure there are no contention issues.

For more information about server resource management see 4.3.9, “NUMA Considerations” on page 68.

### **Database migration tool for SQL Server 7.0 and 2000**

The SQL Server 2005 Upgrade Advisor is a tool that Microsoft offers without charge that can analyze your existing SQL Server 7.0 and 2000 instances and determine their suitability for migration to SQL Server 2005. The Upgrade Advisor highlights feature and configuration changes that might affect an upgrade and provides links to further documentation on the issues identified. We discuss this tool in more detail in 4.3.8, “SQL Server upgrade” on page 68.

### **Database migration tools for Oracle, Sybase, and Access databases**

Microsoft has released SQL Server Migration Assistant (SSMA) tools for Oracle, Sybase Adaptive Server Enterprise (ASE), and Access™ databases. SSMA for Oracle converts Oracle database objects (including stored procedures) to SQL Server database objects, loads those objects into SQL Server, migrates data from Oracle to SQL Server, and then validates the migration of code and data.

SSMA for Access can migrate databases from Microsoft Access versions 97 through 2003 to Microsoft SQL Server 2005. SSMA for ASE can migrate Sybase ASE databases to SQL Server 2005. The SSMA for Access and Sybase ASE convert the database objects to SQL Server database objects, load those objects into SQL Server, and then migrate the data from Access or Sybase ASE to SQL Server.

These tools can be freely downloaded from the following URL:

<http://www.microsoft.com/sql/solutions/migration/default.msp>

## 1.3 64-bit computing and SQL Server 2005

Intel® and AMD™ both offer 64-bit extensions to the long-standing 32-bit x86 architecture. Collectively, these two technologies are referred to as “x64.” Internally, the architectures differ, but they can be considered comparable in what they offer. x64 offers full backwards compatibility for 32-bit applications. For scalability reasons, IBM has chosen to use the Intel 64-bit processor technology in servers such as the IBM System x3950 M2.

Intel also manufactures the Itanium 2™ processor, which is built on an entirely different 64-bit architecture. Itanium 2 is not backwards compatible with x86 applications or cross-compatible with x64 applications. This is because it implements a new instruction set called EPIC, which was designed to be the next-generation processor architecture for mainframe class scalability. However, its adoption has not been widespread and it is not a mass-market technology.

The following sections look at the limitations that are imposed by 32-bit hardware and software, and how 64-bit hardware and software addresses those limitations.

### 1.3.1 Windows, SQL Server, and 32-bit versus 64-bit

SQL Server 2005 is available in three versions: as a 32-bit version, EM64T 64-bit (x64), and Itanium 64-bit versions; however, the IBM System x3950 M2 only runs the 32-bit and x64 versions. Running a 64-bit version of SQL Server also requires a matching version of the operating system.

Table 1-1 shows the combinations of Windows Server 2003 and SQL Server that run on the IBM System x3950 M2.

*Table 1-1 Valid combinations of Windows Server and SQL Server running on the System x3950 M2*

	<b>Windows Server 2003, 32-bit</b>	<b>Windows Server 2003, x64 (EM64T)</b>
<b>SQL Server 2000 32-bit</b>	Supported	Supported
<b>SQL Server 2005 32-bit</b>	Supported	Supported
<b>SQL Server 2005 x64 EM64T</b>	Not valid	Supported

Itanium versions of Windows and SQL Server are not supported on the x3950 M2.

There are three combinations for running 32-bit and x64 versions of both Windows Server 2003 and SQL Server 2005:

- ▶ Both 32-bit
- ▶ x64 Windows and 32-bit SQL Server
- ▶ x64 Windows and x64 SQL Server

#### **Windows and SQL Server, both 32-bit**

When you run 32-bit Windows Server 2003, you can run either 32-bit SQL Server 2000 or 32-bit SQL Server 2005. All processes run in 32-bit, so there is no 64-bit option for SQL Server in this case.

SQL Server 2000 runs as a user process, with the standard 4 GB of address space (see Figure 1-2 on page 9), normally divided into two: 2 GB for the kernel and 2 GB for the user mode portion.

With 4 GB of RAM, you can use the /3GB boot.ini switch to change the split to 1 GB for the kernel and 3 GB for the user portion of the Virtual Address Space (VAS).

If you have more than 4 GB of physical memory installed, 32-bit SQL Server can use it as the database buffer pool. However, you must enable Physical Address Extension (PAE) in Windows (add the /PAE switch to boot.ini) and enable Address Windowing Extensions (AWE) in SQL Server (using `sp_configure`).

**Note:** For systems with more than 16 GB of RAM, you cannot enable both /3GB and /PAE. See Section 9.11 of the IBM Redbook *Tuning IBM System x Servers for Performance*, SG24-5287.

All SQL Server memory objects, binary code, data buffers, database page headers (512 MB for 64 GB of database buffers), sort area, connections, stored procedure caches, open cursors (basically, everything but the unmapped cached database pages) must fit in the 2 GB user mode portion of the process address space. To access the database buffer pool pages above the 4 GB line, SQL Server must map them into the address space below the 4 GB line (Figure 1-2). This mapping incurs some performance overhead. With AWE on, SQL Server 2000 allocates its full “max server memory (MB)” amount and never releases memory until it is shut down.

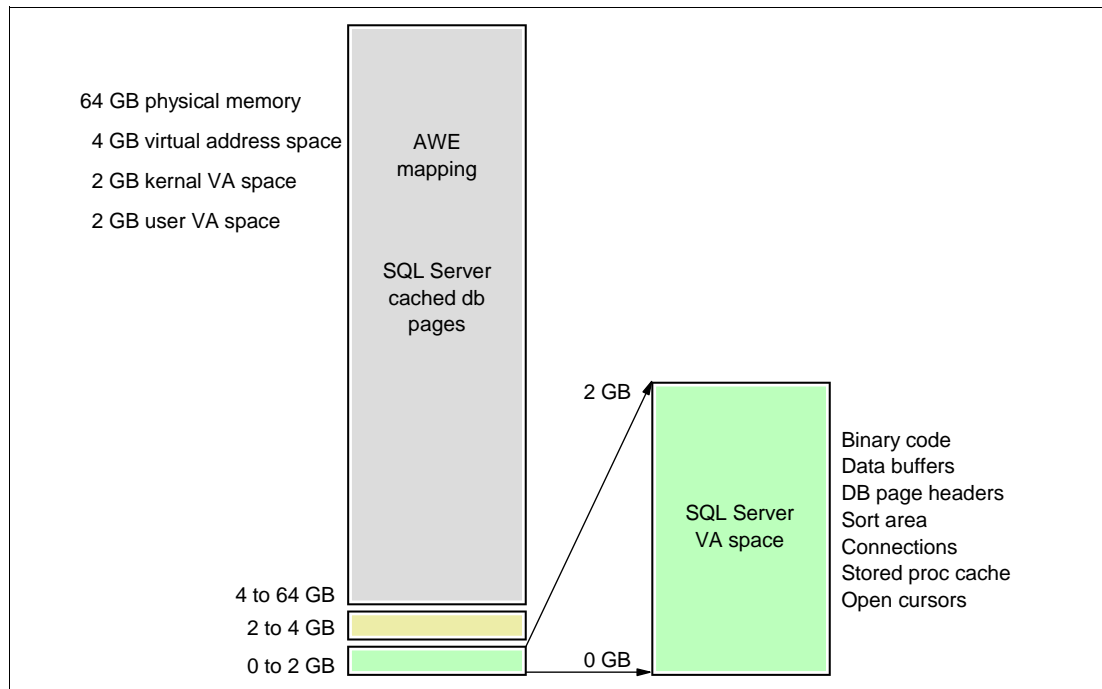


Figure 1-2 SQL Server 2000 uses AWE memory only for buffer pool pages

If your SQL Server workload is constrained by the 2 GB user mode limit for reasons other than requiring more database buffer pages, it will not benefit from having more than 4 GB of physical memory. This is a hardware architectural bottleneck, which is relieved by using 64-bit hardware and software.

SQL Server 2005 32-bit experiences the same hardware architectural bottleneck. SQL Server 2005 manages all of its memory (including the AWE memory) dynamically, releasing and allocating memory in response to internal and external memory pressure.

## **Windows Server x64 and SQL Server 32-bit**

With a Windows Server 2003, x64 Edition installation, you can run either SQL Server 2000 with Service Pack 4 or 32-bit SQL Server 2005.

SQL Server 2000 runs WOW64 (Windows on Windows) in a user mode address space of 4 GB. Unlike running the same application on 32-bit Windows, SQL Server 2000 has a full 4 GB of user mode address space because the 64-bit kernel runs in its own address space. This provides some relief for SQL Server workloads that are constrained by the 2 GB user mode address space of a 32-bit operating system.

SQL Server 2000 can also be configured to use AWE for access up to 64 GB of physical memory on x64 systems. The AWE memory above 4 GB must still be mapped into the lower 4 GB user mode address space to be used.

SQL Server 2005 32-bit benefits in the same way that SQL Server 2000 does with a 4 GB user mode address space that is not shared with the operating system. With AWE enabled, SQL Server 2005 can access up to 64 GB of physical memory on the x3950 M2. This is because the maximum memory for a process running in the WOW64 on Intel EM64T is 64 GB. SQL Server 2005 will manage all of its memory (including the AWE memory) dynamically, releasing and allocating memory in response to internal and external memory pressure.

## **Windows and SQL Server 2005, both 64-bit**

When you run Windows Server 2003 x64, you can (and should) install the x64 version of SQL Server 2005. There is no 64-bit version of SQL Server 2000 for use on EM64T-based servers.

SQL Server 2005 64-bit enjoys the same memory addressability as the 64-bit operating system. The 64-bit user mode address space is not limited to 4 GB, and it can use memory up to the operating system maximum for any purpose, not just for database buffers. AWE mapping is not required because the memory model is flat under 64-bit addressing. This provides the most efficient utilization of resources for SQL Server 2005.

See the IBM Redpaper *Introducing Windows Server x64 on IBM eServer xSeries Servers*, REDP-3982 for more information about 32-bit and 64-bit memory addressing with Windows Server 2003.

# **1.4 Looking forward to SQL Server 2008**

As of November 2007, Microsoft has committed to releasing SQL Server 2008 by the end of the second quarter of 2008. The following sections look at the new features that are slated for inclusion in SQL Server 2008 and cover how its release might influence your migration strategies.

## **1.4.1 New features in SQL Server 2008**

It is clear that, across the board, organizations are experiencing exponential growth in the volume and variety of data that they need to process, analyze, protect, and store. The growing importance of regulatory compliance and increasing globalization dictates that data must be stored securely and be available at all times. Because the costs of disk storage have dropped to record lows, organizations can store more data per dollar. Users must be able to examine and analyze this data quickly and easily on any device using their regular office productivity programs. The management of this information explosion and hike in user expectations creates severe challenges for the enterprise.

Microsoft has positioned their new database platform as the answer to these challenges and we have highlighted some of the new key features here. For a comprehensive review of all the new features in SQL Server 2008, visit:

<http://www.microsoft.com/sql/2008/default.msp>

### **Transparent data encryption**

With SQL Server 2008, it will be possible to encrypt entire databases, data files, or log files without making changes to existing applications. This integrated encryption helps organizations meet the applicable regulatory and business requirements for maintaining data privacy and integrity.

### **Enhanced auditing**

SQL Server 2008 has been enhanced so that activity on your data can be audited. Audit information can be gathered about when data was read and what changes were made. Of particular interest is the fact that you can define audit specifications in each database, which means that the audit configuration moves with the database. Auditing can also be restricted to specific objects so that you can manage the audit overhead and its use of server resources.

### **Database mirroring enhancements**

Database mirroring, a high-availability feature introduced in SQL Server 2005 as an alternative to server clustering, receives some enhancements that improve its reliability, performance, and supportability:

- ▶ The Automatic Page Repair feature relies on keeping the principal and mirror databases in full synchronization. If a page becomes corrupt in the principal database, SQL Server 2008 will be able to read the page from the mirror database and use it to fix the principal database. Similarly, if a page becomes corrupt in the mirror database, the page can be read from the principal database to fix the mirror database.
- ▶ Database mirroring performance has been improved by introducing the compression of the outgoing log stream for minimized network utilization. This helps ensure that the principal and mirror databases are kept synchronized in a shorter time period.
- ▶ The ability to support database mirroring has been improved with additional performance counters that provide finer-grained accounting of the time spent with database log processing. Dynamic Management Views and extensions to existing views reveal new information about database mirroring sessions.

### **Enhanced SQL Server resource management**

With the Resource Governor in SQL Server 2008, administrators can control how resources are consumed in SQL Server. Windows Server Resource Manager (WSRM) provided some ability to control processes within Windows by permitting restrictions on what resources the process `sqlservr.exe` could consume. However, this affected every activity in the SQL Server instance. With Resource Governor, administrators can configure how various workloads can use the available SQL Server-specific resources (only CPU and memory in CTP V5). This is an important feature for SQL Server consolidation because administrators can ensure the best use of available SQL Server resources based on business requirements.

### **Hot add CPU**

Building on the existing support for hot add memory, SQL Server 2008 introduces the installation of additional CPUs into supported server hardware so that you can use the new CPU resources immediately without downtime. This feature extends the ability of SQL Server to scale up the available hardware resources without disrupting the environment.

## **Backup and data compression**

Data and backup compression reduces the storage requirements for databases and database backups. Less I/O is required, which leads to quicker backups and improvements for large I/O-bound workloads, such as those linked with data warehousing.

## **Policy-based management: Declarative Management Framework**

Declarative Management Framework (DMF) is a policy-based management framework for the SQL Server Database Engine. Using policies, you can manage system configuration, monitor and prevent changes to the system, simplify administration tasks, and detect compliance issues in SQL Server Management Studio.

### **1.4.2 Migrate to SQL Server 2005 or SQL Server 2008?**

The answer to the question that this section poses largely depends on the SQL environment you currently have. Therefore, we offer some questions to help you identify the key factors in making this decision. Above all, your decision should be based on your business requirements rather than any technology considerations.

#### **Are your databases critical to the business?**

Most organizations prefer not to deploy mission-critical applications onto new platforms before the first service pack has been released. In the case of SQL Server 2008, the Release To Manufacture (RTM) release is still several months away and the full feature set is not yet available for testing in the Community Technology Preview (CTP) releases.

The first service pack is traditionally available within a few months of the RTM release after the early adopters have had a chance to exercise the software and uncover the new problems. Is your organization willing to accept the risk of using a potentially great, but an as-yet unproven, platform for its critical systems?

#### **Do you explicitly need the new features in SQL Server 2008?**

Are there specific business requirements that can only be met by using some of the new features available in SQL Server 2008? There are certainly some very exciting new features, but it is possible that what you require can still be done by SQL Server 2005, at least for the short term. Therefore, you might be able to wait for Service Pack 1 for SQL Server 2008.

#### **What mix of SQL Server are you currently running?**

Are you running mostly SQL Server 6.x/7/2000 or have you already migrated primarily to SQL Server 2005? If you are already using SQL Server 2005, the likelihood is that this migration was completed fairly recently, so does the feature set of SQL Server 2008 provide a compelling enough business case to migrate again?

If you are mostly running the earlier versions of SQL Server, moving straight to SQL Server 2008 might be a good choice while still considering the release schedule. It is likely to be the third quarter of 2008 before SQL Server 2008 with Service Pack 1 is available, but standard Microsoft support for SQL Server 2000 terminates in April 2008. If you delay upgrading your SQL Server 2000 servers until SQL Server 2008 with Service Pack 1 is available, there will be a gap in support coverage. Is this support gap an acceptable risk for your business?

#### **When do you want to complete your SQL Server migration?**

If you are mostly running SQL Server 6.x/7/2000, how quickly does your organization want to migrate to a Microsoft supported SQL Server environment? If you must move to something supported now, SQL Server 2005 is the only choice. Are you comfortable with a fourth-quarter 2008 or first-quarter 2009 date for a SQL Server 2008 migration?



## Are your applications certified with SQL Server 2008?

Over time, many application vendors will certify SQL Server 2008 for their applications; however, there is normally a gap between the next version release and widespread certification with that new release. Is your business prepared to run applications in a configuration that might not be explicitly supported by your application vendors?

**Note:** For environments consisting mostly of SQL Server 6.x/7/2000, we recommend a migration to SQL Server 2005 as the best strategy for realizing real business savings in the short to medium term. There is huge potential for business cost reductions by deploying SQL Server 2005 for server consolidation and increased performance.

SQL Server 2005 is a mature, well-supported RDBMS platform with numerous features that can offer many benefits over the previous SQL Server versions. However, organizations should begin to make plans for their migration to SQL Server 2008 because Microsoft mainstream support for SQL Server 2005 is currently scheduled to end in 2011.





# Consolidation

This chapter covers the subject of consolidation, the concepts involved, the types of consolidation, consolidation strategies, and considers when virtualization technologies are applicable.

The topics are:

- ▶ 2.1, “Concepts of consolidation” on page 16
- ▶ 2.2, “Forms of consolidation” on page 16
- ▶ 2.3, “Consolidation strategies” on page 17
- ▶ 2.4, “Considering service availability” on page 19
- ▶ 2.5, “Virtualization technologies” on page 24

## 2.1 Concepts of consolidation

The *Compact Oxford English Dictionary* states that the word *consolidation*, a derivative of *consolidate*, means to make stronger or more solid, or to combine into a single unit. The *Cambridge Advanced Learner's Dictionary* suggests it means to combine several things so that they become more effective. Dictionary.com says it is the act of combining into an integral whole. When we apply the word “consolidation” to a computer environment, we are often seeking to improve the current performance, make it more integrated, reduce the physical footprint, and reduce related costs.

When “consolidation” is applied to a SQL Server environment, the desired outcomes are:

- ▶ Reduce the physical number of SQL Servers.
- ▶ Reduce the number of SQL Server instances.
- ▶ Reduce the number of SQL Server databases.
- ▶ Centralize all SQL Servers to one physical location.
- ▶ Centralize SQL Server administration.
- ▶ Improve SQL Server utilization rates.
- ▶ Improve system monitoring and alerting.
- ▶ Implement consistent standards for the whole SQL Server environment.
- ▶ Improve data security and limit user access as the business needs require.

Consolidation is more than just an effort in cost reduction. It is an opportunity to reexamine the needs of your business, to improve the performance and scalability of the environment to meet current and near future needs, and to improve the quality of the whole environment by ensuring that standardization is achieved in all related processes and procedures.

## 2.2 Forms of consolidation

A SQL Server server consolidation exercise involves a combination of five basic forms:

### ▶ **Physical server consolidation**

Physical server consolidation is the act of reducing the actual number of installed physical servers. Because early versions of SQL Server limited a server to a single instance of SQL Server and because it is generally believed that servers cannot be shared safely, many servers are barely used to their full potential. However, after careful analysis of the workload of each server (a requirement), databases can be migrated to refreshed server hardware that is designed to handle the workload being placed on it.

### ▶ **Storage consolidation**

Storage consolidation can involve simply moving the data from a number of servers to a large locally attached disk storage sub-system in a new server. Alternatively, it can also involve the implementation of a centralized disk sub-system that is then shared over a Storage Area Network (SAN) using high-speed network connections. The right solution depends entirely on business requirements and available budget. When physical server consolidation occurs, it follows that there will be some form of storage consolidation.

### ▶ **Geographic consolidation**

Geographic consolidation takes server systems located in different rooms, buildings, cities, or countries and brings them together in a single, or at least reduced, number of sites, which can reduce costs because real estate is expensive. Close attention must be paid to both the server workload and the network requirements. Without sufficient network connectivity and bandwidth, performance could be adversely affected, which might have serious business implications that outweigh any potential real estate cost benefits.

► **Logical server consolidation**

Logical server consolidation relates to the analysis of the SQL Server instances that have been installed on existing SQL servers. Since the release of SQL Server 2000, it has been possible to install multiple instances of SQL Server on the same physical server. Each instance is effectively a logical server because it has its own network interface and server resources. There are only a few situations when a database needs its own SQL Server instance and, in most cases, databases can be moved from their own dedicated SQL Server instances to a SQL Server instance that is shared with many other databases, assuming the workload requirements have been taken into consideration.

► **Database consolidation**

Database consolidation requires in-depth analysis of database designs to determine where databases can be combined, which reduces the number of databases. In practice, this requires a high degree of effort and can require changes to the related application. However, there might be a database that is located in several different countries that shares a common design but contains country-specific data. If these identically structured databases are geographically consolidated into one site, conceivably the databases could be combined and access segregated with views and user permissions controls in SQL Server 2005.

