Dell Compellent Storage Center

Microsoft SQL Server Best Practices

Document Revision

Date	Revision	Comments
10/8/2007	1.0	Initial Release
12/3/2007	1.1	
6/16/2008	1.2	
7/21/2010	1.3	
6/25/2012	1.4	

Contents

Tables

Figures

General Syntax

Table 1. Document Syntax

Item	Convention
Menu items, dialog box titles, field names, keys	**Bold**
Mouse click required	Click:
User Input	Monospace Font
User typing required	Type:
Website addresses	http://www.compellent.com
Email addresses	info@compellent.com

Conventions

Notes are used to convey special information or instructions.

Timesavers are tips specifically designed to save time or reduce the number of steps.

Caution indicates the potential for risk including system or data damage.

Warning indicates that failure to follow directions could result in bodily harm.

Preface

Audience
The audience for this document is database administrators, system administrators, storage administrators, and architects who analyze, design, and maintain robust database and storage systems. Readers should be familiar with Microsoft SQL Server, Microsoft Windows Server, and Dell Compellent Storage Center.

Purpose
This document describes the best practices for running Microsoft SQL Server (2000, 2005, 2008, 2008 R2 and 2012) with a Dell Compellent Storage Center.

SQL Server performance tuning and profiling are beyond the scope of this paper. Due to dependencies on application design and requirements, environments and their recommendations can differ greatly.

Please visit Microsoft Technet for more in depth information on tuning SQL Server databases.
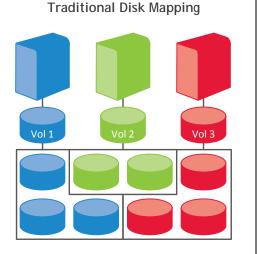
Customer Support
Dell Compellent provides live support 1-866-EZSTORE (866.397.8673), 24 hours a day, 7 days a week, 365 days a year. For additional support, email Dell Compellent at support@compellent.com. Dell Compellent responds to emails during normal business hours.

Understanding Storage Virtualization

With traditional SAN solutions, a LUN is created by carving out a specific amount of disk space from a specified set of spindles. This process is time-consuming and hard to manage because the mapping between servers and storage must be tracked using worldwide names (WWNs). Not only is this method time-consuming, it's also confusing since the storage configuration winds up being managed via a spreadsheet or other means beyond the storage management interface.

Dell Compellent virtualizes storage at the disk level, using all available spindles as a single virtual disk pool. Dell Compellent can create volumes where all drives are used together and performance is improved because the blocks of data are written in parallel to all drives at once.

Figure 1.   Traditional SAN Disk Mapping vs. Dell Compellent

SQL Server on Dell Compellent Storage

A Dell Compellent Storage Center provides a flexible, easy to manage storage platform that helps meet the demanding uptime and performance requirements common in SQL Server environments.

Manageability

SQL Server environments are often very dynamic. A Dell Compellent Storage Center makes it easy to adapt the storage configuration to meet changing database requirements. Volumes can be created or expanded with very little effort. Data can be easily moved between tiers or RAID levels. Disks can be added to increase performance or capacity on the fly. New storage technology can be integrated when it becomes available. All of this can be done while SQL Server remains online and available for applications.

High Availability

The uptime of storage systems is more important than ever, as organizations increasingly grow dependent on SQL Server databases. The Dell Compellent Storage Center was designed with high availability and fault tolerance in mind. Redundant hardware and RAID technology protect against an unplanned outage due to the failure of a component. Maintenance tasks such as controller firmware upgrades, hardware expansions or hardware replacements can all be done "hot" without taking

databases offline. The result is a storage system that rarely needs to be taken down, providing a solid foundation for SQL Server databases.

Performance

SQL Server is an I/O intensive application that requires high performance storage. A Dell Compellent Storage Center makes it easy to leverage a large number of spindles to support SQL Server workloads. Solid state disks (SSD) can also be integrated into a Storage Center for those environments that require extremely low latency. A variety of connectivity options exist, including 8Gbps fibre channel and 10Gbps iSCSI, to ensure that data can be moved quickly and efficiently between the database server and the Storage Center. Increasing performance by adding additional spindles or new storage technology can be done dynamically, ensuring good performance in the future.