## 2.3 Consolidation strategies

The two general approaches to server consolidation are described in this section, together with guidance on determining what mix of strategies to deploy for your consolidation project and how and when to apply them to a SQL Server consolidation. There are two dimensions to consider when consolidating SQL Servers:

- The number of physical servers after consolidation
- The number of SQL Server instances after consolidation

### 2.3.1 Vertical consolidation: Scale up

Vertical consolidation describes the process of building a single, very powerful logical server that has the capacity to assume the databases and processing from a number of smaller, disparate systems.

Traditionally, once a 4-way server had been utilized to its maximum, it was very expensive to add more processing capacity to the same logical server, so more Windows server instances were created, increasing the Windows server management overhead. With the release of Windows Server 2003 supporting NUMA technologies, it became cost-effective and low-risk to transform a 4-way server into an 8, 12, 16, or 32-way server. The IBM System x3950 M2 has been specifically designed for this purpose.

Starting with a single x3950 M2 system (a node), this server platform can grow as your business requirements grow, with the addition of up to four interconnected nodes, to provide mainframe-class levels of processing power and availability.

Building a single logical server limits the administration and maintenance to one Windows server instance, while providing the processing power of many.

## 2.3.2 Horizontal consolidation: Scale out

Horizontal consolidation refers to the act of commoditizing a server function and then building multiple replicas of that server template to meet overall processing demand. Some load balancing technology is required to ensure that the work is evenly spread among all servers and there are limitations on how work can be distributed, making this approach unsuitable for many types of applications. Probably the most common implementation of horizontal consolidation occurs with Web farms, where many identical servers have the same role and the user connection is stateless.

Horizontal consolidation can deliver standardization and uniformity, which simplifies system maintenance and administration. However, there is not necessarily a reduction in the number of physical servers.

With regards to SQL Server, horizontal consolidation can be equated to using multiple instances of SQL Server in place of a single instance.

## 2.3.3 SQL Server instances: Single or multiple?

Using single or multiple instances equates to employing vertical or horizontal consolidation for SQL Server instances. Hosting all databases in a single instance can be considered to be vertical consolidation, while creating multiple SQL instances on a server, or cluster of servers, is a practical application of horizontal consolidation. This section discusses the criteria to assess before a decision can be reached on how many instances of SQL Server are required. “Less is more” is the guiding principle here.

In “Windows and SQL Server, both 32-bit” on page 8, we describe the memory limitations that are inherent with 32-bit versions of Windows Server and SQL Server. One way of addressing this limitation is to install multiple instances of SQL Server so that more of the available RAM can be utilized, albeit with an increased management overhead for each additional instance.

With 64-bit Windows Server and SQL Server, the addressable memory space is flat and the memory limitations are removed, so more databases can be hosted in a single instance. Now the rationale for deciding how many SQL Server instances are required comes down to the following considerations:

- ▶ How many different sort orders and collations are required by your databases? If they all use the same sort order and collation, they can all share the same instance. Each SQL Server instance has a single default sort order and collation.
- ▶ Do you need to limit DBA access to certain databases? A user with DBA rights for a SQL Server instance has DBA access to all databases in that instance. If you must restrict DBA access to certain databases, they must be hosted in a separate SQL instance with user access configured accordingly.
- ▶ What are your maintenance and availability requirements? Each SQL instance must have SQL Server patches applied to it individually. You might choose to group all databases with a similar service window in the same SQL Server instance so that they can all be taken down at the same time without impacting the user availability.
- ▶ Do you have two (or more) SQL user IDs or database names that are identical and cannot be changed? In this case, you must accommodate the databases in separate SQL Server instances. Normally, you resolve the naming conflicts and then host in the same instance, but if this is not possible, you must use separate instances.
- ▶ Do you need to continue running SQL Server 2000 alongside SQL Server 2005? If you have some databases that cannot be migrated for whatever reason (application compatibility, vendor support, and so forth) you might have to continue running SQL

Server 2000. You can install a SQL Server 2000 instance on the same consolidated hardware as your SQL Server 2005 instances.

- ▶ Do you need to run SQL Server with different hotfix or service pack levels installed? You might have some databases that do not work with the latest service packs or you might have a vendor that has only certified their application with a certain level of hotfix or service pack. Using a separate SQL Server instance, you can maintain it at the required level of hotfix or service pack, while other installed instances are kept up to date.

Depending on the licensing model you have chosen, the number of SQL Server instances you arrive at might require you to purchase additional SQL Server licenses. However, there are several options available from Microsoft that are discussed at length at:

<http://www.microsoft.com/sql/howtobuy/multipleinstances.mspx>

## 2.4 Considering service availability

The effect of reducing the number of physical servers with a server consolidation exercise is the concentration of your business systems on the remaining server infrastructure. The consolidation inevitably leads to a consolidation of your risk; for example, a server outage that might have affected one system before consolidation might now affect ten systems.

A crucial part of any consolidation project is to identify and quantify the server infrastructure risks that exist for your business. After they have been identified, the server consolidation architecture must be designed to mitigate those risks. In this section, we cover the risk assessment process and what strategies are available for providing high availability for those systems that need it.

### 2.4.1 Assessing critical systems

Before considering what can be done to protect your service availability, you need to know which systems are worth protecting and what value they add to the business. A comprehensive review of all systems in-scope for consolidation should be performed with representatives from your organization who can identify the applications that serve its core business functions. They must provide a financial value against each critical application, so that the cost of an outage to the organization can be quantified. It is this cost that should drive the decisions about what mitigation is taken in response to the identified risk.

Each organization and application have their own unique risk factors in this equation; however, some of the common elements are:

- ▶ Value of lost orders
- ▶ Value of lost employee productivity and overtime costs for catch-up
- ▶ Cost of recreating lost data
- ▶ Penalty payments for missed service level agreements (SLAs)
- ▶ Loss of customer good will
- ▶ Damage to the brand image and reputation
- ▶ Impact of legal action against organization or executives

You should assess what currently exists in the way of server infrastructure to ensure service availability for the identified critical systems. This might well highlight existing deficiencies in the deployed infrastructure that can be addressed with the consolidation project.

The architectural infrastructure design for the consolidated server environment should reflect the significance of system downtime and surrounding processes should address the data integrity requirements. These are introduced in the next section.

## 2.4.2 Ensuring system availability

There are a number of elements involved in ensuring that a newly consolidated server environment experiences high levels of service availability and this subject could probably be a paper all on its own. We have attempted to highlight the major elements that together contribute to the availability levels of a server infrastructure and we introduce some SQL Server 2005 features that can protect SQL Server service availability.

### Server room environment

When you are consolidating physical servers into a central server room facility, it is crucial that these basics be in place before consolidation begins:

- ▶ Good electrical provision so that you can provide separate power feeds to each server
- ▶ Sufficient cooling for the current servers and those that are in place after the consolidation
- ▶ Secure physical access control to ensure that unauthorized people cannot interfere with the servers
- ▶ Adequate space around each server rack for routine access to the servers during installation or server maintenance

### Hardware fault tolerance

Consider deploying server hardware that contains redundant components that can keep the server running if there is a failure. The following areas offer good protection against individual component failures in a server chassis:

- ▶ Two or more power supplies from two separate power sources
- ▶ Two or more network cards with network card teaming software
- ▶ Hardware RAID protection with a hot spare drive for local disk sub-system
- ▶ Redundant fans in the server chassis for reliable cooling
- ▶ If using a SAN, two or more SAN interface cards (host bus adapters)
- ▶ IBM Chipkill™ memory, ECC memory, redundant bit steering, memory scrubbing, and memory mirroring for protection against memory failure
- ▶ Hot-swap or hot-add capabilities for memory, hard drives, or PCI-X interface slots so that they can be added or replaced without stopping the server

Fault tolerant servers are a key component of a consolidated server platform.

### System monitoring and alerting

Having redundant hardware components is great but the value is diminished if hardware failures are not spotted and acted upon quickly. If a server has two power supplies and one fails, the server stays on and service is not interrupted. However, if the faulty power supply is not quickly replaced, the server would be disabled if the second supply also fails.

There are many tools available for server hardware monitoring; however, all IBM servers include a license for the IBM Director systems management suite. IBM Director, by supporting system management industry standards, can also manage a wide variety of non-IBM systems and operating systems. For more information about system monitoring see 4.3.6, “Monitoring” on page 67.



## Server clustering

Server clustering with Windows Server 2003 and SQL Server 2005 provides a good solution for protecting against the total failure of a physical server.

A cluster built with Windows Server 2003 consists of between two and eight servers, called nodes, using identical hardware. A high-speed shared disk sub-system is also mandatory and, in most situations, requires a SAN.

In a Windows cluster, SQL Server 2005 can be installed on each physical node. It is possible to install SQL Server multiple times on the same physical server or node. These installations are referred to as SQL Server instances and can be configured to run on all or some of the available cluster nodes. However, a SQL Server instance can only run on one node at a time.

Figure 2-1 shows a SQL Server cluster that has four instances of SQL Server installed, all being hosted on Node A. Each SQL Server instance can be considered to be a virtual server because they have their own unique server names and IP addresses and present themselves to users just like standalone physical SQL Servers.




Figure 2-1 Basic two-node SQL Server cluster in Active-Passive configuration

When a cluster node suffers a failure, the cluster detects that there is a problem with the failed node. All SQL Server instances that were running on the failed node are restarted on the other available nodes in the cluster and service is restored. SQL Server 2005 has automatic database recovery features that do not require manual intervention in most cases.

Figure 2-2 on page 22 shows how a basic two-node cluster would look after Node A has failed. The four SQL Server instances that were running on Node A, prior to the failure, have been restarted on Node B. This occurred with no manual intervention. The disruption to service availability equals the length of time that it takes to restart the failed SQL Server instances plus the automatic database recovery time.

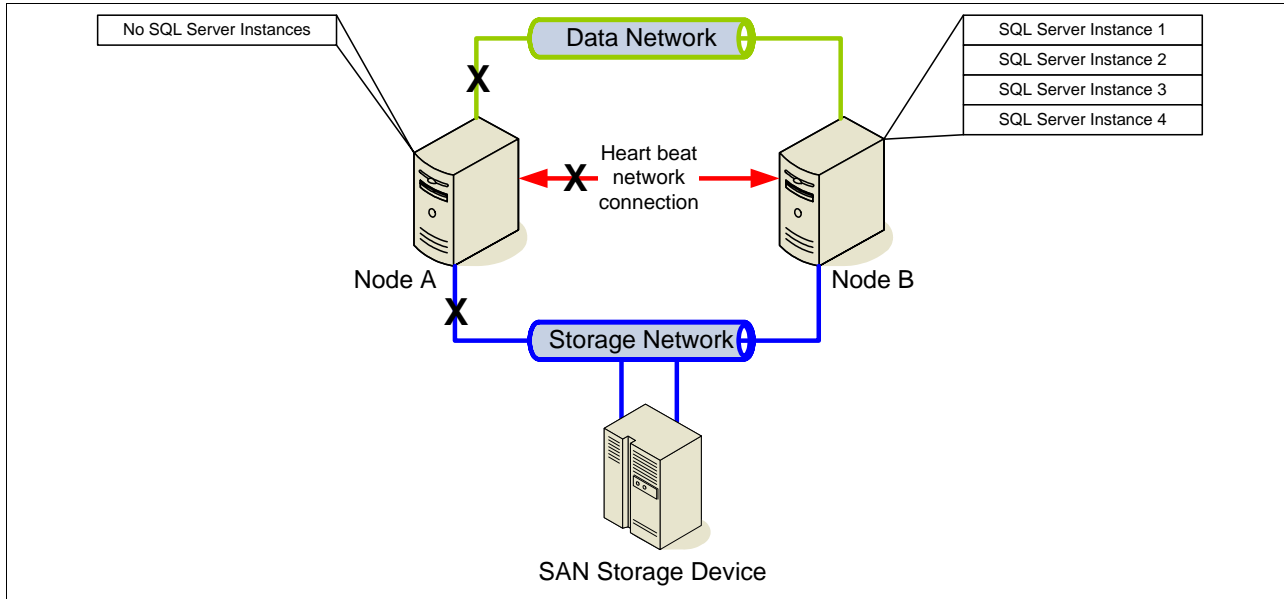


Figure 2-2 Basic two-node SQL Server cluster with failed Node A

### Database mirroring

With database mirroring in SQL Server 2005, you can take a database (the principal) and create a duplicate copy of it (the mirror), on another SQL Server 2005 server. SQL Server 2005 then ensures that the principal and mirror are kept synchronized at all times. There are different levels of synchronization available, but essentially the choice is between more performance and less protection or more protection and less performance.

In its most basic configuration, database mirroring requires two servers, one for the principal and one for the mirror. The two SQL Servers ensure that all updates to the principal are reflected in the mirror. If the principal stops, user access is disrupted until the mirror is made active, which is a manual process. Figure 2-3 illustrates a basic database mirroring setup.




Figure 2-3 Basic database mirroring configuration

Database mirroring can be set up between different server hardware, can be performed between different physical sites (depending on network bandwidth requirements and availability), and does not demand a shared disk sub-system or dedicated heart-beat network connection. As a result, it can be a much cheaper alternative to creating a clustered environment while offering similar levels of protection in certain configurations.

To benefit from automatic activation of the mirror if the principal stops, a third server is required, called the witness server. This server must have SQL Server 2005 installed on it, but it can be the Express Edition, which is a no-charge download. One witness server can be used for multiple database mirroring setups. Figure 2-4 shows how this might look.

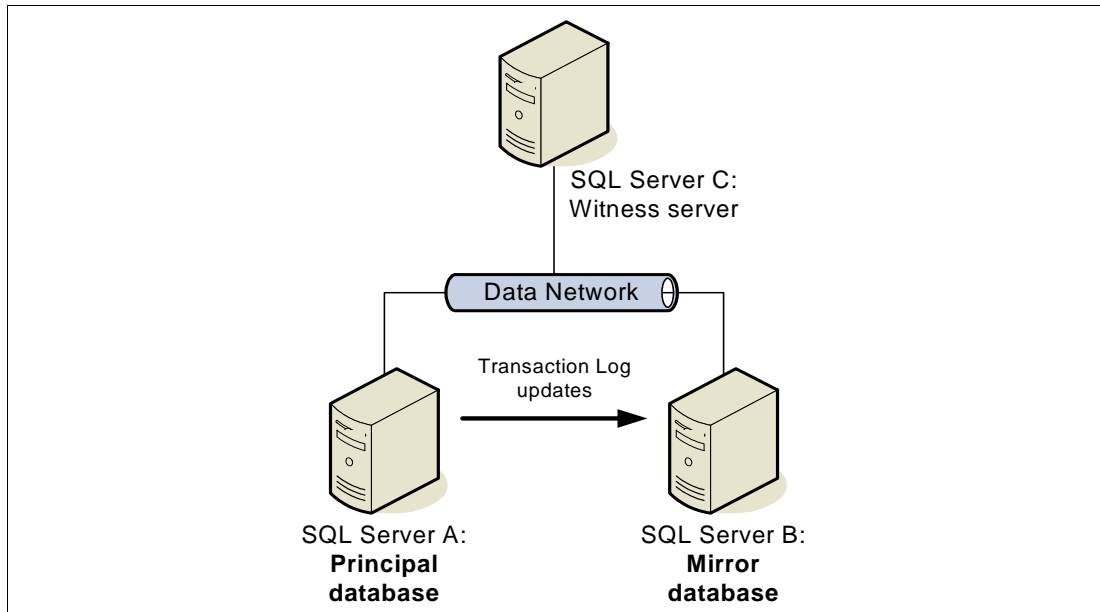


Figure 2-4 Database mirroring with a Witness server

The witness communicates regularly with the principal and mirror SQL Servers. If the witness detects that the principal is not available, it automatically makes the mirror active and accessible, much like a clustered environment would do.

Even with a witness server, if the principal stops, service availability is disrupted because any outstanding transactions must be applied to the mirror before it can be brought online. However, the SQL Server instance that hosts the database will already be running and Microsoft claims that database mirroring can provide quicker failover times than clustering. What determines failover time for mirroring is how long database recovery takes. This is a function of the amount of database update activity (actually, the number of transactions that have not yet been applied to the mirror copy of the database). Failover time does not depend on the database size.

## Data backups

In a consolidated server environment, your data storage is concentrated into a smaller number of repositories, and the data backup infrastructure must be able to process all backup tasks in the available backup window. Failure to complete backup jobs in the prescribed time-frame might infringe on application usage windows and affect performance. Therefore, the backup infrastructure must be sized and configured to reflect the greater volume of data coming through a consolidated infrastructure.

## Network design

We have covered providing two or more network cards in the consolidated server platform for fault tolerance; however, there might be a network throughput justification for using multiple network cards. Network interface cards (NICs) that are configured with teaming software can have greater effective bandwidth that prevent network I/O from being a system bottleneck. Consideration should also be given to building fault tolerance into the network infrastructure.

Implementation of a tiered network infrastructure can provide protection from the loss of network components that maintain service availability in this circumstance. Figure 2-5 illustrates a simplified tiered network design. A server is connected to the red and blue sides of the network. If either server switch fails, the remaining switches maintain server connectivity in the network. The same is true for the core switches; the result is good fault tolerance in the network.




Figure 2-5 Tiered network design with fault tolerance

## 2.5 Virtualization technologies

Server virtualization plays an extremely important role in the area of server consolidation. However, it is important to recognize that virtualization is concerned entirely with physical server consolidation and the removal of actual server boxes. It does not necessarily reduce the logical server count. There is little server management reduction after a successful virtualization project because the number of logical servers largely remains the same.

We recommend considering virtualization technologies for SQL Server in these situations:

- ▶ Test and development SQL Server environments. Virtual servers are ideal for delivering flexible and safe test and development environments for SQL Server and many other applications. You can install all SQL Server hotfixes and service packs on a virtual test SQL Server to test databases before they are installed in the production environment.
- ▶ SQL Server systems that have custom configuration requirements. Sometimes a small number of databases requires settings that are not compatible with other databases on the same physical server. The database might require a particular combination of sort order and collation, or maybe specific regional settings are required that the majority of other databases do not need. If the server resources required are low, it could be a viable solution to use a virtual server to host this type of environment.
- ▶ Replacing old server hardware when new server hardware does not support earlier operating system versions. As server hardware ages, the risk of hardware failure increases. Usually, when server hardware is refreshed, the operating system is upgraded to the current version. In some cases, that is not possible because of application compatibility issues or constraints. However, new server hardware often has limited support for previous versions of Windows Server. With a virtual server hosted on new server hardware, the current operating system can be maintained and the old server hardware could be removed after successful server virtualization.

We suggest that you also apply the following guidelines to your virtualization efforts:

- ▶ Avoid virtualizing applications that need high levels of CPU or I/O resources.
- ▶ Keep consolidated SQL Server environments and virtualized SQL Server environments on separate hardware. In other words, with hosted virtualization products like VMware Server and Microsoft Virtual Server, do not run SQL Server in the host operating system.

A server virtualization project must be approached with the same rigor as a SQL Server consolidation project and the quality of the information gathering has a big influence on the outcome of the virtualization effort.





## IBM System x3950 M2 for consolidation

In this chapter, we identify the main features of IBM System x3950 M2 combined with Microsoft SQL Server 2005 that help make database server consolidation more viable. We outline the main scalability features introduced by the IBM eX4 Architecture, the fourth generation of mainframe-inspired IBM Enterprise X-Architecture®.

The topics in this chapter are:

- ▶ 3.1, “Introduction to IBM System x3950 M2” on page 28
- ▶ 3.2, “Scalability with Windows and SQL Server” on page 38
- ▶ 3.3, “IBM sizing guide tool” on page 39

## 3.1 Introduction to IBM System x3950 M2

The x3950 M2 (Figure 3-1) is a high-performance, 2-way to 16-way scalable server that is powered by Intel Xeon® MP dual-core or quad-core processors with EM64T 64-bit extensions and fourth-generation X-Architecture. With a 64-bit framework for high-performance scalable computing, the x3950 M2 is built on the new IBM eX4 technology and is optimized for ERP applications, high-end databases, and server consolidation.



Figure 3-1 The System x3950 M2

The key features of the x3950 M2 are:

- ▶ Four-way<sup>1</sup> capable server
- ▶ Scalable to two, three, or four NUMA nodes for a server with up to 16 quad-core processors, 512 GB of RAM, and 28 PCI-X slots (see 3.1.5, “NUMA” on page 35)
- ▶ XA-64e fourth-generation chipset (see the next section for more information)
- ▶ Two standard Intel Xeon MP dual-core or quad-core processors, upgradable to four-way that support 64-bit addressing with the Intel 64 Technology architecture (see 3.1.3, “Processors” on page 31)
- ▶ Support for Intel Virtualization Technology (VT)
- ▶ Support for an embedded hypervisor option
- ▶ Support for an internal removable flash drive installed in a dedicated USB connector on the system board
- ▶ 4 GB or 8 GB memory standard expandable to 256 GB (using 8GB DIMMs), using high-performance PC2-5300 ECC DDR2 DIMMs
- ▶ Active Memory™ with Memory ProteXion, memory mirroring, memory hot-swap and hot-add, and ChipKill
- ▶ Seven half-length 64-bit PCI Express x8 slots, two of which are hot-swap
- ▶ Integrated LSI 1078 Serial-Attached SCSI (SAS) controller that supports RAID-0 and RAID-1 standard and an optional ServeRAID-MR10k RAID controller for enabling additional RAID features and a 256 MB battery backed cache
- ▶ Four internal hot-swap drive bays for up to 584 GB of internal storage (with 146 GB disks)
- ▶ Integrated Dual-port Broadcom 5709C PCI Express Gigabit Ethernet controller
- ▶ Onboard Baseboard Management Controller and Remote Supervisor Adapter II adapter
- ▶ Three-year warranty on-site, 9 hours a day, 5 days a week, next business-day response

<sup>1</sup> Four-way means 4 processor sockets. In this document, we use *way* to indicate a processor socket regardless of whether it is a dual-core processor or a quad-core processor.



### 3.1.1 Comparing the x3950 M2 with the x3950

Table 3-1 shows the major differences between the x3950 and the x3950 M2.

*Table 3-1 Major differences between x3950 and x3950 M2*

Feature	x3950 server	x3950 M2 server
Enterprise X-Architecture	Third-generation XA-64e chipset	Fourth generation XA-64e chipset
Dual-core Intel XEON E7210 and quad-core Intel XEON E7300 series processors	Not supported	Supported
Maximum SMP	32-way (using eight chassis) with 64 cores using dual cores processors	16-way (using four chassis) with 64 cores using quad-core processors
Memory	DDR2 PC2-3200 SDRAM two-way interleaved 128 GB max per 4- way chassis (16 DIMM sockets)	DDR2 PC2-5300 SDRAM two-way interleaved 256 GB max per 4-way chassis (32 DIMM sockets)
Internal disks	Six hot-swap bays	Four hot-swap bays
Disk subsystem	Adaptec AIC9410, no external port	LSI 1078, external SAS 4x port
RAID support	Standard not supported; only through optional ServeRAID-8i	Standard RAID-0 and RAID-1, additional RAID features available with optional ServeRAID-MR10k
PCI slots	Six 64-bit PCI-X 2.0 full-length, all 266 MHz	Seven PCI Express 2.5 Gbps x8, half-length
Active PCI slots	Six	Two
Ethernet controller	Broadcom 5704 dual Gigabit Ethernet	Broadcom 5709C dual Gigabit Ethernet
Video controller	ATI™ Radeon™ 16 MB	ATI RN50 16 MB memory on RSA II
Keyboard and mouse connectors	PS/2	USB
Service processor	RSA II SlimLine standard	RSA II standard
Power supply	Two 1300W supplies	Two 1440W supplies
Virtualization	Not supported	Embedded hypervisor support (1Q08)
Mechanical	3U height	4U height
Expansion	x3950 E for additional CPUs, memory, disks, PCI-X slots	No special expansion module

### 3.1.2 IBM XA-64e fourth-generation chipset

The x3950 M2 uses the fourth generation of the IBM XA-64e chipset. The architecture consists of the following components:

- ▶ One to four Xeon MP dual-core or quad-core processors
- ▶ Hurricane 4 Memory and I/O Controller (MIOC)
- ▶ Eight high-speed memory buffers
- ▶ Two PCI Express bridges
- ▶ One South bridge

Figure 3-2 shows a block diagram of the x3950 M2.

**Note:** The unit GB refers to gigabytes while the unit Gb refers to gigabits.

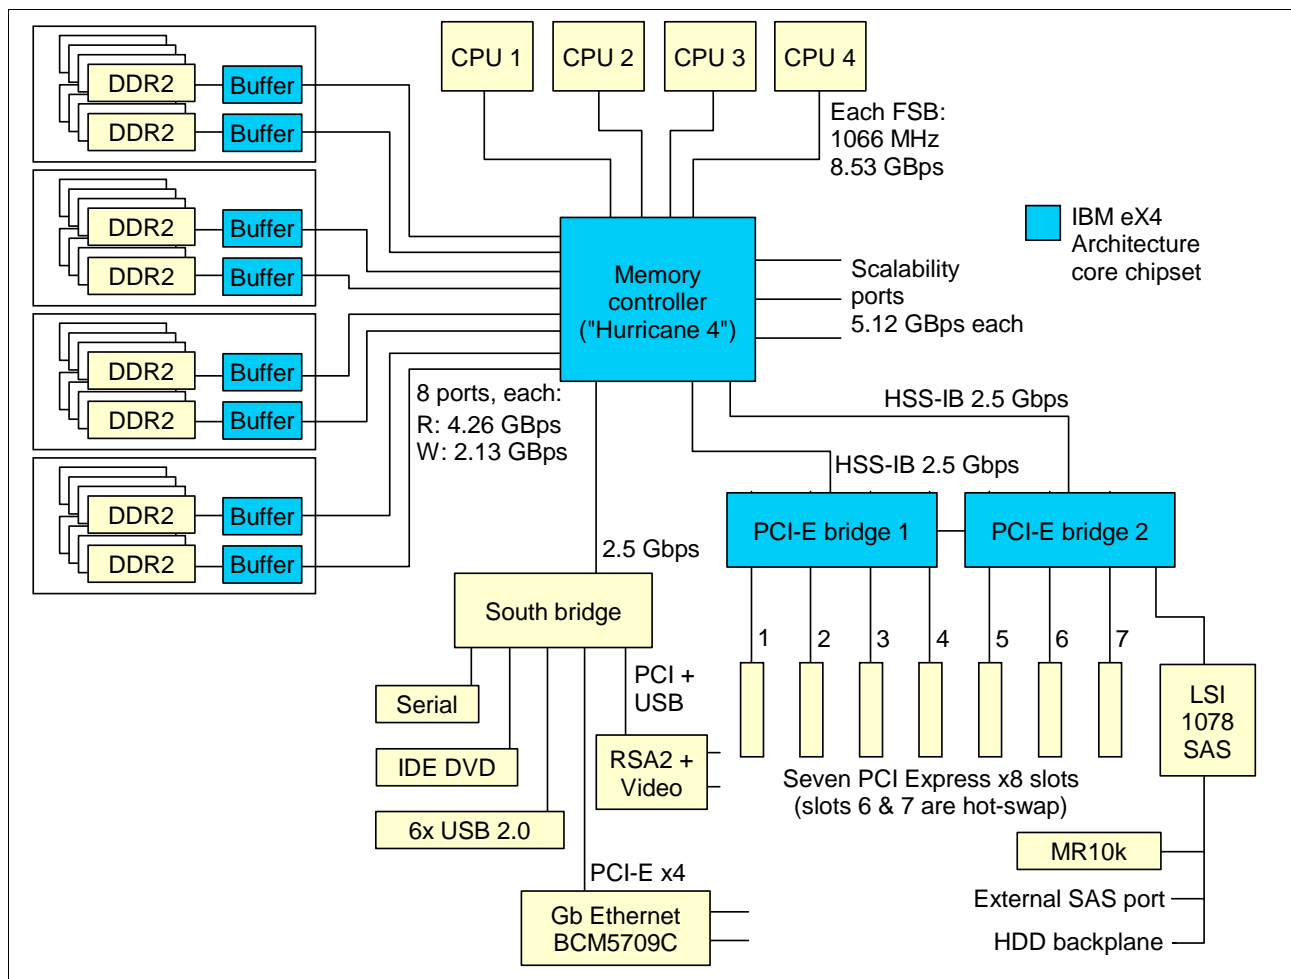


Figure 3-2 x3950 M2 system block diagram

Each memory port out of the memory controller has a peak read throughput of 4.26 GBps and a peak write throughput of 2.13 GBps. DIMMs are installed in matched pairs, two-way interleaving, to ensure that the memory port is fully utilized.

Because there are eight memory ports, spreading installed DIMMs over all four memory ports can improve performance. The eight independent memory ports provide simultaneous access

to memory. With four memory cards installed, and eight DIMMs in each card, peak read memory bandwidth is 34.1 GBps and peak write bandwidth is 17.1 GBps.

The memory controller routes all traffic from the eight memory ports, four microprocessor ports, and the three PCI bridge ports. The memory controller also has embedded DRAM which, in the x3950 M2, holds a snoop filter lookup table. This filter ensures that snoop requests for cache lines go to the appropriate microprocessor bus and not all four of them, which improves performance.

Figure 3-2 on page 30 also shows that PCI bridge 1 supplies four of the seven PCI Express x8 slots on four independent PCI Express buses. PCI bridge 2 supplies the other three PCI Express x8 slots plus the onboard SAS devices, including the optional ServeRAID-MR10k. A separate South bridge supplies all the other onboard PCI devices like the USB ports, onboard Ethernet, and the standard RSA II.

### 3.1.3 Processors

All models of the x3950 M2 have two processors, either the Intel Xeon Processor E7210 dual-core or the Intel Xeon Processor E7300 series quad-core Tigerton, and they support two or four processors. The processors used by the x3950 M2 must be identical in model, speed, and cache size. Three processors in one system is not supported. You can connect multiple models of x3950 M2 to form larger configurations (see “N-way configurations” on page 35).

The x3950 M2 processors include the key features described in the following sections.

#### Models with dual-core or quad-core processors

The Tigerton dual-core processors are similar to a two-way SMP system except that the two processors, or *cores*, are integrated into one silicon die. This realizes the benefits of two-way SMP with less power consumption and faster data throughput between the two cores. The resulting core frequency is lower, but the additional processing capacity means an overall gain in performance. The Tigerton quad-core processors add two more cores onto the same die. Hyper-Threading Technology is not supported.

Each core has separate L1 instruction and data caches, execution units (integer, floating point, and so on), registers, issue ports, and pipelines for each core. A multi-core processor achieves more parallelism than Hyper-Threading Technology, because these resources are not shared between the two cores.

With double and quadruple the number of cores for the same number of sockets, it is critical that the memory subsystem meet the demand for data throughput. The 34.1 GBps peak throughput of the eX4 Architecture with four memory cards is well-suited to dual-core and quad-core processors.

#### 1066 MHz front-side bus

The Tigerton Xeon MP uses two 266 MHz clocks, out of phase with each other by 90°, and uses both edges of each clock to transmit data (Figure 3-3).




Figure 3-3 Quad-pumped front-side bus

A quad-pumped 266 MHz bus results in a 1066 MHz front-side bus. The bus is 8 bytes wide, which means it has an effective burst throughput of 8.53 GBps. This can have a substantial impact, especially on TCP/IP-based LAN traffic.

### **Intel 64 Technology (formerly known as EM64T)**

First introduced in the Xeon DP *Nocona* processor, Intel 64 Technology is a 64-bit extension to the industry standard IA32 32-bit architecture. Intel 64 Technology adds:

- ▶ A set of new 64-bit general purpose registers (GPRs)
- ▶ 64-bit instruction pointers
- ▶ The ability to process data in 64-bit chunks

Even though the names of these extensions suggest that the improvements are simply in memory addressability, Intel 64 Technology is, in fact, a fully functional 64-bit processor.

The Tigerton processors limit memory addressability to 40 bits of addressing.

There are three distinct operation modes available in Intel 64 Technology:

- ▶ 32-bit legacy mode

In this mode, processors with Intel 64 Technology act just like any other IA32-compatible processor. You can install your 32-bit operating system on such a system and run 32-bit applications, but you cannot make use of the new features, such as the flat memory addressing above 4 GB or the additional GPRs. Thirty-two-bit applications run just as fast as they would on any current 32-bit processor.

Most of the time, IA32 applications run even faster because there are numerous other improvements that boost performance regardless of the maximum address size.

- ▶ Compatibility mode

This is an intermediate mode of the full 64-bit mode. To run in compatibility mode, you must install a 64-bit operating system and 64-bit drivers to support both 32-bit applications and 64-bit applications.

With the compatibility mode, you can run a 64-bit operating system and still run unmodified 32-bit applications. Each 32-bit application is limited to a maximum of 4 GB of physical memory; however, the limit is imposed per-process and not system wide. This means that every 32-bit process has its own 4 GB of physical memory space, if sufficient physical memory is installed. This is an improvement over IA32, where the operating system kernel and the application had to share 4 GB of physical memory.

Additionally, the compatibility mode does not support the virtual 8086 mode, so real-mode applications are not supported. 16-bit protected mode applications are, however.

- ▶ Full 64-bit mode

Intel refers to this mode as the *IA-32e mode*. For AMD, it is the *long mode*. This mode is applied when a 64-bit OS and 64-bit application are used. In the full 64-bit operating mode, an application can have a virtual address space of up to 40 bits, equating to one terabyte (TB) of addressable memory. The amount of physical memory is determined by how many DIMM slots the server has and the maximum DIMM capacity supported and available at the time.

Applications that run in full 64-bit mode have access to the full physical memory range, depending on the OS, to the new GPRs, and to the expanded GPRs. However, it is important to understand that this mode of operation requires not only a 64-bit SO (and, of course, 64-bit drivers) but also a 64-bit application that has been recompiled to take full advantage of the various enhancements of the 64-bit addressing architecture.

For more information about the features of the Xeon quad-core processor, visit:

<http://www.intel.com/products/server/processors/index.htm?iid=process+server>

For more information about Intel 64, see:

<http://www.intel.com/technology/architecture-silicon/intel64>

### 3.1.4 System memory

Memory is implemented in the x3950 M2 by memory cards. The server supports up to four memory cards. Each card has eight DIMM sockets, providing a total of up to 32. Using 8GB DIMMs in every socket, the server can hold 256 GB of RAM. With four nodes connected together, this means a single system image has access to 1 TB of memory.

**Note:** For performance reasons, all nodes should have the same amount of memory.

Also, in a multinode scalable system, the XceL4v Dynamic Server Cache dynamically allocates 256 MB of main memory in each node for use as L4 cache. The result is a reduction in overall memory that is available to the operating system of 256 MB per node. Therefore, in a 16-way configuration there is a 1 GB reduction of main system memory.

In a multinode configuration, the memory in all nodes is combined to form a single coherent physical address space. Therefore, for any given region of physical memory, some processors are *closer* to it than other processors. Conversely, for any processor, some memory is considered *local* and other memory is *remote*. The partition descriptor table of the system is used to ensure optimal memory use.

The memory is two-way interleaved, meaning that memory DIMMs are installed in pairs. There are eight ports from the Hurricane 4 memory controller to memory, with each supporting up to 4.26 GBps read data transfers and 2.13 GBps write data transfers (see Figure 3-2 on page 30).

The DIMMs operate at 533 MHz, so that they are in sync with front-side bus. However, the DIMMs are 677 MHz PC2-5300 spec parts because they have better timing parameters than the 533 MHz equivalent. The memory throughput is 4.26 GBps, or 533 MHz x 8 bytes per memory port for a total of 34.1 GBps with four memory cards.

There are a number of advanced features in the x3950 M2 memory subsystem, collectively known as *Active Memory*:

- ▶ Memory ProteXion

The Memory ProteXion feature (also known as *redundant bit steering*) provides the equivalent of a hot-spare drive in a RAID array. It is based in the memory controller, and it enables the server to sense when a chip on a DIMM has failed and to route the data around the failed chip.

Normally, 128 bits out of every 144 are used for data and the remaining 16 bits are used for ECC functions. However, the x3950 M2 needs only 12 bits to perform the same ECC functions, leaving 4 bits free. If a chip failure on the DIMM is detected by memory scrubbing, the memory controller can reroute data around that failed chip through these spare bits.

It can do this automatically without issuing a Predictive Failure Analysis® (PFA) or light path diagnostics administrator alert, although an event is logged to the service processor log. After the second DIMM failure, PFA and light path diagnostics alerts occur on that DIMM as normal.

- ▶ **Memory scrubbing**

Memory scrubbing is an automatic daily test of all the system memory that detects and reports memory errors that might be developing before they cause a server outage.

Memory scrubbing and Memory ProteXion work together and do not require memory mirroring to be enabled to work properly.

When a bit error is detected, memory scrubbing determines if it is recoverable or not. If it is, Memory ProteXion is enabled and the data that was stored in the damaged locations is rewritten to a new location. The error is reported for the purpose of preventative maintenance. As long as there are enough good locations for the proper operation of the server, no further action is taken other than recording the error in the error logs.

If the error is not recoverable, then memory scrubbing sends an error message to the light path diagnostics, which then turns on the proper lights and LEDs to guide you to the damaged DIMM. If memory mirroring is enabled, then the mirrored copy of the data from the damaged DIMM is used until the system is stopped and the DIMM replaced.

Because the x3950 M2 is now capable of supporting a large amount of memory, IBM has added the **Initialization Scrub Control** setting to the BIOS, so that customers can choose when this scrubbing is done and therefore potentially speed up the start process.

- ▶ **Memory mirroring**

Memory mirroring is roughly equivalent to RAID-1 in disk arrays, in that usable memory is halved and a second copy of data is written to the other half. If 8 GB is installed, then the operating system sees 4 GB once memory mirroring is enabled. It is disabled in the BIOS by default. Because all mirroring activities are handled by the hardware, memory mirroring is operating-system independent.

When memory mirroring is enabled, there are certain restrictions on placement and size of memory DIMMs and the placement and removal of memory cards.

- ▶ **Chipkill memory**

Chipkill is integrated into the XA-64e chipset, so it does not require special Chipkill DIMMs and is not visible to the operating system.

When Chipkill is combined with Memory ProteXion and Active Memory, the x3950 M2 provides high reliability in the memory subsystem. When a memory chip failure occurs, Memory ProteXion handles the rerouting of data around the failed component as described previously. However, if a further failure occurs, the Chipkill component in the memory controller reroutes data.

The memory controller provides memory protection similar to disk array striping with parity, writing the memory bits over multiple memory chips on the DIMM. It can reconstruct the missing bit from the failed chip and continue working as usual. One of these additional failures can be handled for each memory port for a total of eight Chipkill recoveries.

- ▶ **Hot-add and hot-swap memory**

The x3950 M2 supports the replacement of failed DIMMs while the server is still running. This hot-swap support works with memory mirroring. The server also supports adding additional memory while the server is running. Adding memory requires OS support. These two features are mutually exclusive. Hot-add requires that memory mirroring be disabled and hot-swap requires that memory mirroring be enabled.

In addition, to maintain the highest levels of system availability, if a memory error is detected during POST or memory configuration, the server can automatically disable the failing memory bank and continue operating with reduced memory capacity. You can manually re-enable the memory bank after the problem is corrected with the Setup menu in the BIOS.

Memory mirroring, Chipkill, and Memory ProteXion provide multiple levels of redundancy to the memory subsystem.

Combining Chipkill with Memory ProteXion allows up to two memory chip failures for each memory port on the x3950 M2, for a total of eight failures sustained. For example, the first failure detected by the Chipkill algorithm on each port does not generate a light path diagnostics error because Memory ProteXion recovers from the problem automatically. Each memory port can then sustain a second chip failure without shutting down. As long as memory mirroring is enabled, the third chip failure on that port sends the alert and takes the DIMM offline, but keeps the system running out of the redundant memory bank.

### 3.1.5 NUMA

The x3950 M2 is the initial base building block, or *node*, for a scalable system. At their most basic, these nodes are comprised of a 4-way SMP-capable system with processors, memory, and I/O devices.