Sizing for SQL Server

The I/O subsystem is a critical component of any SQL Server environment. Sizing and configuring a storage system without understanding the I/O requirements can have disastrous consequences. Monitoring performance in an existing environment using a tool like the Dell Performance Analysis Collection Kit (DPACK) can help define the I/O requirements. For best results, capture performance statistics for a time period of at least 24 hours that includes the system's peak workload.

Key I/O Performance Metrics

A good understanding of the key metrics used to describe I/O performance is required to successfully define I/O requirements and size storage systems.

IOPS

I/O operations per second, or IOPS, describe the number of reads and writes occurring each second. This is the primary metric used to determine the number of disks required in a storage system. The number of IOPS is reported by Windows Performance Monitor using the "Disk Reads/sec", "Disk Writes/sec" and "Disk Transfers/sec" physical disk counters.

Throughput

Typically expressed in megabytes per second (MB/s), throughput describes the amount of data transferred between the server and the storage. Throughput can be calculated using the number of IOPS and the average I/O size (Throughput = IOPS * Average I/O size). This is the primary metric used to design the path between the server(s) and the storage. Throughput is reported by Windows Performance Monitor using the "Disk Read Bytes/sec", "Disk Write Bytes/sec" and "Disk Bytes/sec" physical disk counters.

Latency

Typically expressed in milliseconds (ms), latency describes the amount of time an I/O operation takes to complete. When a server experiences high latency, this generally indicates an I/O bottleneck in the system. Latency is reported by Windows Performance Monitor using the "Avg. Disk sec/Read", "Avg. Disk sec/Write" and "Avg. Disk sec/Transfer" physical disk counters.

The Write Penalty

RAID technology is used to protect data on a Dell Compellent Storage Center. For each write sent to a Storage Center volume, additional I/O is performed on the physical disks to provide the redundancy

needed to protect the incoming data. The amount of additional I/O is determined by the RAID configuration for the volume. The following table shows the number of I/Os required when performing a write on a Dell Compellent Storage Center. The number of I/Os required to complete a write is referred to as the "Write Penalty".

Table 2. Write Penalty by RAID Type

RAID type	Write penalty	I/O Description
RAID 10	2	2 writes
RAID 10 DM	3	3 writes
RAID 5	4	2 reads, 2 writes
RAID 6	6	3 reads, 3 writes

By default, writes to Storage Center volumes will use RAID 10. Since RAID 10 has the lowest write penalty, it is strongly recommended to always use RAID 10 for writes. All storage profiles shipped with the Storage Center will write to RAID 10.

When sizing a Storage Center, it is important to factor in the write penalty when determining the number of disks required. The number of IOPS reported by the server when gathering I/O statistics does not include the write penalty. If the IOPS target for the Storage Center design does not include the write penalty, it will be undersized. To adjust the number of IOPS to include the write penalty, use the following formula:

*Number of IOPS + (Number of IOPS * Write Percentage * (Write Penalty – 1))*

Where:
Number of IOPS   = Number of IOPS generated by the server
Write Percentage = Percent of the I/O attributed to writes
Write Penalty   = Number of I/Os required to complete a write

As an example, consider a server that is performing a workload of 10,000 IOPS with a mix of 70% reads and 30% writes. Since the default storage profile will be used, writes will use RAID 10, which has a write penalty of 2. For this workload, the Storage Center will need enough disks to support 13,000 IOPS (10,000 + (10,000 * .30 * (2-1)) = 13,000).

Random I/O vs. Sequential I/O

Spinning disks can perform more sequential I/O than random I/O. This is because sequential I/O requires less head movement on the disk. While I/O statistics show how much I/O is occurring, they do not show how much is sequential I/O versus random I/O. The virtualized storage architecture of a Dell Compellent Storage Center may cause sequential I/O requests from the server to become large block random I/O on the physical disks. Performance will still be good, as large block random I/O provides comparable performance to sequential I/O. For best results when sizing a Dell Compellent Storage

Center, assume all of the I/O is random. This will ensure good performance for both random and sequential operations.

OLTP Workloads

While every environment is unique, an OLTP (Online Transaction Processing) workload typically consists of small, random reads and writes. A storage system servicing this type of workload is primarily sized based on the number of IOPS required. Different drive types have different performance capabilities. Faster spinning drives are able to provide more IOPS with lower latency than slower spinning drives. For best performance, 15K drives are recommended for OLTP workloads. For environments with very low latency requirements, consider using SSDs (See the "Using SSDs" section below).