You can form a multinode configuration by adding one or more x3950 M2 servers to another. The following configurations are possible:

- ▶ A two-node complex that consists of two x3950 M2 servers, with four or eight processors and up to 512 GB RAM
- ▶ A three-node complex that consists of three x3950 M2 servers, with six or 12 processors and up to 768 GB RAM
- ▶ A four-node complex that consists of four x3950 M2 servers, with eight or 16 processors and up to 1 TB RAM

#### N-way configurations

Unlike the x3950 and xSeries® 460, there is no special modular expansion enclosure for the x3950 M2. The multinode configuration is simply formed by another x3950 M2 or an x3850 M2 with the ScaleXpander Option Kit.

The ScaleXpander Option Kit includes the cable necessary to join two nodes together plus a small component (Figure 3-4) that enables the formation of a multinode complex.




Figure 3-4 ScaleXpander Option Kit component to enable multi-node configurations

**Note:** In this paper, when we refer to an x3950 M2, we mean either an x3950 M2 or an x3850 M2 with the ScaleXpander Option Kit.

## Multinode scalable systems

A scalable system consists of one, two, three, or four x3950 M2 systems (Figure 3-5). Each node can have either two or four processors.

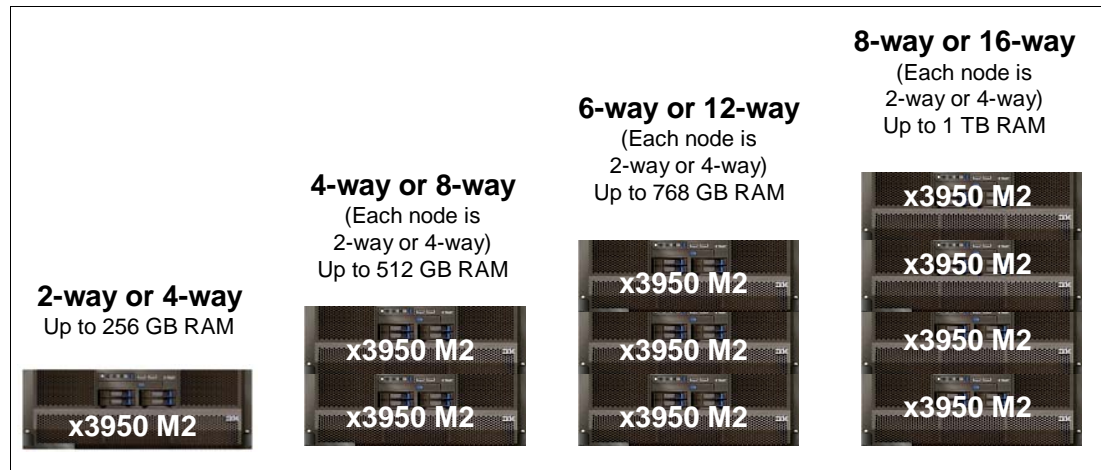


Figure 3-5 The multinode configurations supported

The following configuration rules apply:

- ▶ Combinations of x3950 M2s

You can connect one, two, three, or four x3950 M2 servers. Specific cabling is required.

- ▶ Processors

Each node must have two or four processors. In a multinode configuration, every node must have either two processors or four processors. (In other words, you cannot have some 2-way nodes and some 4-way nodes in the same complex). All processors must be the same speed and cache size.

- ▶ Memory

For performance reasons, you should have the same amount of memory in each node. A minimum of 2 GB of RAM is required in each node.

- ▶ Firmware

All system firmware, including the system BIOS, diagnostics, BMC firmware, and RSA II firmware, must be at the same level for all systems.

Updating the system BIOS in every node in a scalable system can be performed from the primary node. The server diagnostics and the BMC and RSA II firmware must be individually updated on each node, but this can be performed remotely: The RSA II firmware can be updated using the RSA II Web interface or IBM Director. The BMC firmware can be updated with an RSA II remote console session using the remote diskette function.

Disk drives in any of the x3950 M2s are seen by the OS as normal disk drives, and all PCI-Express slots and onboard Gigabit Ethernet ports in the x3950 M2 are visible.

A fully configured, four-node, scalable system with quad-core processors has 64 cores, 1TB of memory (using 8GB DIMMs), 28 PCI-Express adapters, 2.3 TB of raw disk space, and eight Gigabit Ethernet connections.



### 3.1.6 Broadcom dual Gigabit Ethernet controller

The x3950 M2 offers a dual Gigabit Ethernet controller integrated standard. The x3950 M2 includes one dual-port Broadcom BCM5709C 10/100/1000 BASE-T MAC (Media Access Controller) on the PCI Express x4 bus. The BCM5709C has the following features:

- ▶ Full and half-duplex performance at all speeds (10/100/1000 Mbps, auto-negotiated)
- ▶ Two IEEE 802.3 Ethernet MAC addresses
- ▶ Integrated on-chip memory for buffering data transmissions
- ▶ Dual onboard DMA engines to maximize bus throughput and minimize CPU overhead
- ▶ IPMI support for system management

The Broadcom controller also includes software support for failover, layer-3 load balancing, and comprehensive diagnostics.

**Note:** Category 5 or better Ethernet cabling is required with RJ-45 connectors. If you plan to implement a Gigabit Ethernet connection, ensure that your network infrastructure is capable of the necessary throughput to match the I/O capacity of the server.

### 3.1.7 SAS disk subsystem

The x3950 M2 has a disk subsystem that is comprised of an LSI Logic 1078 Serial Attached SCSI (SAS) controller and four internal 2.5-in SAS hot-swap drive bays. The x3950 M2 supports internal RAID-0 and RAID-1. The optional ServeRAID-MR10k, part number 43W4280, provides additional RAID levels and a 256 MB battery backed cache.

SAS is the logical evolution of SCSI. SAS uses much smaller interconnections and offers longer cabling distances, smaller form factors, and greater addressability. The x3950 M2 has an external SAS x4 port that is used with the optional ServeRAID-MR10k. This port supports SAS non-RAID disk enclosures such as the EXP3000 and has an SFF-8088 connector.

### 3.1.8 PCI subsystem

There are five half-length full-height PCI Express x8 slots and two half-length full-height active PCI Express x8 internal to the x3950 M2 (see Figure 3-2 on page 30), and all are vacant in the standard models. All seven slots have the following characteristics:

- ▶ Separate bus from the other slots and devices
- ▶ PCI Express x8
- ▶ 40 Gbps full duplex
- ▶ 64-bit, each supporting 32-bit adapters as well

Slots 6 and 7 also support Active PCI hot-swap adapters. The optional ServeRAID-MR10k adapter does not use a PCI slot because it has a dedicated slot on the motherboard.

The PCI subsystem also supplies these I/O devices:

- ▶ LSI 1078 Serial-attached SCSI (SAS) controller
- ▶ Broadcom dual port 5709C 10/100/1000 Ethernet
- ▶ Six USB ports, two on the front panel, three on the rear, one onboard
- ▶ Remote Supervisor Adapter II adapter in a dedicated socket on the I/O board
- ▶ ATI ES1000™ 16MB video controller
- ▶ EIDE interface for the DVD/ROM drive
- ▶ Serial port
- ▶ Trusted Platform Module (TPM)

### 3.1.9 Redundancy

The x3950 M2 has the following redundancy features to maintain high availability:

- ▶ There are six hot-swap, multi-speed fans. These fans provide cooling redundancy and enable individual fan replacement without stopping the server. Each of the three groups of two fans is redundant. If a fan fails, the other fans speed up to continue to provide adequate cooling until the fan can be hot-swapped by the IT administrator. In general, failed fans should be replaced in 48 hours.
- ▶ The two Gigabit Ethernet ports can be configured as a team to form a redundant pair.
- ▶ The memory subsystem has a number of redundancy features, including memory mirroring and Memory ProteXion (see 3.1.4, “System memory” on page 33).
- ▶ Support is available for RAID disk arrays, both with the onboard LSI 1078 for RAID-0 and RAID-1. The optional ServeRAID-MR10k provides additional RAID features and a 256 MB battery backed cache. The x3950 M2 has four internal, hot-swap disk drive bays.
- ▶ The two, standard 1440 W hot-swap power supplies are redundant in all configurations at 220 V. At 110 V, the second power supply is not redundant.

## 3.2 Scalability with Windows and SQL Server

Use Table 3-2 as a reference if you are trying to choose between SQL Server 2005 and Windows Server 2003 versions for your consolidated environment.

Table 3-2 Scalability for x3950 M2, Windows Server 2003, and SQL Server 2005.

	IBM System x3950 M2	Windows Server 2003 Standard Edition (x64)	Windows Server 2003 Enterprise Edition(x64)	Windows Server 2003 Datacenter Edition(x64)	SQL Server 2005 Standard Edition (x64)	SQL Server 2005 Enterprise Edition (x64)
<b>Processors</b>	2-16	1-4	1-4	1-32	1-4	1-64
<b>Memory</b>	1TB	32GB	1TB	1TB	OS and hardware Maximum	OS and hardware Maximum
<b>Failover clustering</b>	Provided a OS level	N/A	1-8 nodes	1-8 nodes	1-2 nodes	1-8 nodes
<b>Windows System Resource Manager (WSRM)</b>	Provided by OS	No	Yes	Yes	Soft NUMA (See 4.3.9, “NUMA Considerations” on page 68)	Soft NUMA (See 4.3.9, “NUMA Considerations” on page 68)
<b>Hot-add memory</b>	Yes	No	Yes	Yes	No	Yes

We recommend combining x3950 M2, Windows Server 2003 (Enterprise or Datacenter edition), and SQL Server 2005 Enterprise Edition for a scalable consolidated environment.

### 3.3 IBM sizing guide tool

IBM offers an online workload-based sizing tool for estimating System x servers for SQL Server 2000 and SQL Server 2005. To access the guide, go to:

<http://www.developer.ibm.com/servers/sizing/portal/search.jsp>

Enter Microsoft SQL Server WLE-based Sizing Guide in the **4. Refine search by limiting results to solution names containing** field. Click **Run** to launch it. An alternate direct link to the guide is:

[http://www.developer.ibm.com/servers/sizing/sizingguide/guide\\_redirect.jsp?guide\\_id=sgq80502835091106002](http://www.developer.ibm.com/servers/sizing/sizingguide/guide_redirect.jsp?guide_id=sgq80502835091106002)

Before you size your system, you need this information:

- ▶ Version of SQL Server that you are using
- ▶ Total of your database sizes and percent growth in the database size per year
- ▶ Desired server form factor for this sizing, such as rack-optimized, tower, and blade
- ▶ RAID support for this workload, such as RAID-1, RAID-5, and so forth
- ▶ Number of active user connections and peak transactions per hour per connection
- ▶ Percent growth per year in the number of users
- ▶ The percent of total DB transactions per hour that update the database

To show you how to use the sizing tool, we will go through the process, with these values:

- ▶ We are using SQL Server 2005 version.
- ▶ 500GB is total size for the data that will populate the consolidated server.
- ▶ 5% is the estimated annual growth for the data.
- ▶ We want a Rack optimized server.
- ▶ We will use RAID-5 for fault tolerance.

1. Launch the sizing guide. Figure 3-6 shows the values to enter.

Version: 2007.5.1 04-Dec-07 www-912 **IBM Systems Workload Estimator**

Workload Selection | **Workload Definition** | Selected System | Help/Tutorials

- SQL
- Options

• Reset This Workload • Save This Workload • Edit Workload Name  
• Refresh This Workload • Modify Intervals

**SQL** Partition: [SQL](#)  
SQL Workload Definition Windows - Server 2003  
Native Mode

**Global Selections**

Please provide global parameters using the selections below for this database server sizing.  
Note: You can return to this page and modify these parameters as you wish.

1. Which [version](#) of MS SQL will you be using?
2. What is the [size](#) of your database in Gigabytes?  GBs
3. What percent growth in the [database](#) size do you expect?  Percent
4. Select the [desired](#) server form factor for this sizing.
5. [RAID Support](#) for this workload:

Figure 3-6 SQL Workload Definition

2. Click **Continue**.
3. Enter the OLTP input (Figure 3-7), considering the following information:
  - 1000 active user connections
  - Up to 10,000 transactions per hour during peak usage
  - 5% is the estimated annual user increase
  - The server will handle 66% write queries and 34% read queries

The screenshot shows the 'IBM Systems Workload Estimator' interface. At the top, it displays 'Version: 2007.5.1 04-Dec-07 www-912' and the title 'IBM Systems Workload Estimator'. Below this are navigation tabs: 'Workload Selection', 'Workload Definition', 'Selected System', and 'Help/Tutorials'. The 'Workload Definition' tab is active, showing a list of workloads including 'SQL'. Under 'SQL', there are options: 'Options', 'Reset This Workload', 'Save This Workload', 'Edit Workload Name', 'Refresh This Workload', and 'Modify Intervals'. Below this, the current workload is identified as 'SQL' with a 'Partition: SQL' and 'Windows - Server 2003 Native Mode'. The main section is titled 'OLTP input' and contains a descriptive paragraph: 'Database systems are highly scalable and highly customizable, results therefore vary widely based how they are used. This sizing is based on OLTP (On-Line Transaction Processing) applications are based on interactive transactions such as entering an order, or looking up the inventory on hand for a particular item.' Below the text are four numbered input fields:
 

1. Enter the number of active user connections. Input: 1000 connections.
2. Specify the number of peak transactions per hour per connection. Input: 10 Thousands.
3. What percent growth in the number of users do you expect? Input: 5 Percent.
4. Indicate the percent of total DB transactions that update the database. A slider bar shows 'Updates 34%' and 'Queries 66%'.

 At the bottom, there are 'Back' and 'Continue' buttons.

Figure 3-7 OLTP input

4. Click **Continue**.
5. Figure 3-8 on page 41 shows the estimated server configuration for this consolidated environment.

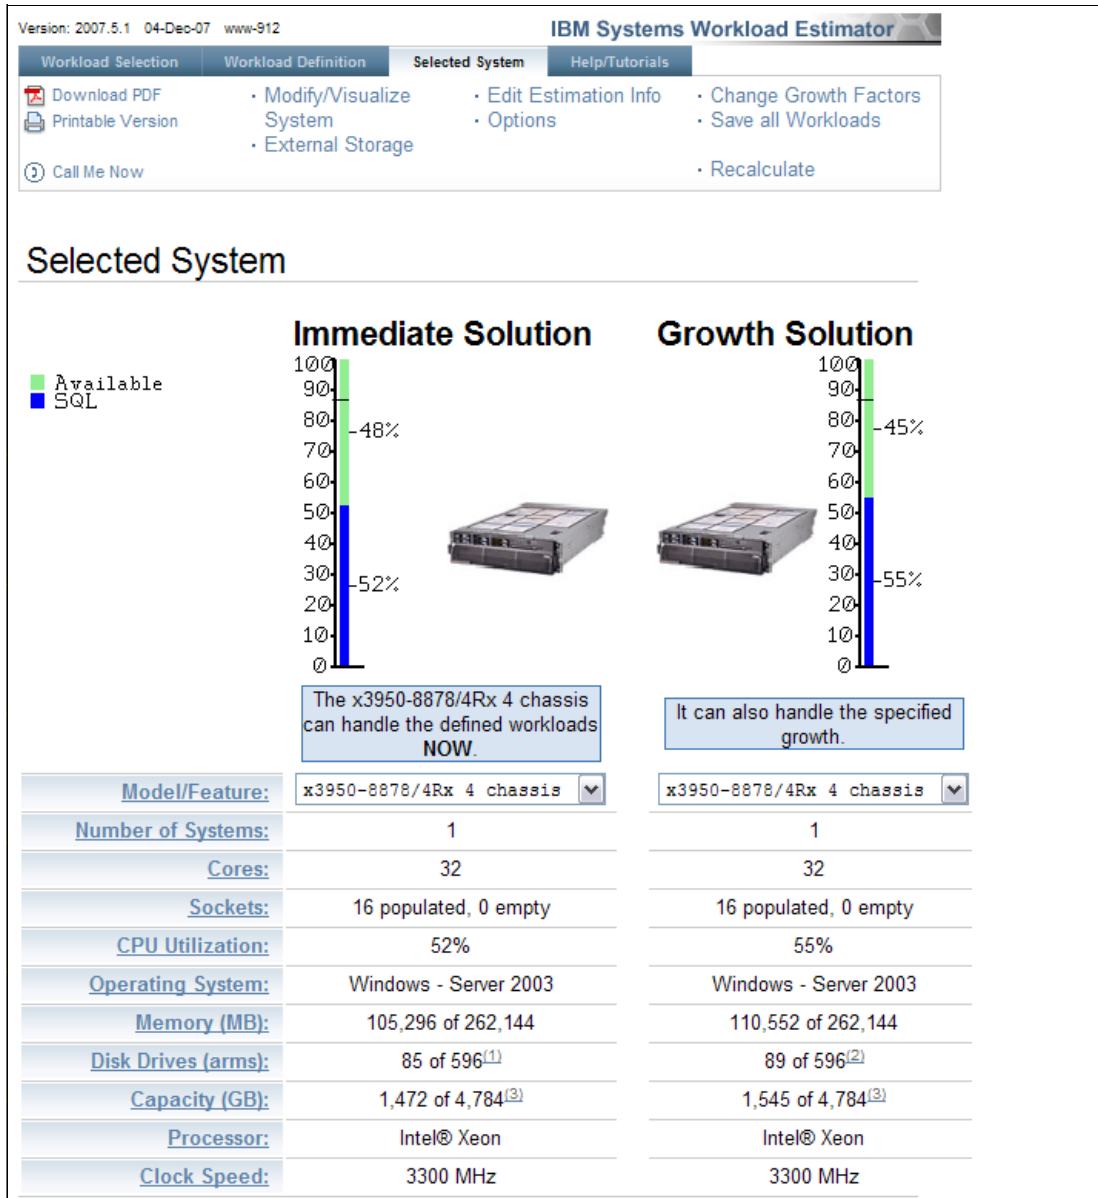


Figure 3-8 Estimated server configuration

**Note:** This example uses hypothetical data. You must have the workload information prior to estimating the server.

Also, the sizing tool considers the information for one SQL Server instance, so make sure that you consider adding more memory for additional SQL Server instances.





# Consolidation strategy and implementation

There are various consolidation concepts, but in general, you can divide them into physical and logical consolidation. For the purpose of this paper, we consider consolidation to be physical server consolidation (see 2.2, “Forms of consolidation” on page 16).

By the end of the 1990s, there was a period of extreme enthusiasm and overinvestment in IT infrastructure. Investments were made in IT just for the sake of technology and there was little or no control over the IT environment. IT infrastructure today has become a commodity and rarely represents a competitive advantage by itself.

In a mature company, IT investments must translate to lower total cost of ownership and a sustainable return on investment. In this scenario, server consolidation helps because it can centralize management, optimize resource usage, and standardize the environment. This can result in lower administration, staffing, and licensing costs.

There is no one strategy for database consolidation. Instead, consolidation must consider the particular aspects of each company. In this chapter, we discuss best practices for database consolidation.

The topics are:

- ▶ 4.1, “The five phases of consolidation” on page 44
- ▶ 4.2, “Assessment and analysis” on page 44
- ▶ 4.3, “Consolidation project” on page 56
- ▶ 4.4, “Testing” on page 71
- ▶ 4.5, “Implementation and stabilization” on page 72
- ▶ 4.6, “Best practices” on page 72
- ▶ 4.7, “Capacity management” on page 73
- ▶ 4.8, “Services offerings” on page 75

## 4.1 The five phases of consolidation

The objective of consolidation is to have the least number of servers and the least number of SQL Server instances as possible while maintaining the existing functions to decrease administration overhead. Our proposition is to implement the consolidation in five phases:

▶ **Phase 1: Assessment and analysis**

This is the most important phase in the project. Careful planning and documentation determine its success. In this phase, you determine which databases are eligible for consolidation and which are not. You determine the number of instances, SQL Server versions, and so forth that you need to meet business and technical requirements.

▶ **Phase 2: Consolidation project**

This is the development of the project plan. Using your environment information and business and technical constraints, you design the technical solution for the implementation of the consolidated environment.

▶ **Phase 3: Testing**

Creating a lab environment to test the intended configuration setting helps you identify any gaps in the project and is very important.

**6. Phase 4: Implementation and stabilization**

After you have successfully tested your plan and addressed what you found, you are ready to implement in production. Start with less critical databases so that you gain expertise in the migration process.

**7. Phase 5: Capacity management**

After the consolidation project is in place and stabilized, you must manage the hardware resource consumption to make sure that the capacity meets current and future business requirements.

## 4.2 Assessment and analysis

The initial steps toward consolidation are the assessment and documentation of your environment. You must gather and document all the existing information for the current environment, such as number of instances, SQL server versions, and so on. The level of detail determines the success of the consolidation process. In this phase, you must identify all incompatibilities to determine which servers and databases are eligible for consolidation.

After the analysis of the information, you identify business and technical constraints, such as:

- ▶ User databases that are not eligible for consolidation. Some restrictions, such as vendor or business requirements, might require a dedicated database server.
- ▶ User databases that make heavy usage of the tempdb database. Heavy usage of tempdb by one of few databases affects the performance of the rest. Consider installing a separate instance for these databases.
- ▶ Isolation of sysadmin access. Logins that are assigned the sysadmin server role have unrestricted access to the SQL Server instance. If you must isolate specific databases for a specific administrator group, you must have a separate SQL Server instance installation.
- ▶ Other constraints. You must try to take into consideration any other constraint that can generate an impact on how the solution is designed such as application incompatibility to SQL Server 2005 and so on.



## 4.2.1 Assessment guidelines

The information that you must gather during the assessment includes:

- ▶ Existing servers and databases
- ▶ Application-specific information
- ▶ Analysis Services
- ▶ External dependencies
- ▶ Network

The sections that follow provide more details about what you should collect.

**Note:** Most of this information can be obtained with the script in Example 4-2 on page 49.

### Existing servers and databases

For existing servers, you must document and take into consideration the following information:

- ▶ Name  
This is important to help you identify any naming conflict.
- ▶ OS version, language, and regional settings  
You must be very careful when consolidating servers with different OS languages. You must verify if the application relies on OS defaults, such as date and time formats.
- ▶ CPU  
You must document the CPU model, frequency, and number of CPUs used. This is important for estimating the CPUs for the new hardware in the consolidated environment.
- ▶ Memory  
Basically, you must document the amount of physical memory for future hardware estimation.
- ▶ Windows domain name  
When consolidating servers from different domains, you must ensure that the proper level of trust is in place before the testing phase.
- ▶ Geographical and physical location  
This determines how the data will be moved during the process. You might want to move the servers to the same physical location before implementation or you might decide that you want to have a backup sent over in some type of media.
- ▶ SQL server version  
This information is used during the analysis to decide on the migration strategy. You must document the exact version. For example:
  - Is it version 2000, 7.0, or 2005?
  - What is the service pack level?
  - Is it Enterprise or Standard edition?
  - Is it 32 or 64 bits?

For each SQL server instance, you need the following information:

- ▶ Default instance or named instance  
You must document whether it is a default instance installation or a named instance. With this information, you can identify possible name conflicts.

- ▶ **Default collation**

Collation refers to a set of rules that determine how data is sorted and compared. Character data is sorted using rules that define the correct character sequence, with options for specifying case-sensitivity, accent marks, kana character types, and character width. When a database is created in SQL Server, it uses the default collation unless specified differently.

For more information about collations, see SQL Server 2005 Books Online topic “about collations.” Microsoft SQL Server 2005 Books Online can be accessed from:

<http://msdn2.microsoft.com/en-US/library/ms130214.aspx>

- ▶ **Port number**

SQL Server can make use of manual or dynamic port numbering. By default, the default instance installation uses TCP port 1433 and named instances are configured with dynamic ports that are assigned upon instance start up. For the application to connect to the named instance, you either have to hard code the port number in the connection string or the SQL browser service has to be running so that it can be used for port number discovery by the application.

Make sure you document how your port numbering is set up, especially when you have named instances installed on the server. See 4.6, “Best practices” on page 72 for more details.

- ▶ **Service account**

Document the accounts that SQL services run under and what level of access those accounts have.

- ▶ **Linked servers**

Document any linked servers that might exist.

- ▶ **User objects in master database**

In some cases, you see user stored procedures or user-defined functions created in the master database. Document the existence of these objects. During the analysis phase, confirm if they really must be located in the master database. Otherwise, consider moving each off the master database.

- ▶ **Replication**

Document all the technical configurations for replicated environments, such as publisher, subscribers, replication type, replication folder location, and so on.

- ▶ **DTS/SSIS packages**

Document the name and a brief description of the packages. Deeper investigation during the analysis phase is necessary for determining their compatibility with the new environment.

- ▶ **Non-default model database**

The model database is a system database that can be used to set a standard for the creation of default user databases. At this point, you must document the existence of a non-default model database. SQL Server supports one model database per instance.

- ▶ **Heavy usage of tempdb**

Tempdb is a system database that is used for internal SQL Server functionality, such as storing temporary objects and some sort operations. A user database that uses tempdb intensively might affect the performance of the remaining user databases. Consider segregating databases that make intensive use of tempdb into separate SQL Server instances. For more information about the tempdb system database, see SQL Server 2005 Books Online topic “about tempdb.”

- ▶ Backup information

Document the backup requirements for all databases in the environment that you plan to consolidate, such as time window, backup type, frequency, and so forth.

- ▶ Authentication mode

Document whether the instance runs on mixed or windows authentication mode. In a consolidated environment, you are most likely to run your instances in mixed mode. For more information about authentication modes, see SQL Server 2005 Books Online topic “authentication.”

For existing databases, document and take into consideration the following information:

- ▶ Name

This is important because it can help you identify any naming conflict.

- ▶ Type

Document whether the database is an OLTP or OLAP database.

- ▶ Recovery model

Verify the database recovery model. Every user database created uses the recovery model from the model database, which by default, is *Full*. If the database does not make use of transaction log backups, consider changing the recovery model to *Simple*. This way you reduce administration overhead by not having to truncate the database transaction log.

**Note:** Database mirroring requires a *Full* recovery model. See “Database mirroring” on page 62 for details.

- ▶ Collation

Collation refers to a set of rules for sorting and comparing data. Character data is sorted using rules that define the correct character sequence, with options for specifying case-sensitivity, accent marks, kana character types, and character width.

When a database is created in SQL Server, it uses the default collation unless specified differently. For more information about collations, see SQL Server 2005 Books Online topic “about collations.”

- ▶ Sort order

Specify the sort order used for columns that are returned in a SELECT statement. For more information, see SQL Server 2005 Books Online topic “sort orders.”

- ▶ Storage

Document current storage consumption and the estimated growth. This information is important to the design of the storage infrastructure for the consolidated environment.

For each SQL Server login, gather and consider this information:

- ▶ Name

It is important to identify any login name conflict.

- ▶ Type

Is the login a SQL Server or Windows login?

- ▶ Permissions

Document the effective permissions and roles assigned to the login. You must recreate the logins in the consolidated environment.

## Application-specific information

The application-specific information to document during assessment and analysis is:

- ▶ Object linking and embedding database (Oledb) providers  
Oledb providers are programs that access data in the database. To obtain a list of the oledb providers on the server run the query in Example 4-1 in Microsoft SQL Server 2005 Management Studio or Microsoft SQL Server 7.0/2000 Query Analyzer.

*Example 4-1 Listing installed oledb providers with the Management Studio query window*

---

```
use master
go
exec sp_enum_oledb_providers
```

---

- ▶ Applications that do not support usage of SQL Server named instances  
SQL Server 2000 and 2005 support the installation of only one default instance. In a consolidated environment, you are likely to use more than one SQL Server instance per server. Some older applications cannot connect to an SQL Server named instance. The databases for such applications must be placed in a default instance.
- ▶ Application and database on the same OS  
Check for servers that mix applications and databases together. This is a common scenario for a sprawl environment. The best option in this case is to migrate just the database to the consolidated server and leave all the application-specific programs on the current server but make them point to the new database location.
- ▶ Compatibility with SQL Server 2005  
For databases using previous versions of SQL Server, we recommend that you migrate to SQL Server 2005. Check with the application vendors for compatibility with SQL Server 2005. If you cannot upgrade your application so that it is compatible with SQL Server 2005, there is an option in SQL Server where you can run user databases under the compatibility mode. See 4.3.8, “SQL Server upgrade” on page 68 for more information.

## Analysis Services

The Analysis Services information to document during assessment and analysis is:

- ▶ Instance name  
Document the existence of Analysis Service installation and its instance name.
- ▶ Storage  
Document current storage consumption and the estimated growth. This information is important for designing the storage infrastructure for the consolidated environment.

## External dependencies

The external dependencies information to document during assessment and analysis is:

- ▶ Drive lettering constraints  
Assess any drive letter dependency for things such as replication, batch scripts, and Visual Basic® scripts. Drive lettering must be carefully planned for consolidation.
- ▶ Windows Task scheduler  
Record existing jobs involving SQL Server so that you can create them for consolidation.
- ▶ COM objects  
Document any COM objects on the database server. During the analysis phase, you must determine if they will be moved to the consolidated environment.

## Network

The network information you need to gather are the firewall rules. Look for firewall rules that permit the proper access to and from the database servers. You must make sure that the firewall modifications are made before going to production.

### 4.2.2 Sample script to gather information

Example 4-2 shows a script to help you gather the information for your consolidation project.

**Note:** You can download the source of this script at the following FTP site:

<ftp://www.redbooks.ibm.com/redbooks/REDP4385>

*Example 4-2 Transact SQL script to help gather information*

---

```
set nocount on
set dateformat dmy
set nocount on
set concat_null_yields_null off

use master
go

DECLARE @SysInfo INT
DECLARE @DomainName VARCHAR(128)

EXEC sp_OACreate 'ADSystemInfo', @SysInfo OUTPUT
EXEC sp_OAGetProperty @SysInfo, 'DomainShortName', @DomainName OUTPUT

print ''
print '1. General Info'
print '*****'
print ''

print 'Server Name.....: ' + convert(varchar(100), serverproperty('ServerName'))
print 'Physical Server Name.....: ' + convert(varchar(100), serverproperty('MachineName'))
print 'Domain Name.....: ' + isnull(@DomainName, 'NOT ON DOMAIN')
print 'Instance Name.....: ' + convert(varchar(30), @@SERVICENAME)
print 'SQL Server version.....: ' + convert(varchar(100), serverproperty('productversion'))
print 'Edition.....: ' + convert(varchar(100), serverproperty('edition'))
print 'Service Pack Level.....: ' + convert(varchar(100), serverproperty('productlevel'))
print 'Authentication Mode.....: ' + case when serverproperty('IsIntegratedSecurityOnly') = 1 then
'Windows Integrated' else 'Mixed Mode' end
print 'Is Clustered.....: ' + case when serverproperty('IsClustered') = 1 then 'Yes' else 'No'
end
print 'License Type.....: ' + convert(varchar(100), serverproperty('LicenseType'))
print 'Number of Licenses.....: ' + convert(varchar(100), serverproperty('NumLicenses'))
print 'Default Instance Collation.: ' + convert(varchar(100), serverproperty('collation'))
go

print ''
print '2. System Information.'
print '*****'
print ''

exec master..xp_msver
go
```

```

print ''
print '3. Other information'
print '*****'
print ''

print 'Number of connections..: ' + convert(varchar(30),@@connections)
print 'Language.....: ' + convert(varchar(30),@@language)
print 'Lock Timeout.....: ' + convert(varchar(30),@@LOCK_TIMEOUT)
print 'Maximum of connections.: ' + convert(varchar(30),@@MAX_CONNECTIONS)
print 'CPU Busy.....: ' + convert(varchar(30),@@CPU_BUSY/1000)
print 'CPU Idle.....: ' + convert(varchar(30),@@IDLE/1000)
print 'IO Busy.....: ' + convert(varchar(30),@@IO_BUSY/1000)
print 'Packets received...: ' + convert(varchar(30),@@PACK_RECEIVED)
print 'Packets sent.....: ' + convert(varchar(30),@@PACK_SENT)
print 'Packets w errors...: ' + convert(varchar(30),@@PACKET_ERRORS)
print 'TimeTicks.....: ' + convert(varchar(30),@@TIMETICKS)
print 'IO Errors.....: ' + convert(varchar(30),@@TOTAL_ERRORS)
print 'Total Read.....: ' + convert(varchar(30),@@TOTAL_READ)
print 'Total Write.....: ' + convert(varchar(30),@@TOTAL_WRITE)
go

print ''
print '4. Database Information'
print '*****'
print ''

exec master..sp_helpdb
go

PRINT ''
Print '5. Lists of users/groups with DBA privileges'
print '*****'
PRINT ''
DECLARE @Statement VARCHAR(255)
SELECT @Statement = 'SELECT ISNULL(name,loginname) as ' +
    ''System Admin Role Users'' +
    'FROM master..syslogins where sysadmin = 1'

EXEC (@Statement)

IF @@rowcount =0
PRINT '!!!! No One has System Admin rights !!!!'
PRINT ''

print ''
print '6. List of special users per database'
print '*****'
print ''

declare @name sysname,
        @SQL nvarchar(600)

if exists (select [id] from tempdb..sysobjects where [id] = OBJECT_ID ('tempdb..#tmpTable'))
drop table #tmpTable

CREATE TABLE #tmpTable (
    [DATABASE_NAME] sysname NOT NULL ,
    [USER_NAME] sysname NOT NULL,
    [ROLE_NAME] sysname NOT NULL)

```

```

declare c1 cursor for select name from master.dbo.sysdatabases
open c1
fetch c1 into @name
while @@fetch_status >= 0
begin
    select @SQL =
        'insert into #tmpTable
        select N'''+ @name + ''', a.name, c.name
        from ' + QuoteName(@name) + '.dbo.sysusers a
        join ' + QuoteName(@name) + '.dbo.sysmembers b on b.memberuid = a.uid
        join ' + QuoteName(@name) + '.dbo.sysusers c on c.uid = b.groupuid
        where a.name != 'dbo'''
        exec(@SQL)
    fetch c1 into @name
end
close c1
deallocate c1

select * from #tmpTable

drop table #tmpTable
go

print ''
print '7. Datafiles list'
print '*****'
print ''

if exists (select [id] from tempdb..sysobjects where [id] = OBJECT_ID ('tempdb..#TempForFileStats '))
DROP TABLE #TempForFileStats

if exists (select [id] from tempdb..sysobjects where [id] = OBJECT_ID ('tempdb..#TempForDataFile'))
DROP TABLE #TempForDataFile

if exists (select [id] from tempdb..sysobjects where [id] = OBJECT_ID ('tempdb..#TempForLogFile'))
DROP TABLE #TempForLogFile

DECLARE @DBName nvarchar(40)
DECLARE @SQLString nvarchar (2000)
DECLARE c_db CURSOR FOR
    SELECT name
    FROM master.dbo.sysdatabases
    WHERE status&512 = 0

CREATE TABLE #TempForFileStats([Database Name]          nvarchar(40),
                                [File Name]              nvarchar(128),
                                [Usage Type]            varchar (6),
                                [Size (MB)]             real,
                                [Space Used (MB)]       real,
                                [Growth Type]           varchar (12),
                                [File Id]              smallint,
                                [Group Id]             smallint,
                                [Physical File]         nvarchar (260),
                                [Date Checked]         datetime)

CREATE TABLE #TempForDataFile ([File Id]             smallint,
                                [Group Id]            smallint,
                                [Total Extents]       int,
                                [Used Extents]        int,
                                [File Name]          nvarchar(128),

```

```

                                [Physical File]          nvarchar(260))

CREATE TABLE #TempForLogFile ([File Id]          int,
                               [Size (Bytes)]      real,
                               [Start Offset]      varchar(30),
                               [FSeqNo]           int,
                               [Status]           int,
                               [Parity]           smallint,
                               [CreateTime]       varchar(20))

OPEN c_db
FETCH NEXT FROM c_db INTO @DBName
WHILE @@FETCH_STATUS = 0
  BEGIN
    SET @SQLString = 'select ''' + @DBName + ''' + '          as ''Database'', ' +
                    '      f.name, ' +
                    '      CASE ' +
                    '        WHEN (64 & f.status) = 64 THEN ''Log'' ' +
                    '        ELSE ''Data'' ' +
                    '      END          as ''Usage Type'', ' +
                    '      f.size*8/1024.00          as ''Size (MB)'', ' +
                    '      NULL          as ''Space Used (MB)'', ' +
                    '      CASE ' +
                    '        WHEN (1048576&f.status) = 1048576 THEN ''Percentage'' ' +
                    '        ELSE ''Pages'' ' +
                    '      END          as ''Usage Type'', ' +
                    '      f.fileid, ' +
                    '      f.groupid, ' +
                    '      filename, ' +
                    '      getdate() ' +
                    ' FROM ' + @DBName + '.dbo.sysfiles f'

    INSERT #TempForFileStats
    exec(@SQLString)

    SET @SQLString = 'USE ' + @DBName + ' DBCC SHOWFILESTATS'
    INSERT #TempForDataFile
    exec(@SQLString)

    UPDATE #TempForFileStats
    SET [Space Used (MB)] = s.[Used Extents]*64/1024.00
    FROM #TempForFileStats f,
         #TempForDataFile s
    WHERE f.[File Id] = s.[File Id]
          AND f.[Group Id] = s.[Group Id]
          AND f.[Database Name] = @DBName

    TRUNCATE TABLE #TempForDataFile

    SET @SQLString = 'USE ' + @DBName + ' DBCC LOGINFO'
    INSERT #TempForLogFile
    exec(@SQLString)

    UPDATE #TempForFileStats
    SET [Space Used (MB)] = (SELECT (MIN(l.[Start Offset]) +
                                   SUM(CASE
                                       WHEN l.Status <> 0 THEN l.[Size (Bytes)]
                                       ELSE 0
                                       END))/1048576.00
                              FROM #TempForLogFile l
                              WHERE l.[File Id] = f.[File Id])

```



```

FROM #TempForFileStats f
WHERE f.[Database Name] = @DBName
      AND f.[Usage Type]   = 'Log'

TRUNCATE TABLE #TempForLogFile

      FETCH NEXT FROM c_db INTO @DBName
END
DEALLOCATE c_db

SELECT * FROM #TempForFileStats

DROP TABLE #TempForFileStats
DROP TABLE #TempForDataFile
DROP TABLE #TempForLogFile
go

```

Example 4-3 Shows the output for the script on Example 4-2 on page 49:

*Example 4-3 Transact SQL script output*

1. General Info

```

*****
Server Name.....: X3950-SQL\SQLTEST
Physical Server Name.....: X3950-SQL
Domain Name.....: NOT ON DOMAIN
Instance Name.....: SQLTEST
SQL Server version.....: 9.00.3042.00
Edition.....: Developer Edition (64-bit)
Service Pack Level.....: SP2
Authentication Mode.....: Windows Integrated
Is Clustered.....: No
License Type.....: DISABLED
Number of Licenses.....:
Default Instance Collation.: SQL_Latin1_General_CP1_CI_AS

```

2. System Information.

```

*****

```

Index	Name	Internal_Value	Character_Value
1	ProductName	NULL	Microsoft SQL Server
2	ProductVersion	589824	9.00.3042.00
3	Language	1033	English (United States)
4	Platform	NULL	NT AMD64
5	Comments	NULL	NT AMD64
6	CompanyName	NULL	Microsoft Corporation
7	FileDescription	NULL	SQL Server Windows NT - 64 Bit
8	FileVersion	NULL	2005.090.3042.00
9	InternalName	NULL	SQLSERVER
10	LegalCopyright	NULL	© Microsoft Corp. All rights reserved.
11	LegalTrademarks	NULL	Microsoft is a registered trademark of Microsoft Corporation. Windows(TM) is a trademark of Microsoft Corporation
12	OriginalFilename	NULL	SQLSERVER.EXE
13	PrivateBuild	NULL	NULL
14	SpecialBuild	199360512	NULL
15	WindowsVersion	248381957	5.2 (3790)

```

16 ProcessorCount          16          16
17 ProcessorActiveMask    16          ffff
18 ProcessorType          8664        NULL
19 PhysicalMemory         32767      32767 (34358390784)
20 Product ID             NULL        NULL

```

3. Other information

\*\*\*\*\*

```

Number of connections..: 226
Language.....: us_english
Lock Timeout.....: -1
Maximum of connections.: 32767
CPU Busy.....: 13
CPU Idle.....: 9678
IO Busy.....: 3
Packets received...: 713
Packets sent.....: 102789
Packets w errors...: 13
TimeTicks.....: 31250
IO Errors.....: 0
Total Read.....: 47478
Total Write.....: 1744629