OLAP/DSS Workloads

An OLAP (Online Analytic Processing) or DSS (Decision Support System) workload is typically dominated by large, sequential reads. A storage system servicing this type of workload is primarily sized based on throughput. When designing for throughput, the performance of the entire path between the server and the Storage Center needs to be considered. For best throughput, consider using 8Gbps Fiber Channel (FC) or 10 Gbps iSCSI connectivity to the Storage Center (aka front-end connectivity), and 6 Gbps SAS connectivity from the controllers to the disk enclosures. To meet high throughput requirements, multiple HBAs may be required in the server, the Storage Center or both. While 15K drives offer the best performance, 10K drives may be sufficient to meet throughput requirements and offer more capacity per spindle.

Sizing Tiers

For Storage Centers with two tiers, the rule of thumb is to size Tier 1 for all of the IOPS and at least 20% of the capacity, and Tier 3 for 80% of the capacity, and at least 20% of the IOPS. (In a two tier system, the fastest disks will be Tier 1 and the slower disks Tier 3.) When sizing for SQL Server, there are a couple of things to consider when sizing each tier.

New pages on a volume are not eligible to move to a lower tier until they are at least 12 days old. In environments where indexes are rebuilt frequently, active pages on the database volume may never be older than 12 days. All active pages will always be in Tier 1. In those environments, Tier 1 will need to be sized for 100% of the capacity required by SQL Server.

The automated tiering feature of a Dell Compellent Storage Center automatically moves infrequently accessed data to slower, cheaper storage (Tier 3), reducing storage costs without impacting performance. For environments where the access pattern is even across the entire data set, all data may need to reside on Tier 1 to ensure good performance. In those environments, Tier 1 will need to be sized for 100% of the capacity required by SQL Server.

Using SSDs

Solid state disks (SSD) offer significantly better performance than traditional hard drives for most I/O patterns. From a performance perspective, SSDs are suitable for storing any type of SQL Server file. While SSDs can support a large number of IOPS, it is easy to exceed throughput limits of the drive or other components in the path with the large I/Os that can be generated by SQL Server. It is important to understand the performance characteristics of SSDs as well as the I/O requirements of the SQL Server environment before implementing SSDs.

When using SSDs as a tier that participates in automated tiering, there needs to be enough SSDs in the tier to handle 100% of the IOPS and at least 20% of the capacity of the volumes that will use the SSD

tier. Since SSDs can perform a large number of IOPS, a small number of drives may meet the IOPS requirement. However, the capacity requirement may require many drives. It's very important to have enough SSDs to meet both the IOPS and capacity requirements. If the SSD tier becomes full, performance will degrade significantly.

Given the relatively high cost of SSDs, it may be more cost effective to place a subset of the SQL Server files on SSDs, instead of using SSDs with automated tiering. Storage profiles allow an administrator to easily control which volumes reside on SSD. While transaction log files and tempdb files require high performance and tend to be small enough to fit on a small number of SSDs, they are not always the best choice for SSDs. In some cases, it may be better to place data files on SSD. Analyze the access patterns of the database files to determine the best fit. If a data file is too large to fit on the available SSD drives, the portion of the data requiring high performance can be moved to a separate file group and placed on SSD.

Validating the Storage Design

Once the I/O requirements have been defined, it is easy to get a feel for whether the hardware can provide the desired performance by running some simple tests. Both SQLIO and IOMeter can simulate I/O patterns generated by SQL Server. SQLIO is available from Microsoft as a free download. IOMeter is available via Open Source and can be found by using your favorite Internet search engine.

Validating the I/O Path

The first thing to test on a new configuration is the path between the server and the Storage Center. Running a large block sequential read test using small files will saturate the path between the server and storage. This test will verify that all HBA ports in the server are active and can be used for I/O traffic.

> ⚠️ **Warning**
>
> This test should be run on a server and Storage Center that is dedicated to this test. Running this test on a live system will cause significant performance issues.

To validate the I/O path using SQLIO, perform the following steps:

1) Download the SQLIO utility and install it on the server.

2) On the Storage Center, create volumes and map them to the server. If testing a single controller system, create one volume. If testing a dual controller Storage Center, create two volumes, with a volume owned by each controller.

3) On the server, format each volume using a 64KB allocation unit.

4) Create a SQLIO parameter file called "paramPath.txt". This file will contain the number of files to use, along with the number of threads for each file and the name of the file. Use one 10MB file per volume. Use a total of 8 threads.