```

4. Database Information

\*\*\*\*\*

name	db_size	owner	dbid	created	status	compatibility_level
master	4.50 MB	sa	1	Apr 8 2003	Status=ONLINE, Updateability=READ_WRITE, UserAccess=MULTI_USER, Recovery=SIMPLE, Version=611, Collation=SQL_Latin1_General_CP1_CI_AS, SQLSortOrder=52, IsAutoCreateStatistics, IsAutoUpdateStatistics	90
model	1.69 MB	sa	3	Apr 8 2003	Status=ONLINE, Updateability=READ_WRITE, UserAccess=MULTI_USER, Recovery=FULL, Version=611, Collation=SQL_Latin1_General_CP1_CI_AS, SQLSortOrder=52, IsAutoCreateStatistics, IsAutoUpdateStatistics	90
msdb	7.44 MB	sa	4	Oct 14 2005	Status=ONLINE, Updateability=READ_WRITE, UserAccess=MULTI_USER, Recovery=SIMPLE, Version=611, Collation=SQL_Latin1_General_CP1_CI_AS, SQLSortOrder=52, IsAutoCreateStatistics, IsAutoUpdateStatistics, IsFullTextEnabled	90
tempdb	9.00 MB	sa	2	Dec 6 2007	Status=ONLINE, Updateability=READ_WRITE, UserAccess=MULTI_USER, Recovery=SIMPLE, Version=611, Collation=SQL_Latin1_General_CP1_CI_AS, SQLSortOrder=52, IsAutoCreateStatistics, IsAutoUpdateStatistics	90
testdb	9301.00 MB	X3850-SQL\Administrator	5	Dec 6 2007	Status=ONLINE, Updateability=READ_WRITE, UserAccess=MULTI_USER, Recovery=FULL, Version=611, Collation=SQL_Latin1_General_CP1_CI_AS, SQLSortOrder=52, IsAutoCreateStatistics, IsAutoUpdateStatistics	90

5. Lists of users/groups with DBA privileges

\*\*\*\*\*

System Admin Role Users

-----

```

sa
BUILTIN\Administrators
NT AUTHORITY\SYSTEM
X3950-SQL\SQLServer2005MSSQLUser$X3950-SQL$SQLTEST
X3950-SQL\SQLServer2005SQLAgentUser$X3950-SQL$SQLTEST
X3950-SQL\Administrator

```

6. List of special users per database

\*\*\*\*\*

DATABASE_NAME	USER_NAME	ROLE_NAME
msdb	SQLAgentReaderRole	SQLAgentUserRole
msdb	SQLAgentOperatorRole	SQLAgentReaderRole

7. Datafiles list

\*\*\*\*\*

Database Name	File Name	Usage Type	Size (MB)	Space Used (MB)	Growth Type	File Id	Group Id
Physical File					Date Checked		
master	master	Data	4	2.75	Percentage	1	1
E:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\master.mdf					2007-12-06 16:17:11.920		
master	mastlog	Log	0.5	0.4921875	Percentage	2	0
E:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\mastlog.ldf					2007-12-06 16:17:11.920		
tempdb	tempdev	Data	8	1.625	Percentage	1	1
E:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\tempdb.mdf					2007-12-06 16:17:11.950		
tempdb	templog	Log	1	0.5	Percentage	2	0
E:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\templog.ldf					2007-12-06 16:17:11.950		
model	modeldev	Data	1.1875	1.1875	Pages	1	1
E:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\model.mdf					2007-12-06 16:17:11.950		
model	modellog	Log	0.5	0.4921875	Percentage	2	0
E:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\modellog.ldf					2007-12-06 16:17:11.950		
msdb	MSDBData	Data	6.4375	4.9375	Percentage	1	1
E:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\MSDBData.mdf					2007-12-06 16:17:11.967		
msdb	MSDBLog	Log	1	0.5	Percentage	2	0
E:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\MSDBLog.ldf					2007-12-06 16:17:11.967		
testdb	testdb	Data	5400	5374.813	Pages	1	1
E:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\testdb.mdf					2007-12-06 16:17:11.997		
testdb	testdb_log	Log	3901	3863.5	Pages	2	0
F:\Program Files\Microsoft SQL Server\MSSQL.3\MSSQL\DATA\testdb_log.ldf					2007-12-06 16:17:11.997		

## 4.3 Consolidation project

Now that you have all the detailed information for your database environment, it is time to start planning for the migration. During this phase, you create the technical design for the new consolidated environment. Your technical solution must fulfill two requirements:

- ▶ Business constraints
- ▶ Technical constraints

We recommend that you divide the implementation phase in two. Start with non-mission critical applications before you consolidate the mission-critical ones. This way, your team gains expertise and can identify gaps in the assessment phase that could affect the consolidation.

Your solution must have all the technical information for the new environment, such as:

- ▶ Number of physical servers
- ▶ Number of SQL Server 2005 instances (both database engine and analysis services)

Despite the security advantages offered by separating instances of SQL Server, you should not install multiple instances of SQL Server unless absolutely necessary. Installing multiple instances requires more hardware resources (especially RAM) than simply installing multiple databases in a single instance.

The following situations, however, force you to consolidate on different instances:

- User databases with irreconcilable name conflicts
- User databases that make heavy use of tempdb
- User databases that come from SQL Servers that have nonstandard collations
- Isolation of sysadmin access

- ▶ Storage design

Design the storage technology that you will use for your consolidated environment and the fault tolerance level. See 4.3.3, “Storage” on page 59.

- ▶ Backup strategy

Design your backup strategy so that you meet the recoverability requirements for your business. See 4.3.4, “Backup” on page 61.

- ▶ High availability solution

Design the high availability architecture for the consolidated environment. See 4.3.5, “High availability” on page 62 for more information.

- ▶ Monitoring

Monitoring your consolidated environment for availability and performance is critical for meeting SLAs and delivering a stable and reliable environment. Many professional monitoring applications have a per-server license fee. By consolidating your databases on one server, you can reduce these licensing costs. See 4.3.6, “Monitoring” on page 67.

- ▶ Testing plan

Consolidation should not negatively affect business application operations. Therefore, create procedures to verify that your applications will run as expected in the consolidated environment.

- ▶ Rollback plan

Unrecoverable errors can occur during the consolidation process. As part of your consolidation plan, create a method to revert to the original state before consolidation. A disaster recovery plan should be designed and tested for all mission-critical servers.

### 4.3.1 Critical success factors

The critical success factors for a successful consolidation project include:

- ▶ Not apparent to users

The reasons for a database consolidation project are business oriented. A successful consolidation project should not affect users.

- ▶ Detailed documentation

The documentation must be highly detailed. To minimize the risks that are inherent with implementing such a project, try to map as many potential issues as you can.

- ▶ Careful planning

All steps in the consolidation project must be well planned. Validate your implementation plan during the testing phase and change it to reflect all lessons learned.

- ▶ Enough testing

During the testing phase, you put your plan to proof. Make sure that your test environment is as close as possible to your intended production environment so that you can identify any constraint, gap, or bottleneck. Update your plan with the lessons learned.

- ▶ Management commitment

Management must support the project. In some situations, business decisions will be needed to resolve issues that might come up during the process.

- ▶ Maintained or improved performance

Performance should not be degraded in the new environment. We recommend that you do a baseline on the overall performance of your current application so that you can compare it to the performance after you consolidate your environment. This is not an easy task because direct comparison between CPU, disk, and memory from the old and new environment might not represent relevant data because you have completely different hardware. You should have an initial baseline for the candidate servers so that you have data to compare with after the consolidated environment is in place. See 4.7, "Capacity management" on page 73 for thresholds you should gather.

- ▶ The same level of support

In many cases, departments are unwilling to turn over management of their databases to another group because they feel that they will lose control over their application. Maintaining the same level of support is essential.

- ▶ Realistic SLAs

Determine SLAs so that expectations are set.

- ▶ Identification of mission-critical applications

Mission-critical applications must be clearly identified. A different SLA might be necessary.

### 4.3.2 SQL Server 2005 features for consolidation

There are five editions of SQL Server 2005:

- ▶ SQL Server 2005 Enterprise Edition (32-bit and 64-bit)
- ▶ SQL Server 2005 Standard Edition (32-bit and 64-bit)
- ▶ SQL Server 2005 Workgroup Edition (32-bit only)
- ▶ SQL Server 2005 Express Edition (32-bit only)
- ▶ SQL Server 2005 Developer Edition (32-bit and 64-bit)

Greater use of server resources and product features is possible with these systems, starting with the Express Edition up through Workgroup, Standard, and the Enterprise editions. The Developer Edition contains all the features of the Enterprise Edition, but is licensed for development and testing, not for production use.

Other editions of Microsoft SQL Server 2005 have a limit on the number of physical processors in a server, but Enterprise Edition has no restriction and can take advantage of all processors available from the Windows operating system. Enterprise Edition supports up to 64 CPUs on a single server running on Windows Server 2003 Data Center Edition. It is also available in a 180-day Evaluation Edition. For a database consolidation project, it is most likely that SQL Server 2005 64-bit Enterprise Edition is the best option.

Table 4-1 compares the scalability features between SQL Server 2005 Enterprise and Standard editions.

*Table 4-1 Comparison between SQL server standard and enterprise editions.*

<b>Feature</b>	<b>Standard Edition</b>	<b>Enterprise Edition</b>
Number of CPUs	4	64
Memory	OS Maximum (64 bit)	OS Maximum (64 bit)
Multi-instance support	16	50
AWE memory management	Standard	Standard
Table and index partitioning	Not Available	Standard
Database snapshots	Not Available	Standard
Indexed views	Not Available	Standard
Parallel index operations	Not Available	Standard
Parallel consistency checks (DBCC)	Not Available	Standard
Failover clustering	Two-node only	Up to 8 nodes
Database mirroring (see "Database mirroring" on page 62)	High Availability and High Protection modes only	High Availability, High Protection and High Performance modes.
Fast recovery	Not Available	Standard
Online index operations	Not Available	Standard
Online page and file restore	Not Available	Standard
Parallel index operations	Not Available	Standard
Hot-add memory support	Not Available	Standard
Mirrored backup media	Not Available	Standard
Multi-server administration	Not Available	Standard

For a detailed description of which features are supported in which edition, see:

- ▶ Comparison Between SQL Server 2005 Standard and Enterprise Editions:  
<http://www.microsoft.com/sql/editions/enterprise/comparison.msp>

- ▶ SQL Server 2005 Features Comparison:

<http://www.microsoft.com/sql/prodinfo/features/compare-features.msp>

See 1.2, “SQL Server 2005 and server consolidation” on page 3 for new SQL Server 2005 features that support server consolidation.

**Note:** SQL Server 2005 Enterprise Edition is the most comprehensive edition of SQL Server 2005 and is ideal for the largest organizations and the most complex requirements. We recommend the Enterprise Edition for database consolidation because of its scalability features. If you decide to consolidate using SQL Server 2005 Standard Edition, review the data in Table 4-1 on page 58 to make sure it suits your requirements.

### 4.3.3 Storage

Consolidation of your storage must be carefully planned so that you meet your performance, availability, and scalability requirements. There are currently three types of storage available on the market:

- ▶ Direct attached storage (DAS)
- ▶ Storage Area Network (SAN)
- ▶ Network Attached Storage (NAS)

Your storage plan should take into consideration these different types of storage, along with fault tolerance, file placement, and drive letter assignment.

#### DAS

DAS consists of disks that are directly connected to the server. RAID can be configured at the software level through the host operating system or at hardware level by using a RAID capable controller. We recommend you use hardware-based RAID with a battery backup for the onboard cache.

The main advantages of using DAS are:

- ▶ Less expensive when compared to SAN or NAS
- ▶ Good overall performance

The main disadvantages are:

- ▶ It is more difficult to back up a local drive through a network connection.
- ▶ Moving the local drive to a different server can be difficult if a server fails.
- ▶ It is very difficult to increase a drive size without causing an outage to the system.
- ▶ Scalability is limited.

#### SAN

SANs are normally attached to the servers by Fibre Channel, and the OS sees it as a locally attached device. By providing this capability, OSs can start directly from the SAN itself. It is the best solution for large enterprises for the following reasons:

- ▶ It is very easy to replace a faulty server by assigning the LUNs to a new server.
- ▶ Current SANs can back up live files at byte level directly rather than over the network.
- ▶ They are highly scalable.
- ▶ Drives can be dynamically resized without having to cause an outage.
- ▶ Administration is centralized.

SAN storage combined with Windows clustering is the most reliable and scalable solution for large mission-critical database environments. Its main disadvantage is the high price when compared to DAS.

### Network Attached Storage (NAS)

NAS is somewhat similar to SAN, but it does not provide a dedicated drive to the SQL Server. Instead, it provides a share through a network connection. NAS is primarily used for centralization of large amounts data that do not perform write-intensive operations, such as a file share. For large enterprise database environments and those that make use of Windows and SQL Server clustering, we recommend SAN storage. We do not recommend using NAS to place the SQL Server 2005 database data and log files.

**Tip:** For large-capacity and high-performance storage options, there are many IBM System Storage™ options available including DAS products, SAN products, and NAS products. For information, see:

<http://www.ibm.com/storage/snetwork>

### Fault tolerance

All current storage technologies provide fault tolerance with a RAID configuration. There are several levels of RAID configurations, but the most common for database servers are RAID levels 1, 5, and 10.

RAID-1 is also called a mirror. It combines two physical hard drives into one logical unit and replicates the data to both physical drives. In case one drive fails, the data is still available. It is very fast and robust and is commonly recommended for OS files and database log files.

RAID-5 is a striped set that uses a minimum of three disks. It is fault tolerant because it relies on distributed parity. It can operate with all but one drive, but it will be vulnerable until the failed drive is replaced and the data is rebuilt on it. This RAID level is not recommended for I/O intensive files, such as database files, because it needs to write parity to all disks.

RAID-10 is a striped set of mirrors (RAID-1), and it requires a minimum of four disks to operate. With RAID 10, all your drives can fail and the data will still be available. It is recommended for I/O intensive files such as SQL Server data files.

Table 4-2 compares the RAID array types:

Table 4-2 RAID Array comparison

RAID type	Fault tolerance	Performance
RAID-1 (mirrored volume)	One disk may fail	Good read, acceptable write
RAID-5 (stripe set with parity)	One disk may fail	Excellent read, relatively poor write
RAID-10 (stripe of mirrors)	Half of the disks could fail (but not both disks in any mirrored pair)	Excellent read and write

**Note:** Although RAID level 10 is the best option for placing your data or log files, it might be very expensive if you plan to deploy a SAN storage environment. Our experience is that RAID level 5 normally performs well and has good reliability with SAN storage.

Also, we do not recommend using RAID level 5 on DAS for placing data or log files in a consolidated environment.



## File placement and drive letter assignments

For a consolidated environment, you must plan the placement of your data and log files carefully so that you balance I/O performance. For example, place data files, log files, and tempdb database data and log files on separate logical and physical drives. This eases administration and increases performance.

Table 4-3 presents drive lettering suggestions.

Table 4-3 Recommendation for drive letter assignments

Logical drive letter	Description
C	Use for operating system and SQL Server binary files.
D	This is normally a CD-ROM drive.
E	Reserve for a system file such as a paging file.
F through H	Use for tempdb data and log files.
I through P	Use for data files.
Q	Use for quorum (clustering).
R, S	Use for SQL Server executables and system databases.
T through V	Use for transaction log files.
X through Z	Reserve for working area. Place files such as temporary backup files or temporary data for bulk import.

### 4.3.4 Backup

Recoverability is the ability to restore data after a data loss incident such as a server failure. Designing a recoverable database server requires adequate planning and preparation so that you can meet all the business requirements for data recovery.

**Important:** Keep in mind that a recoverability strategy is not just the backup of data. You need to make sure that your backups work. Perform test restores of your data to ensure that your procedures work.

These are what we consider best practices for a backup solution on a consolidated environment:

- ▶ Have the minimum backup jobs as possible. Consolidate the backups whenever business rules allow. This way, you reduce administrative overhead by reducing the job configuration complexity.
- ▶ Only back up transaction logs for databases that require them. If business does not require transaction log backups for a certain database then do not back it up. Having the least backup jobs possible reduces administrative overhead.

**Note:** If you are not going to back up the transaction logs for a certain database, then set its recovery model to *Simple* so that you do not have problems caused by excessive log growth. See the Microsoft article, “A transaction log grows unexpectedly or becomes full on a computer that is running SQL Server” for the most common reasons why transaction logs become full. The article is available from:

<http://support.microsoft.com/kb/317375/en-us>

- ▶ Make sure you back up master, model, and msdb databases. They contain information relevant to the instance installation, SQL jobs, and defaults when you create a new user database.
- ▶ Make sure you back up any additional data. You might have additional services on your database server, such as SQL Server Notification Services (SSNS), SQL Server Analysis Services (SSAS), and SQL Server Reporting Services (SSRS). All these additional services have different backup requirements.
- ▶ Schedule backups during low periods of activity. Identify the lowest period of activity for the consolidated server so that you schedule backups that do not impact business.
- ▶ Select an adequate medium for storing backup information. Choosing hardware for backups is essentially a matter of selecting either tape or hard disk. Magnetic tape generally offers greater storage capacity and lower cost when compared to hard disk. Hard disks offer better performance.
- ▶ Store backups off site as a secondary measure. Storing your backups in the same location as your database server protects you from server failures, accidental deletions, and other problems that occur on the database server itself. When the backups are in an easily accessible location, you can perform a fast restore, which is important to disaster recovery. However, disasters such as fires or floods can occur that affect the entire site. If such a disaster occurs, the backups stored at this site can easily be destroyed along with your original data. For this reason, we recommend storing a copy of your most important data off site. Consider a site that is less susceptible to floods or other natural disasters than your primary site but has the same level of security.
- ▶ Create a restore routine that constantly tests your backups. It is very common for organizations to discover that their backups do not work when they need to recover from a disaster because they did not test the process.

### 4.3.5 High availability

High availability is more critical in consolidated environments than in regular environments. Because most of your databases are on a single system, a hardware failure can have a greater effect on your business. Therefore, your consolidated environment requires a high availability solution.

SQL Server 2005 has two fault tolerance features for high availability:

- ▶ Database mirroring
- ▶ Clustering

#### Database mirroring

Database mirroring is a new feature with SQL Server 2005. In database mirroring, an originating SQL Server 2005 instance continuously sends database transaction log records to a copy of the database on another standby SQL Server instance. The originating database and server have the role of principal, and the receiving database and server have the role of mirror. The principal and mirror servers must be separate instances of SQL Server 2005.

The benefits are:

- ▶ There is no reliance on matching hardware.
- ▶ The database that is being mirrored does not have to be located on shared storage. There are two separate databases on two separate servers, so data and log files are not shared.
- ▶ There is no distance limitation, because the mirroring is done over an Ethernet network.

- ▶ There is automatic failover capability when the mirroring is operating in High Availability mode.
- ▶ Setup is straightforward and does not require previous OS and hardware setup.

Here are some important items to note about database mirroring:

- ▶ The principal database must be in the *Full* recovery mode. Log records that result from bulk-logged operations cannot be sent to the mirror database.
- ▶ The mirror database must be initialized from a restore of the principal database with NORECOVERY, followed by restores in the sequence of principal transaction log backups. If you copy the backup files through a LAN or WAN connection, make sure you have enough network bandwidth.
- ▶ The mirror database must have the same name as the principal database.
- ▶ Because the mirror database is in a recovering state, it cannot be accessed directly. You can create database snapshots on the mirror to indirectly read the mirror database at a point in time. Database snapshots are only available in SQL Server 2005 Enterprise Edition. See Table 4-1 on page 58.
- ▶ It is not possible to back up a snapshot database. If there is a media failure, it might not be possible to recreate the snapshot at the same point-in-time for reporting purposes. For more information about database snapshots, see SQL Server 2005 Books Online topic “Using database snapshots with database mirroring.”

There are three possible operating modes when setting up database mirroring:

- ▶ High Availability
- ▶ High Protection
- ▶ High Performance

The operating mode is determined by the existence or non-existence of a witness server, the transaction safety level, and the type of failover. Table 4-4 shows the characteristics for the three possible operating modes for a database mirroring configuration.

Table 4-4 Possible modes for setting up database mirroring.

Operating mode	Transaction safety	Transfer mechanism	Witness server	Failover type
High Availability	Full	Synchronous	Required	Automatic or manual
High Protection	Full	Synchronous	Not Required	Manual only
High Performance	Off	Asynchronous	Not Applicable	Forced only

The *witness server* is a third SQL Server 2005 installation that is used by the mirroring session in the High Availability operating mode to form a quorum with the principal database. It is responsible for monitoring the health of the principal database and starting the automatic failover process if it identifies that the current principal database is unavailable.

**Tip:** SQL Server 2005 Express Edition can be used as a witness server for configuring a database mirroring session in High Availability mode.

*Transaction safety* determines whether the communication between the principal and the mirror is synchronous or asynchronous. Synchronous transfer guarantees that all transactions in the mirror database transaction log are synchronized with the principal database transaction log, and so the transactions are considered safely transferred.

If transaction safety is set to *Full*, the principal and mirror servers operate in a synchronous transfer mode. As the principal server writes its principal database log records to disk, it also sends them to the mirror. The principal then waits for a response from the mirror server, which responds when it has written those same log records to its log disk.

When safety is set to *Off*, the communication between the principal and the mirror is asynchronous. The principal server does not wait for an acknowledgment from the mirror that the mirror has written a block of transaction records. The mirror tries to keep up with the principal by recording transactions as quickly as possible, but some transactions might be lost if the principal suddenly fails and you force the mirror to become the principal database. Therefore, when the safety is set to *Off*, the principal and mirror might not be fully synchronized if the mirror does not keep up with the principal.

By default, the transaction safety level is set to *Full* when database mirroring is initially configured. To set database mirroring to High Performance mode, you set transaction safety to *Off* with TSQL. Example 4-4 shows the command.

*Example 4-4 Turning transaction safety off through Management Studio query window.*

---

```
ALTER DATABASE <dbname> SET PARTNER SAFETY OFF  
go
```

---

**Note:** The command in Example 4-4 does not work in SQL Server 2005 Standard Edition. See Table 4-1 on page 58.

The process of *failover* is when the mirror database becomes the principal and a writable database. There are three failover types for a database mirroring session:

► Automatic

For this type of failover, the witness server queries the principal database regularly to check for its responsiveness. If the principal database becomes unresponsive, the mirror database forms a quorum with the witness server and will become principal. No data is lost. Note that automatic failover is only available in High Availability operating mode.

► Manual

A DBA can manually failover the principal role to the mirror database so that it becomes the principal and no data is lost. Available for both High Availability and High Protection operating modes, it is the *only* option for High Protection operating mode.

► Forced

The forced type of failover is the only one capable of recovering a mirrored database running in High Performance operating mode. Because transaction safety is off, SQL Server cannot guarantee that all transactions are recorded in the database log file by the time of the failover. Example 4-5 on page 65 shows how to execute a forced failover.

### **High Availability operating mode**

The High Availability operating mode supports maximum database availability with automatic failover to the mirror database if the principal database fails. It requires that you set safety to *Full* and define a witness server as part of the database mirroring session.

The High Availability mode is well-suited to a configuration with fast network communication between the servers that requires automatic failover for a single database. When safety is *Full*, the principal server must wait briefly for responses from the mirror server, and therefore the performance of the principal server can be affected.

In the High Availability mode, database mirroring is self-monitoring. If the principal database becomes unavailable, or the server for the principal is down, then the witness and the mirror form a quorum of two and the mirror SQL Server performs an automatic failover.

At that point, the mirror server instance becomes the new principal and recovers the database. The mirror server can become available quickly because the mirror has been replaying the transaction logs for the principal and its transaction log has been synchronized with the principal.

### **High Protection operating mode**

The High Protection operating mode has transactional safety set to *Full*, but has no witness server as part of the mirroring session. In this mode, only a manual failover is possible, because there is no witness to determine if the primary server has become unresponsive.

We recommend this mode in two situations:

- ▶ When you do not want to install the third SQL Server 2005 instance that is required for the High Availability mode.
- ▶ When you do not want your principal databases to failover automatically.

### **High Performance operating mode**

In the High Performance operating mode, transactional safety is *Off*, and the transfer of log records is asynchronous. The principal server does not wait for an acknowledgement from the mirror that all transaction log records have been recorded on the mirror. The mirror does its best to keep up with the principal, but it is not guaranteed that all the most recent transactions from the principal will have been written in the transaction log for the mirror.

Because there is no witness server configured in the High Performance mode, automatic failover is not possible. Because transaction safety is set to *Off*, manual failover is not possible either. The only type of failover in this case is forced failover. The forced failover causes immediate recovery of the mirror database.

Example 4-5 shows how to execute a forced failover.

*Example 4-5 Performing forced failover.*

---

```
ALTER DATABASE <dbname> SET PARTNER FORCE_SERVICE_ALLOW_DATA_LOSS  
go
```

---

**Note:** Forced failover might create data loss on the mirror when it is recovered, if some of the transaction log blocks from the principal have not yet been received by the mirror. We recommend using High Performance mode is for transferring data over long distances or for mirroring very active databases where some potential data loss is acceptable.

For more information about forced failover, see SQL Server 2005 Books Online topic “Forced service.”

## **Clustering**

Failover clustering provides hardware redundancy by transferring mission-critical resources from a failing machine to an equally configured server automatically. It is configured at the OS level and provides a virtual layer for cluster-aware applications such as SQL Server 2005.

The main new features in SQL Server 2005 failover clustering are:

- ▶ SQL Server 2005 Analysis Services is now cluster aware.

**Note:** You can install both Microsoft SQL Server 2005 Database Engine and Analysis Services on a failover cluster using the Microsoft SQL Server 2005 setup. You can also configure Notification Services and Integration Services to run on a failover cluster manually. Refer to the next section for setup documentation.

- ▶ Both SQL Server 2000 and SQL Server 2005 can be installed in failover clustering configurations side-by-side on the same Windows Server cluster.

**Note:** When you use SQL Server 2000 and SQL Server 2005 on the same cluster, SQL Server 2000 must be installed before you install SQL Server 2005 instances locally or as failover cluster instances. All failover cluster instances use the SQL Server 2005 Server Cluster resource. This includes the instances of SQL Server 2000

- ▶ SQL Server 2005 Standard Edition supports a two-node cluster configuration. Enterprise Edition supports up to eight nodes in a cluster configuration. See Table 4-1 on page 58.
- ▶ Depending on your available resources (such as disk) and edition of SQL Server 2005, up to 23 clustered instances of SQL Server 2005 can be deployed.

**Note:** Each instance of SQL Server 2005 in a clustered environment must have at least one dedicated drive. Disks cannot be shared by instances because they can only exist in a single cluster group, and only one SQL Server 2005 installation can be placed in a single cluster group. Therefore, you are limited to a maximum of 22 or 23 usable drive letters. (You must disregard the local system drive, CDROM if exists, and shared quorum drive for the cluster resources.) Unlike that of SQL Server 2000, SQL Server 2005 failover clustering does support the use of mounted volumes. All disks must be formatted as Basic disks. Dynamic disks are not supported in a cluster installation.

- ▶ WOW mode of 64-bit Windows Server 2003 (x64) is supported for 32-bit failover clusters.
- ▶ The new Database Mail feature of SQL Server 2005 replaces SQL Mail and makes e-mail support in a clustered configuration possible.
- ▶ Microsoft Virtual Server is supported for use with failover clustering (two node configuration only).
- ▶ You can rename a failover clustering instance without uninstalling/reinstalling SQL Server.
- ▶ Each installation of SQL Server 2005 in a clustered configuration has its own full-text resource.

### ***Failover Cluster configuration documentation reference***

The following documents can help you configure failover and clustering:

- ▶ Guide to Creating and Configuring a Server Cluster under Windows Server 2003:  
<http://www.microsoft.com/downloads/details.aspx?familyid=A5BBB021-0760-48F3-A53B-0351FC3337A1&displaylang=en>
- ▶ Server Clusters: Cluster Configuration Best Practices for Windows Server 2003:  
<http://www.microsoft.com/downloads/details.aspx?familyid=98BC4061-31A1-42FB-9730-4FAB59CF3BF5&displaylang=en>

- ▶ SQL Server 2005 Failover Clustering White Paper:  
<http://www.microsoft.com/downloads/details.aspx?FamilyID=818234DC-A17B-4F09-B282-C6830FEAD499&displaylang=en>
- ▶ Microsoft support policy for clustering:  
<http://support.microsoft.com/kb/327518/en-us>.
- ▶ SQL Notification Services cluster setup:  
<http://msdn2.microsoft.com/en-us/library/ms171407.aspx>
- ▶ SQL Integration Services cluster setup:  
<http://msdn2.microsoft.com/en-us/library/ms345193.aspx>

**Important:** SQL Reporting services is not cluster aware and we recommend that you install it on a separate server.

Failover clustering is a powerful high availability solution and we recommend using it for consolidated environments.

### 4.3.6 Monitoring

Centralized monitoring is one of the main advantages of a consolidated environment. You have fewer servers and SQL Server instances to monitor. At the same time, it is more critical because you have more SQL Server instances and databases in the same system.

IBM Tivoli® Monitoring is an excellent tool for managing performance and availability of your IT environment, and it can help you identify performance bottlenecks, services availability, and more. For more information, visit:

<http://www-306.ibm.com/software/tivoli/products/monitor>

IBM Director is another option. This centralized administration software for IBM systems can help you maximize system availability. For more information, see *Implementing IBM Director 5.20*<<Citation>, SG24-6188:

<http://www.redbooks.ibm.com/abstracts/sg246188.html>

### 4.3.7 Online analytical processing and online transaction processing

SQL Server 2005 Analysis Services is a Microsoft tool for online analytical processing (OLAP) and data mining. OLAP databases provide multidimensional views of data for the purpose of data analysis. Online transaction processing (OLTP) databases are transaction based and provide very detailed view of data.

Table 4-5 compares the differences between OLTP and OLAP databases:

Table 4-5 Differences between OLTP and OLAP databases

Characteristic	OLTP	OLAP
Operation	Transaction	Analysis
Data granularity	Detailed	Aggregated
Number of users	High	Low
Query type	Pre-defined	Ad-hoc

Because of the different characteristics of both database types, we recommend that you put your OLTP and OLAP databases on separate servers. If consolidating your OLAP and OLTP environments is not a viable option, we recommend that you keep OLAP and OLTP database files on separate physical drives for better I/O throughput.

### 4.3.8 SQL Server upgrade

A sprawl SQL environment is most likely to have different versions of SQL Server. For the consolidated database server, the ideal is to have only SQL Server 2005 instance. SQL Server 2005 provides features and tools for upgrading databases from previous versions of SQL Server, along with two upgrade methods: the in-place method and side-by-side method.

The in-place method installs SQL Server 2005 directly over a previous version of SQL Server. This type of upgrade is possible for the following SQL Server versions:

- ▶ SQL Server 7.0 Service Pack 4 (SP4) or later.
- ▶ SQL Server 2000 SP3, SP4, or later.

**Note:** SQL Server 6.5 databases and instances must be manually upgraded.

The side-by-side method installs a SQL Server 2005 instance so that it runs alongside the previous SQL Server version and you manually migrate data files and other necessary database objects.

See the Microsoft publication, “SQL Server 2005 Upgrade Technical Reference Guide” for detailed technical guidance on SQL Server upgrade, available from:

<http://www.microsoft.com/downloads/details.aspx?FamilyID=3d5e96d9-0074-46c4-bd4f-c3eb2abf4b66&DisplayLang=en>

#### Microsoft SQL Server 2005 Upgrade Advisor

This Microsoft tool analyzes possible upgrade problems from versions 7.0 and 2000 and provides links to documentation that can help you resolve the issues. For download and information about system requirements, visit:

<http://www.microsoft.com/downloads/details.aspx?familyid=1470e86b-7e05-4322-a677-95ab44f12d75&displaylang=en>

#### Backward compatibility mode

In addition to migrating to SQL Server 2005, you can run databases from any version of SQL Server under compatibility mode, although we do not recommend running databases under 6.0 and 6.5 in compatibility mode in SQL Server 2005. Most of the features that you need to administer 6.0 and 6.5 are not available in SQL Server 2005 Management Studio and the next versions of SQL Server will not support 6.0 and 6.5 compatibility modes. For more information about backward compatibility mode, see SQL Server 2005 Books Online topic “About backward compatibility.”

### 4.3.9 NUMA Considerations

NUMA is an architectural design where more than one system bus accesses a small set of processors. Each set of processors is called a NUMA node and it has its own memory and, in some cases, its own I/O channels. Accessing the local memory for the node is faster than accessing memory from another node.



In SMP, all memory access is posted to the same shared memory bus. This design works well for a small number of CPUs, but not when you have a large number of CPUs accessing the shared memory bus simultaneously. NUMA reduces these bottlenecks by limiting the number of CPUs on a single memory bus and by using high-speed interconnection to connect the NUMA nodes. There are two possibilities for NUMA configurations: hardware NUMA and soft-NUMA.

x3950 M2 delivers the latest IBM NUMA architecture by connecting the scalability ports of each node (see Figure 1-1 on page 6) directly to the memory controller in each node and providing high-speed communication between processors in different nodes. This is called *hardware NUMA* because the ports act as though they were hardware extensions to the CPU local buses. They direct read and write requests to the appropriate memory or I/O resources, and they also maintain cache coherency between the processors. Refer to 3.1.5, “NUMA” on page 35 for detailed information about the NUMA features of x3950 M2.

A new SQL Server 2005 feature can logically partition one physical node into multiple NUMA nodes. This is referred to as *soft-NUMA*. The benefits of soft-NUMA include reducing I/O and lazy writer bottlenecks on computers with many CPUs and no hardware NUMA. When you configure a soft-NUMA node, it will have one I/O thread and one lazy writer thread for each node. This configuration can increase performance for heavily used databases.

For more detailed information about soft-NUMA, see:

- ▶ “How SQL Server 2005 Supports NUMA”:  
<http://msdn2.microsoft.com/en-us/library/ms180954.aspx>
- ▶ “How to: Configure SQL Server to Use Soft-NUMA”:  
<http://msdn2.microsoft.com/en-us/library/ms345357.aspx>

**Important:** The following information about soft-NUMA should be taken into consideration:

- ▶ SQL Server 2005 soft-NUMA configuration is typically only used on non-NUMA hardware with a large number of processors. SQL Server 2005 is a NUMA-aware application and automatically configures NUMA features when running on NUMA hardware such as IBM System x3950 M2.
- ▶ On servers such as the x3950 M2, soft-NUMA can still be useful when contention for CPU resources arises in a single instance. Using soft-NUMA to partition CPUs by workload in a single SQL Server instance can resolve this issue.
- ▶ When SQL Server 2005 starts, it automatically configures NUMA features based on the hardware NUMA boundary exposed by Windows. Informational messages describing the NUMA configurations are added to the SQL Server error logs.
- ▶ Soft-NUMA only applies to SQL Server database engines. SQL Server Analysis Services before SP2 and SQL Server Integration Services are not NUMA aware.

## Processor and I/O affinity

By default, no specific processor affinity is set in SQL Server 2005, and all processors can be scheduled to perform all tasks. SQL Server 2005 still attempts to localize the use of resources to take advantage of the NUMA design of x3950 M2.

To define how system resources are used precisely, you can enable hard processor affinity in SQL Server 2005. You can configure the affinity setting per SQL Server instance.

Refer to the following Microsoft documents for detailed information about setting up processor affinity:

- ▶ How to: Configure the Affinity Mask (SQL Server Management Studio):  
<http://msdn2.microsoft.com/en-us/library/ms186255.aspx>
- ▶ Setting Server Configuration Options  
<http://msdn2.microsoft.com/en-us/library/ms189631.aspx>

You can also set affinity to disk I/O with I/O affinity, which can associate SQL Server disk I/O to a specified subset of processors so that they handle all disk I/O. This option can work effectively in the OLTP environment with high loads, especially in cases of a server with 16 or more processors.

I/O affinity can enhance the performance of SQL Server threads that are issuing I/O. For example, you could assign CPU 1 for I/O affinity and all the other CPUs for processor affinity. However, it is important to remember that I/O affinity does not always improve performance and it does not support hardware affinity for individual disks or disk controllers.

### Network affinity

Network affinity is a new feature in SQL Server 2005. This feature can connect clients to specific NUMA nodes. The settings can be configured using SQL Server Configuration Manager GUI or Windows registry.

To configure network affinity in SQL Server 2005, use a binary mask that represents nodes. The mask has a bit for each node as 76543210 with the first node as zero (Figure 4-1).

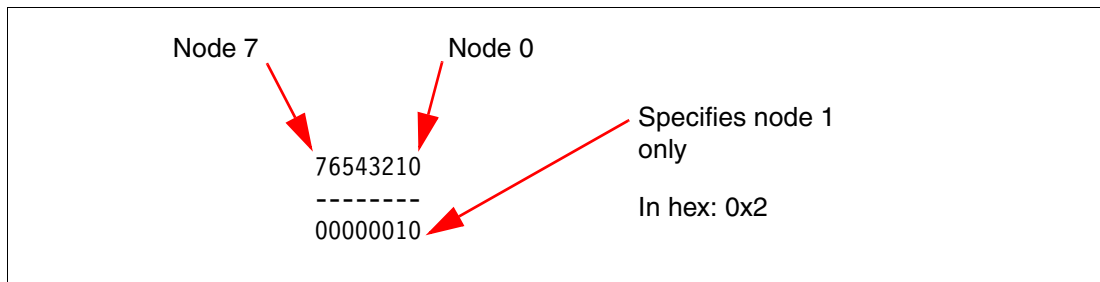


Figure 4-1 Determining the affinity mask to use for network affinity

Set each node bit to 1 to select a node or 0 to not select a node. For example:

- ▶ To specify only node 0, use mask 00000001 or hex 0x1.
- ▶ To specify only node 1, use mask 00000010 or hex 0x2.
- ▶ To specify nodes 0, 2, and 4, use mask 00010101 or hex 0x15.

To specify all nodes, use a mask of -1 or leave the field blank.

**Tip:** You do not have to use different port numbers.

Refer to the following Microsoft document for detailed information about setting up network affinity: How to: Map TCP/IP Ports to NUMA Nodes, available at:

<http://msdn2.microsoft.com/en-us/library/ms345346.aspx>

## Windows System Resource Management (WSRM)

WSRM is a tool that provides resource management of processors and memory resources. With WSRM, you can specify which CPUs you want each process to run on.

You can install WSRM from the Windows Server 2003 CD-ROM. WSRM is included with Windows Server 2003 Enterprise Edition and Datacenter Edition and is covered by the license for these products.

Alternatively, you can download it from:

<http://www.microsoft.com/windowsserver2003/downloads/wsrn.msp>

Table 4-6 lists the editions of Windows Server 2003 and whether they can install WSRM.

Table 4-6 Windows 2003 Server editions that include WSRM

Edition	x86 (32-bit)	x64 (64-bit)	IA64 (64-bit)
Windows 2003 Web Edition	No	No	No
Windows Server 2003 Standard Edition	No	No	No
Windows Server 2003 Enterprise Edition	Yes	Yes	Yes
Windows Server 2003 Datacenter Edition	Yes	Yes	Yes
Windows 2000 Family	No	No	No

**Restriction:** Do not configure CPU affinity or SQL Server memory in Windows with WSRM and also configure the affinity mask in SQL Server. These settings attempt to achieve the same thing, and if the configurations are inconsistent, you might have unpredictable results. Do not use WSRM to control SQL Server memory for the same reasons.

## 4.4 Testing

Careful testing is a critical success factor for your consolidation project. After you have the project plan, create a testing environment that matches your actual environment as closely as possible and use the a load for the consolidated test server that is the same for that of the production environment.

Remember that you must test your application functionality on the consolidated server and not just SQL Server itself. By doing this, you can identify any problems with your project plan and identify issues not considered during assessment and analysis.

**Important:** Make sure you modify your plan along the way with the lessons learned during the testing phase.

## 4.5 Implementation and stabilization

During this phase, you deploy your designed and tested consolidated environment and stabilize the applications in the production environment. We recommend that you not perform all your consolidations at once, or even in a short period. Start with non-mission critical databases and applications and then move to mission-critical. Your team will gain experience and the business impact will not be as great if you must rollback your changes.

Here are some important best practices for this phase:

- ▶ Give enough time for stabilization. Make sure the databases and application you migrate to the consolidated environment are stable before you move towards the next ones.
- ▶ Have a rollback plan for every step of your migration and test it before implementation. Even after assessing, developing, and testing your plan, an unforeseen event might occur or a hole might be found in your design.
- ▶ Communicate clearly. Make sure that users are notified of any possible outages and that people are aware of all actions that will be taken during the deployment.