   Here's an example of the parameter file for a single controller test:

   ```
   L:\testfile.dat 8 0x0 10
   ```

Here's an example of the parameter file for a dual controller test:

```
L:\testfile.dat 4 0x0 10
M:\testfile.dat 4 0x0 10
```

Note: The drive letters should be adjusted to match the server being tested.

5) Run a sequential read test. Use a block size of 512KB with 32 outstanding IOs. This SQLIO command will run the test for 2 minutes:

```
sqlio –kR –s120 –fsequential –o32 –b512 –BH –LS –FparamPath.txt
```

If the throughput shown in the results matches the expected throughput for the number of HBA ports in the server, the path between the server and Storage Center is setup correctly.

Validating the Disks

Once the I/O path has been validated, it's time to test the disks. For best results when testing disks on a Storage Center, use the following guidelines when configuring the test:

1) In a dual controller system, create at least one volume per controller. This will ensure that I/O will be distributed across both controllers. Using both controllers will more closely simulate real world activity. For best results, create the same number of volumes on each controller.

2) When performing I/O tests on any storage platform, it's important to use files that are larger than the controller's cache. For more accurate results, use a file size that matches the amount of data being stored. In an environment where that's not practical due to a large data set, use a file size of at least 100GB.

3) Some I/O test tools, including SQLIO, will generate files full of zeros. By default, Storage Center will track pages full of zeroes in the metadata, but not store them on disk. This behavior, known as thin writing, will cause inaccurate results when testing with files containing only zeros. Pre-allocate the volumes being used for I/O testing to prevent thin writing from skewing the test results.

Thin writing is beneficial during normal operations. Discard the test volumes and create new volumes for SQL Server once testing is complete.

The purpose of this type of testing is to validate that the storage design will provide the required throughput and IOPS with acceptable latency. It's important that the test does not exceed the designed capacity of the Storage Center. For example, a Storage Center designed for a workload of 5,000 IOPS is likely to perform poorly with a workload of 10,000 IOPS. If a test is generating a workload higher than the designed capacity, adjust the workload being generated by reducing the number of threads and/or outstanding IOs.

The results of the DPACK analysis will provide an I/O target to simulate using these tests. To get an idea of the performance capabilities of the Storage Center, run I/O tests with a range of I/O sizes commonly seen with SQL Server. When testing random I/O, test with an I/O size of 8KB and 64KB. When testing sequential I/O, start with an I/O size of 8KB and 64KB. Since processes like read ahead scans and backups can issue much larger sequential I/Os, it's a good idea to also test block sizes up to 1024KB.

Storage Setup & Configuration

Single vs. Multiple Disk Pools

Dell Compellent recommends using a single virtual disk pool when implementing SQL Server. This provides better performance by leveraging the aggregate I/O bandwidth of all disks to service I/O requests from SQL Server. A single disk pool is also easier to manage, allowing an administrator to quickly and easily adapt the storage system to satisfy the ever changing workloads that are common in SQL Server environments.

Creating Volumes

A Dell Compellent Storage Center virtualizes its storage to take advantage of all spindles in the disk pool. Characteristics such as RAID level and storage tier can have a big impact on performance and are configured at the volume level. Data Instant Replay™ is also configured at the volume level. There are many different types of files that are part of a SQL Server instance. Those different types of data have different performance and Replay requirements. Dell Compellent recommends creating at least 5 volumes for an instance of SQL Server as shown in the following table:

Table 3. Volume Provisioning Recommendations

File type	Number of volumes	Volume performance requirements	Volume Replay requirements
User DB data	At least 1 volume per instance	Lower performance may be acceptable	Frequent Replays — same schedule as transaction log
User DB transaction log	At least 1 volume per instance	High performance required	Frequent Replays — same schedule as data files
Data root directory (includes system DBs)	1 volume per instance	Lower performance may be acceptable	Infrequent Replays — independent schedule
Tempdb data and transaction log	1 volume per instance	High performance required	No Replays
Native SQL Server backup	1 volume per instance	Lower performance may be acceptable	Replays optional — independent schedule

Make each volume a unique size. If a server needs three 100GB volumes, create a 100GB, 101GB and 102GB volume instead. This will make it much easier to correlate the volumes between the Storage Center and the database server.

Replay Considerations

Even though Replay Manager™ allows specific databases to be selected, Replays are taken on the entire volume. If all databases on a given set of volumes are not backed up together, disk space will be wasted on the Storage Center. Consider placing the databases for each backup set on a separate set of volumes.