## 4.6 Best practices

We have gathered the following best practices for consolidating a database environment:

- ▶ Check for problematic databases before consolidation. Otherwise, you might be accountable for any problems that might have been happening before consolidation.
- ▶ Separate data files from log files for better performance. Disk I/O will probably become a bottleneck, so make sure that the files do not compete for resources. Put the tempdb database on a fast I/O subsystem, stripe it over multiple disks for better performance, and move it to disks that are different from those for user databases.
- ▶ Do not back up databases directly to the same disk where data and log files are located because it degrades performance.
- ▶ Run SQL services under a domain account with limited privileges, not as “Local system,” and provide the minimum rights necessary for running SQL Server. Grant “Lock pages in memory” privilege to the SQL account to prevent the OS from paging SQL Server buffer pool memory during periods of high memory utilization. For more information, see:

- Lock pages in memory

<http://technet2.microsoft.com/windowsserver/en/library/e72dcd66-fe17-49dd-a382-02baad31a1371033.msp?mfr=true>

- How to: Enable the Lock Pages in Memory Option (Windows)

<http://msdn2.microsoft.com/en-us/library/ms190730.aspx>

- ▶ Use Microsoft Best Practices Analyzer to analyze SQL Server 2000 and SQL Server 2005 for best practices compliance. It tells you whether your SQL Server instance is compliant with Microsoft best practices and how to deal with any non-compliance. Download it from:

<http://www.microsoft.com/downloads/details.aspx?FamilyID=b352eb1f-d3ca-44ee-893e-9e07339c1f22&displaylang=en>

Download Microsoft Best Practices Analyzer for SQL Server 2005 from:

<http://www.microsoft.com/downloads/details.aspx?familyid=da0531e4-e94c-4991-82fa-f0e3fbd05e63&displaylang=en>

- ▶ Do not install previous SQL Server versions on the same server, especially SQL Server versions 6.5 or 7.0. The more different SQL Server versions you have on the same system, the greater the potential for stability issues.  
Consider running previous versions of SQL Server databases in compatibility mode in SQL Server 2005. If there is no better option, combine SQL Server 2005 and SQL Server 2000 in the same system. See 4.3.8, “SQL Server upgrade” on page 68 for more details about upgrading your databases.
- ▶ Do not install application-specific programs on the same server as the database server. If you are migrating the database from a dedicated application server that has several other applications, move only the SQL Server.
- ▶ Named instances of the SQL Server database engine are configured for dynamic ports, so define static TCP ports for SQL Server 2005 named instances. If enabled, the default instance of the SQL Server database engine listens for connections on UDP port 1434 and responds to requests on TCP port 1433.
- ▶ If users connect to a named instance through a firewall, configure its database engine to listen on a specific port so that you can open the appropriate port in the firewall. (You use the SQL Server Configuration Manager to assign port numbers.) Consider using a value above 32000 because many of those below 32000 are reserved for well-known services.

## 4.7 Capacity management

According to Microsoft Operations Framework (MOF), which is available from the Microsoft Web site,<sup>1</sup> capacity management is:

The process of planning, analyzing, sizing, and optimizing capacity to satisfy demand in a timely manner and at a reasonable cost. This process should be proactive and responsive to business needs because the business cannot add resources after a capacity problem has happened without impacting performance.

Effectively monitoring the size of your databases is a key element in the management of your storage space capacity.

Here are some aspects that you need to cover:

- ▶ Disk drive capacities
- ▶ Size of databases, size of free space inside the database
- ▶ Difference between the database size compared to the disk drive space
- ▶ The rate of database growth
- ▶ Location (drives) where data files are located (\*.MDF, \*.NDF, \*.LDF)

After you implement the consolidated server, you must monitor the resource use so that you can identify possible bottlenecks and plan for hardware upgrade ahead of time.

Table 4-7 on page 74 shows the most common performance counters for SQL Server capacity management. You can monitor these resources using Windows Performance console.

<sup>1</sup> See <http://www.microsoft.com/technet/solutionaccelerators/cits/mo/smf/smfcapmg.aspx>

Table 4-7 Performance counters for SQL Server capacity management

Counter	Recommended threshold
Memory: Available Bytes	Should remain above 50 MB.
Physical Disk: % Disk Time	Should average below 50%.
Physical Disk: Avg. Disk Queue Length	Should average below 2 per disk. For example, for an array of 5 disks, this figure should average below 10.
Physical Disk: Avg. Disk Reads/sec Physical Disk: Avg. Disk Writes/sec	Used to size the disk and CPU, this should be below 85% of the capacity of the drive. Of the I/Os handled by OLTP systems, an average of 66% are reads and 33% are writes (although this percentage varies considerably from system to system). In a typical decision support system (DSS), 90% are reads, and only 10% are writes. If you know only the total number of I/Os on a system, you can use these averages to help estimate the number of reads and writes on the system.
Network Interface: Bytes Total/sec	Used to size the network bandwidth.
SQL Server: Buffer Manager: Buffer Cache Hit Ratio	Should exceed 90% (and ideally approach 99%).
SQL Server: Buffer Manager: Page Life Expectancy	Used to size memory. Should remain above 300 seconds.
SQL Server: Databases: Data File(s) Size KB	Used to size the disk subsystem. This threshold can tell you how much the data files are growing so you can project the growth and prepare for disk acquisition if necessary.
SQL Server: Databases: Log File(s) Size KB	Used to size the disk subsystem. This threshold can tell you how much the data files are growing so you can project the growth and prepare for disk acquisition if necessary.

Besides collecting statistics with performance counters, you can learn the current growth, such as 100 GB a year or 10% per month, by interviewing application owners. With this type of information you can calculate the linear and geometric growth rates.

### Linear growth

Linear growth calculates a static growth rate for a specific amount of time, such as 100GB a year. Projecting this growth is a simple matter of multiplying the amount of increases by the number of time periods. The equation for calculating linear growth is:

$$\text{Future usage} = \text{current usage} + (\text{growth amount} \times \text{number of periods})$$

Suppose the database currently handles 100 transactions per minute and increases 20 transactions per minute per year and you want to know how many transactions per minute it will be handling in 5 years. Example 4-6 shows how to calculate this linear growth.

#### Example 4-6 Calculating linear growth

---


$$x = 100 + (20 \times 5)$$

$$x = 100 + (100)$$

$$x = 200$$


---

## Geometric growth

Geometric growth is characterized by an increase in a steady percentage per time period, such as an increase of the database size of 3% each month. The equation for calculating geometric growth is:

$$\text{future\_usage} = \text{current\_usage} \times (1 + \text{growth\_rate})^{\text{number\_of\_periods}}$$

Be sure to express the growth rate as a decimal value. For example, if the database is currently 200 GB and grows at the rate of 3% per month, you can calculate the size of the database in 2 years (24 months) by using the formula in Example 4-7.

*Example 4-7 calculating geometric growth*

---

$$\begin{aligned}x &= 200 \times (1 + 0,3)^{24} \\x &= 200 \times (1,03)^{24} \\x &= 200(2,03) \\x &= 406\end{aligned}$$

---

## 4.8 Services offerings

IBM Integrated Technology Services has several fee-based offerings relevant to the IBM System x3950 M2:

**Note:** Not all services listed are Business Partner enabled. Services vary by country.

► **Operational Support Services : Support Line**

Support Line gives you access to IBM technical support specialists by telephone or electronically. The offering covers most IBM operating systems and the most popular multivendor platforms, including Microsoft Windows, Linux®, and ESX Server. Additionally, you can get usage and configuration assistance for selected hardware products. Customized options make it easy to choose what you want covered and the level of support required.

The technical support specialists can address a wide variety of questions and concerns, including:

- Usage and installation questions
- Product compatibility and interoperability questions
- Interpretation of product documentation
- Diagnostic information review to help isolate the cause of a problem
- Configuration samples
- IBM and multivendor database searches
- Planning information for software fixes
- Assistance with problem source identification and problem resolution
- Defect support

For the list of products that are supported, see:

<http://www.ibm.com/services/s1>

► **High Availability Services**

These services identify availability issues, inhibitors, and exposures, and then recommend improvements and tailored solutions. For information, see:

<http://www.ibm.com/services/us/index.wss/offerfamily/bcrs/a1026936>

- ▶ Enterprise Application Integration Services for Microsoft Technologies

Skilled Microsoft Certified Professionals can assist you with migration, assessment, planning, architecture, design, and deployment services for the implementation of Datacenter and other Microsoft products. Testing Services include proof of concept, assessment, design, planning, and automated testing to help determine IT problem areas before system applications are launched. For more information, go to:

<http://www.ibm.com/services/us/index.wss/so/its/a1002580>

- ▶ IBM Business Continuity And Consulting Services

IBM Business Continuity and Consulting Services consultants work with you to determine risks, vulnerabilities, and financial impacts that a disruption can have on your business and help you plan for emergency situations. They can also operate and manage your continuity program for you, so you can focus on your core business. For complete information about the full range of services, visit:

<http://www.ibm.com/services/us/index.wss/itservice/bcrs/a1000411>

For information about other available IBM services, see:

<http://www.ibm.com/services/>



# Abbreviations and acronyms

<b>AMD</b>	Advanced Micro Devices™	<b>ITSO</b>	International Technical Support Organization
<b>AS</b>	Australian Standards	<b>IWS</b>	International Warranty Service
<b>ASE</b>	Adaptive Server Enterprise	<b>LAN</b>	local area network
<b>AWE</b>	Address Windowing Extensions	<b>LED</b>	light emitting diode
<b>BIOS</b>	basic input output system	<b>LUN</b>	logical unit number
<b>BMC</b>	baseboard management controller	<b>MAC</b>	media access control
<b>CD-ROM</b>	compact disc read only memory	<b>MB</b>	megabyte
<b>COM</b>	Component Object Model	<b>MIOC</b>	Memory and I/O Controller
<b>CPU</b>	central processing unit	<b>MOF</b>	Managed Object Format
<b>CTP</b>	composite theoretical performance	<b>MP</b>	multiprocessor
<b>DAC</b>	dual address cycle	<b>NAS</b>	network attached storage
<b>DAS</b>	Direct Attached Storage	<b>NUMA</b>	Non-Uniform Memory Access
<b>DB</b>	database	<b>OLAP</b>	online analytical processing
<b>DBA</b>	database administrator	<b>OLTP</b>	online transaction processing
<b>DBCC</b>	database consistency checker	<b>OS</b>	operating system
<b>DIMM</b>	dual inline memory module	<b>PAE</b>	Physical Address Extension
<b>DMA</b>	direct memory access	<b>PCI</b>	Peripheral Component Interconnect
<b>DMF</b>	Declarative Management Framework	<b>PFA</b>	Predictive Failure Analysis
<b>DP</b>	dual processor	<b>POST</b>	power on self test
<b>DRAM</b>	dynamic random access memory	<b>RAID</b>	redundant array of independent disks
<b>DTS</b>	Data Transformation Services	<b>RAM</b>	random access memory
<b>ECC</b>	error checking and correcting	<b>RDBMS</b>	relational database management system
<b>EIDE</b>	enhanced IDE	<b>ROI</b>	return on investment
<b>EPIC</b>	Explicitly Parallel Instruction Computing	<b>RSA</b>	Remote Supervisor Adapter
<b>ERP</b>	enterprise resource planning	<b>RTM</b>	release to manufacturing
<b>GB</b>	gigabyte	<b>SAN</b>	storage area network
<b>GPR</b>	general purpose register	<b>SAS</b>	Serial Attached SCSI
<b>GUI</b>	graphical user interface	<b>SCSI</b>	Small Computer System Interface
<b>I/O</b>	input/output	<b>SLA</b>	service level agreement
<b>IBM</b>	International Business Machines	<b>SMP</b>	symmetric multiprocessing
<b>ID</b>	identifier	<b>SQL</b>	Structured Query Language
<b>IEEE</b>	Institute of Electrical and Electronics Engineers	<b>SSAS</b>	SQL Server Analysis Services
<b>IP</b>	Internet Protocol	<b>SSIS</b>	SQL Server Integration Services
<b>IPMI</b>	Intelligent Platform Management Interface	<b>SSMA</b>	SQL Server Migration Assistant
<b>IT</b>	information technology	<b>SSNS</b>	SQL Server Notification Services
<b>ITS</b>	IBM Integrated Technology Services	<b>SSRS</b>	SQL Server Reporting Services
		<b>TB</b>	terabyte
		<b>TCO</b>	Total Cost of Ownership

<b>TCP</b>	Transmission Control Protocol
<b>TCP/IP</b>	Transmission Control Protocol/Internet Protocol
<b>TPM</b>	Trusted Platform Module
<b>TSQL</b>	Transact-SQL
<b>UDP</b>	user datagram protocol
<b>URL</b>	Uniform Resource Locator
<b>USB</b>	universal serial bus
<b>VAS</b>	Virtual Address Space
<b>VT</b>	Virtualization Technology
<b>WAN</b>	wide area network
<b>WOW</b>	Windows on Windows
<b>WSRM</b>	Windows System Resource Manager
<b>WSU</b>	Warranty Service Upgrade

# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

## IBM Redbooks

You can search for, view, or download books, papers, Technotes, draft publications, and additional materials, as well as order hardcopy Redbooks, at the IBM Redbooks Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

Related publications from IBM Redbooks include the following:

- ▶ *Tuning IBM System x Servers for Performance*, SG24-5287
- ▶ *Implementing IBM Director 5.20*, SG24-6188
- ▶ *Introducing Windows Server x64 on IBM eServer xSeries Servers*, REDP-3982

## Other publications

These publications are also relevant as further information sources:

- ▶ Microsoft whitepaper: Guide to Creating and Configuring a Server Cluster under Windows Server 2003  
<http://www.microsoft.com/downloads/details.aspx?familyid=A5BBB021-0760-48F3-A53B-0351FC3337A1&displaylang=en>
- ▶ Microsoft whitepaper: Server Clusters: Cluster Configuration Best Practices for Windows Server 2003  
<http://www.microsoft.com/downloads/details.aspx?familyid=98BC4061-31A1-42FB-9730-4FAB59CF3BF5&displaylang=en>
- ▶ Microsoft whitepaper: SQL Server 2005 Failover Clustering  
<http://www.microsoft.com/downloads/details.aspx?FamilyID=818234DC-A17B-4F09-B282-C6830FEAD499&displaylang=en>
- ▶ Microsoft whitepaper: Microsoft SQL Server 2005 Upgrade Advisor  
<http://www.microsoft.com/downloads/details.aspx?familyid=1470e86b-7e05-4322-a677-95ab44f12d75&displaylang=en>
- ▶ Microsoft whitepaper: SQL Server 2005 Upgrade Technical Reference Guide  
<http://www.microsoft.com/downloads/details.aspx?FamilyID=3d5e96d9-0074-46c4-bd4f-c3eb2abf4b66&DisplayLang=en>
- ▶ Microsoft whitepaper: Best Practices Analyzer Tool for Microsoft SQL Server 2000 1.0  
<http://www.microsoft.com/downloads/details.aspx?FamilyID=b352eb1f-d3ca-44ee-893e-9e07339c1f22&displaylang=en>
- ▶ Microsoft whitepaper: SQL Server 2005 Best Practices Analyzer  
<http://www.microsoft.com/downloads/details.aspx?familyid=da0531e4-e94c-4991-82fa-f0e3fbd05e63&displaylang=en>

## Online resources

These Web sites are also relevant as further information sources.

### IBM pages

- ▶ IBM Sizing Guides  
<http://www.developer.ibm.com/servers/sizing/portal/search.jsp>
- ▶ IBM Sizing Guide for SQL Server  
[http://www.developer.ibm.com/servers/sizing/sizingguide/guide\\_redirect.jsp?guide\\_id=sgq80502835091106002](http://www.developer.ibm.com/servers/sizing/sizingguide/guide_redirect.jsp?guide_id=sgq80502835091106002)
- ▶ Integrated Technology Services: Business continuity and resiliency  
<http://www.ibm.com/services/us/index.wss/itservice/bcrs/a1000411>
- ▶ Support Line  
<http://www.ibm.com/services/sl>
- ▶ High availability services  
<http://www.ibm.com/services/us/index.wss/offerfamily/bcrs/a1026936>
- ▶ Implementation Services - enterprise integration services for Microsoft technologies  
<http://www.ibm.com/services/us/index.wss/so/its/a1002580>
- ▶ IBM Storage area network products  
<http://www.ibm.com/storage/snetwork>

### Microsoft pages

- ▶ Comparison Between SQL Server 2005 Standard and Enterprise Editions  
<http://www.microsoft.com/sql/editions/enterprise/comparison.msp>
- ▶ SQL Server 2005  
<http://www.microsoft.com/sql/2005>
- ▶ SQL Server 2008  
<http://www.microsoft.com/sql/2008>
- ▶ Migrate to SQL Server 2005  
<http://www.microsoft.com/sql/solutions/migration>
- ▶ Licensing issues regarding Virtualization and Multiple Instances  
<http://www.microsoft.com/sql/howtobuy/multipleinstances.msp>
- ▶ SQL Server 2005 Books Online  
<http://msdn2.microsoft.com/en-US/library/ms130214.aspx>
- ▶ How to: Configure the Affinity Mask (SQL Server Management Studio)  
<http://msdn2.microsoft.com/en-us/library/ms186255.aspx>
- ▶ Setting Server Configuration Options  
<http://msdn2.microsoft.com/en-us/library/ms189631.aspx>
- ▶ How to: Map TCP/IP Ports to NUMA Nodes  
<http://msdn2.microsoft.com/en-us/library/ms345346.aspx>

- ▶ SQL Server support policy for Microsoft Clustering  
<http://support.microsoft.com/kb/327518>
- ▶ Walkthrough: Clustered Single-Server Deployment  
<http://msdn2.microsoft.com/en-us/library/ms171407.aspx>
- ▶ Configuring Integration Services in a Clustered Environment  
<http://msdn2.microsoft.com/en-us/library/ms345193.aspx>
- ▶ How SQL Server 2005 Supports NUMA  
<http://msdn2.microsoft.com/en-us/library/ms180954.aspx>
- ▶ How to: Configure SQL Server to Use Soft-NUMA  
<http://msdn2.microsoft.com/en-us/library/ms345357.aspx>
- ▶ Windows System Resource Manager download  
<http://www.microsoft.com/windowsserver2003/downloads/wsrp.msp>
- ▶ Lock pages in memory  
<http://technet2.microsoft.com/windowsserver/en/library/e72dcdf6-fe17-49dd-a382-02baad31a1371033.mspx?mfr=true>
- ▶ How to: Enable the Lock Pages in Memory Option (Windows)  
<http://msdn2.microsoft.com/en-us/library/ms190730.aspx>
- ▶ A transaction log grows unexpectedly or becomes full on a computer that is running SQL Server  
<http://support.microsoft.com/kb/317375/en-us>

### Intel pages

- ▶ Intel 64 Architecture  
<http://www.intel.com/technology/architecture-silicon/intel64>
- ▶ Intel Server and Workstation Processors  
<http://www.intel.com/products/server/processors/index.htm?iid=process+server>

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)







# Consolidating Microsoft SQL Server on the IBM System x3950 M2



**Suggests strategies and techniques to consolidate SQL Server systems**

Database servers have become a vital part of the IT infrastructure for organizations of all sizes. This has been fueled by the requirement for those organizations to capture, analyze, and report on an ever-increasing volume of data of all types and sizes.

**Explains when consolidation is and isn't practical**

The proliferation of systems running Microsoft SQL Server in many organizations, while tactically solving many business problems, has brought with it a number of management issues and costs. This has been accentuated by the growth in importance of the regulatory and compliance controls that have imposed new demands on most businesses around the globe.

**Helps you prepare for a consolidation activity**

The concept of server consolidation is not new and the cost reduction benefits that a physical server consolidation project can bring to an organization are well understood. This paper explores how the server consolidation concept is best applied to a SQL Server environment by suggesting strategies and techniques that apply specifically to SQL Server. We also explain why the IBM System x3950 M2 server is an ideal platform for a SQL Server consolidation exercise.

The paper is written for the technical decision makers, infrastructure architects, solution designers, and technical managers and is positioned between a business overview and a technical user guide.

**INTERNATIONAL  
TECHNICAL  
SUPPORT  
ORGANIZATION**

**BUILDING TECHNICAL  
INFORMATION BASED ON  
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:  
[ibm.com/redbooks](http://ibm.com/redbooks)**