Placing a large database on its own set of volumes will allow additional recovery flexibility when using Replay Manager™. If a database is on its own volumes and in its own backup set, the database can be recovered very quickly using the Fast Recovery feature without affecting other databases. This can be a big time saver when recovering a large database with Replays and transaction log backups.

Performance Considerations
If there is a group of databases that require high performance and another group that does not, consider creating a set of volumes for each group of databases. Even if there is only one tier of storage, this strategy will make it easier to adjust the storage configuration in the future. For example, having the high performance databases on separate volumes will make them very easy to move them to a higher performance tier in the future.

Having multiple volumes has an additional benefit in dual controller systems. All I/O requests for a given volume are processed by the controller that owns the volume. While a volume can be owned by either controller, a volume is only owned by one controller at a time. Having many volumes makes it easier to distribute the I/O load evenly across both controllers.

Flexibility vs. Manageability
For ultimate flexibility, create a volume for each user database file. This provides the ability to independently optimize the storage and Replay configuration for each individual database. With thin provisioning, there is no space penalty for creating a lot of volumes. However, a large number of volumes can be difficult to manage. It is up to the DBA and/or storage administrator to find the right balance between flexibility and maintainability when determining the number of volumes to create.

Storage Profiles
Storage profiles are used to define the RAID level used to protect data on a volume and the tiers on which that data is stored. The information in the profile is used by Data Progression™ when moving existing pages as well as by the Storage Center when creating new pages. In most environments, using the "Recommended" storage profile (which is the default) will provide good I/O performance for all types of database volumes. It is strongly recommended to use the "Recommended" profile first and evaluate the suitability of the profile before attempting to change storage profiles.

User Database Data File Considerations
On volumes storing the data files, the "Recommended" profile will keep the highly active parts of the database on Tier 1 and the less active parts on lower tiers. In most cases, this will provide the best performance. However, there are some environments where the entire dataset is highly active and needs the performance of Tier 1. Before changing the storage profile to force the data volume to live only on Tier 1, consider the frequency of index rebuilds. A new page has to exist for 12 days before it is eligible to be moved to a lower tier. If indexes are rebuilt more often than every 12 days, most of the database, if not all, will be on Tier 1. Using the "Recommended" profile allows inaccessible Replay pages to be moved to Tier 3, freeing up Tier 1 resources for active pages. Moving inaccessible Replay pages to Tier 3 provides better performance on Tier 1.

User Database Transaction Log File Considerations
The transaction log requires good performance and should reside on Tier 1. Writes to the transaction log start at the beginning of the file, go to the end of the file and then start again at the beginning of the file. If the entire transaction log is written within 12 days, all of the active pages will be in Tier 1 using the "Recommended" profile. The "Recommended" profile will allow inaccessible Replay pages to be move to Tier 3, freeing up Tier 1 resources for active pages.

Tempdb Database File Considerations

Tempdb files require high performance and should reside on Tier 1. Since Replays should not be enabled on tempdb volumes, all pages will stay on Tier 1 using the "Recommended" profile in a two Tier system. In a three tier system, use a storage profile to ensure the tempdb volume stays on the desired Tier. In most cases, this will be Tier 1. If Tier 1 is an SSD tier, some environments are better off with tempdb on the tier containing 15K drives instead of SSD drives.

System Database File Considerations

The system databases (master, model, msdb) generally do not have special performance requirements. Use the "Recommended" profile for the system database volumes.

SSD Considerations

There are usually not enough SSDs to function as Tier 1 storage for all volumes in the Storage Center. Custom storage profiles can be used to control which volumes will use the SSD Tier and which volumes will be restricted to spinning drives. The type of files that benefit the most from SSDs will vary from environment to environment. Monitor the performance of the database volumes to determine the best fit for SSD.

> Manual storage mode must be enabled to use custom storage profiles. See the Storage Center System Manager Admin Guide for more information.

Data Instant Replay

Data Instant Replay™ is very powerful when combined with SQL Server. By leveraging Replays, database backup and recovery is very fast and space efficient, even for large databases.

Dell Compellent strongly recommends using Replay Manager™ to take Replays of a database volume while SQL Server is online. Using Replay Manager™ is the only way to take a Replay while I/O is frozen inside of SQL Server. Although it is rare, Replays of active database volumes taken without Replay Manager™ may not be usable. Replay Manager™ also allows a Replay to be used as the starting point of a restore chain, as databases can be restored without recovery.

When taking Replays of database volumes without Replay Manager™, it is recommended to use a consistency group. This will increase the odds of a successful database recovery using those Replays. Using a consistency group on the Storage Center allows Replays to be taken on multiple volumes at the same point in time. For information on how to setup a consistency group, see the "Creating a Consistent Replay Profile" section in the Storage Center System Manager Admin Guide.

Dell Compellent recommends taking a Replay of each user database volume at least once per day, with a retention period of 25 hours or longer. Data Progression™ will move Replay data (i.e: read-only data) according to settings in the storage profile. If using the "Recommended" storage profile (which is the default), Replay data will be moved to RAID5, while writes will still occur on RAID10. For more information on how Data Progression™ works with Replays, see the document Data Progression™ Configuration for Replay Enabled Volumes.

> It is not necessary to take Replays on the Storage Center if Replays are being taken with Replay Manager™.

Tasks ideally suited for using Replays with SQL Server include:

- Restoring a copy of a database to quickly recover data or objects after errant updates or deletes.

- Restoring a copy of a database on another server for development, testing or reporting.

- Replicating databases to a DR location. During a DR test, databases can be recovered without interrupting replication.

For more information about using Replays with SQL Server, see the Replay Manager™ Users Guide.

Windows Setup & Configuration

Allocation Unit Size

Use a 64KB allocation unit size when formatting volumes that will contain database files (log and data) or database backups.

MPIO

Set the MPIO policy to "round robin" for all database volumes. This is the default for Windows 2008 and newer. This will allow all paths to be used, allowing higher throughput between the server and the Storage Center.

Partition Alignment

An offset of 64KB, or a multiple of 64KB, is recommended for volumes created on Windows 2003 and earlier. New volumes created on Windows Server 2008 or newer should already be aligned, as 1024KB is the default offset.

Server Setup & Configuration

HBA Considerations

Be sure that the HBA firmware and drivers are up-to-date. As with any update, it is important to verify functionality in a test or QA environment before implementing in production.

SQL Server is a very I/O intensive application. It is important to configure the queue depth of the HBA properly to optimize performance. See the Storage Center Release Notes for HBA configuration settings and best practices.

iSCSI Considerations

Use at least 1-Gbit Ethernet to connect to the Storage Center. For best performance, use 10-Gbit Ethernet.

To provide consistent performance, create a separate network for iSCSI traffic.

If using the Microsoft iSCSI Initiator, Dell Compellent recommends the following:

- Make the SQL Server service dependent on the Microsoft iSCSI Initiator Service.

- Use dedicated network interface cards (NIC) for iSCSI connections to the Storage Center.

- On each of the dedicated NICs, disable all features, protocols and clients except for TCP/IPv4.

- Explicitly create each of the iSCSI connections, rather than using the "Quick Connect" button.

Additional Resources

Dell Compellent Knowledge Center

Use the whitepapers, tips, and best practice guides in the Dell Compellent Knowledge Center to help get the most out of your Dell Compellent Storage Center. The Knowledge Center can be accessed at http://kc.compellent.com or through the Dell Compellent customer portal at http://customer.compellent.com.

Online Resources

There is a tremendous amount of SQL Server information available online. While not a complete list, the table below contains several useful links. Additional information can be found by using your favorite Internet search engine.

Table 4. Useful Links

Description	Link
SQL Server I/O Basics	http://technet.microsoft.com/en-us/library/cc966500.aspx
SQL Server I/O Basics, Chapter 2	http://technet.microsoft.com/en-us/library/cc917726.aspx
Predeployment I/O Best Practices	http://technet.microsoft.com/en-us/library/cc966412.aspx
Analyzing I/O Characteristics and Sizing Storage Systems for SQL Server Database Applications	http://msdn.microsoft.com/en-us/library/ee410782(v=SQL.100).aspx
Microsoft SQL Server homepage	http://www.microsoft.com/sql
SQL Server Customer Advisory Team	http://sqlcat.com
Microsoft online SQL Server forum	http://social.msdn.microsoft.com/Forums/en-US/category/sqlserver
Professional Association for SQL Server	http://sqlpass.org
Online SQL Server resources	http://www.sqlservercentral.com http://www.sqlserverpedia.com http://www.sqlteam.com http://www.sql-server-performance.com