

Chapter 1

Introducing SQL Server 2005

In This Chapter

- ▶ Figuring out which SQL Server Edition is best for you
 - ▶ Discovering what SQL Server 2005 is all about
-

In this chapter, I introduce you to SQL Server 2005. SQL Server 2005 is a multi-component relational database management system centered around a high-performance, highly available database engine.

The quality of the database engine in SQL Server is crucial to the reliability of SQL Server 2005 in handling large quantities of data. However, SQL Server 2005 is much more than a database engine and consists of a suite of tools and components that support you in designing, managing, maintaining, and programming a SQL Server 2005 installation and its associated data. In addition, there are powerful new or improved tools for business intelligence.

I introduce you to many of the important features and tools that you can find in SQL Server 2005. SQL Server 2005 is such an extensive suite of programs that I can only touch briefly on each of these many features. I show you how to put many of these features and tools to work in later chapters.

Getting to Know SQL Server 2005

You can use SQL Server 2005 to store information for personal use, for departmental use, for mid-size company use, or for enterprise use. SQL Server 2005 has editions (Microsoft provides a full comparison of the editions online at www.microsoft.com/sql/2005/productinfo/sql2005features.mspx) to meet the needs in each of those scenarios:

- ✓ **Enterprise:** Provides a relational database to meet the exacting needs of the largest enterprises and busiest online databases. The Enterprise Edition includes high-end business intelligence support and clustering. I introduce you to business intelligence in Chapters 20 through 22. This book does not cover clustering.

- ✓ **Standard:** Meets the needs of medium-sized companies or large departments in larger companies.
- ✓ **Workgroup:** Meets the needs of small- to medium-sized businesses that don't require the features of Standard Edition.
- ✓ **Mobile:** Formerly called SQL Server CE. This book does not cover Mobile Edition.
- ✓ **Express:** A lightweight edition intended for use by application programmers. To find out more about Express Edition, see *Microsoft SQL Server 2005 Express For Dummies*, by Robert D. Schneider (Wiley).

A Client-Server Database

SQL Server 2005 is a client-server database. Typically, the SQL Server 2005 database engine is installed on a server machine to which you connect anything from a few machines to many hundreds or thousands of client machines.

A client-server architecture can handle large amounts of data better than a desktop database such as Microsoft Access. The SQL Server instance provides security, availability, and reliability features that are absent from databases such as Access. A client-server architecture also can reduce network traffic.

The server side of a SQL Server installation is used for two broad categories of data processing: Online Transaction Processing (OLTP) and Online Analytical Processing (OLAP).

OLTP

Online Transaction Processing is the kind of processing that the databases of Amazon.com or any other large online retailer needs to do. A large number of orders come in every minute and the information from each of those orders needs to be written to the database quickly and reliably.

With OLTP, you can group certain actions together. For example, the different aspects of a bank transfer between accounts would be carried out together, so that if money is moved out of one account, it is also moved into another account. Actions such as these that must be done together are called a *transaction*. In the account transfer, either both the transfers take place or neither do. The all or nothing characteristic of a transaction ensures that the data remains in a consistent state. An OLTP database is tuned to support high

volumes of transactions that frequently change the data in the database. SQL Server 2005 performs well as an OLTP database management system.

The *transaction log* stores information about transactions and the data changes made in transactions, which are not rolled back. The transaction log is an important container for information about recent changes made to a database.

OLAP

An *Online Analytical Processing* database is intended to process large amounts of data that doesn't change often. For example, an online retailer might want to store summary data about sales by month, by region, by product category, and so on. In SQL Server 2005, the OLAP functionality is carried out in Analysis Services. In Analysis Services, you create *cubes* that allow you to examine *dimensions* of a cube. I describe Analysis Services in more detail in Chapter 21.

OLAP often takes place in a *data warehouse*. Getting large amounts of data into good shape before putting it into a data warehouse is a major task, an important aspect of which is to maximize data quality.

After all the data is aggregated, it is unlikely to change in the future. However, you can query it in complex ways, so an OLAP database is typically optimized to support fast querying.

A Secure Database

If the data on which your business depends is stored in SQL Server, you need to keep the wrong people from accessing the data or, worse, changing or deleting the data. Imagine if a hacker could change the price for certain goods and then buy a huge quantity for a nominal amount — your business could soon be a former business. Similarly, you don't want your competitors to be able to access information about the performance of your business.

SQL Server 2005 implements Microsoft's recent emphasis on security. Unlike its predecessor (SQL Server 2000), SQL Server 2005 is much more secure by default. Many potential attack points are turned off until you explicitly turn them on, so reducing the exposed risk of a default installation. The Surface Area Configuration tool is one way to configure this.

SQL Server 2005 builds on the authentication and authorization features present in SQL Server 2000. If you want to allow users to use SQL Server 2005 in particular ways, you — as the administrator for SQL Server — can allow them access.

Table 1-1 summarizes some key security features in SQL Server 2005.

| Table 1-1 Key Security Features of SQL Server 2005 | | | | |
|--|-------------------------------|---------------------------------|--------------------------------|----------------------------------|
| <i>Feature</i> | <i>Express Edition</i> | <i>Workgroup Edition</i> | <i>Standard Edition</i> | <i>Enterprise Edition</i> |
| Authentication and authorization | Yes | Yes | Yes | Yes |
| Data encryption and key management | Yes | Yes | Yes | Yes |
| Best Practices Analyzer | Yes | Yes | Yes | Yes |
| Integration with Microsoft Baseline Security Analyzer | Yes | Yes | Yes | Yes |
| Integration with Microsoft Update | Yes | Yes | Yes | Yes |

As you can see from the table, all editions of SQL Server 2005 support an extensive range of important security features. I discuss security issues in Part IV.

A Programmable Database

SQL Server 2005 is a great environment for programmers. If you're a developer, you might even find that your DBA (database administrator) is afraid of the programmability of SQL Server 2005 because it offers so many options. The good thing is that the increased programmability is accompanied by a detailed system of security permissions that greatly reduces the chances of rogue code doing harm to SQL Server itself.

Table 1-2 summarizes some key programmability support features in SQL Server 2005.

| Table 1-2 Key Programmability Support Features | | | | |
|---|------------------------|--------------------------|-------------------------|---------------------------|
| <i>Feature</i> | <i>Express Edition</i> | <i>Workgroup Edition</i> | <i>Standard Edition</i> | <i>Enterprise Edition</i> |
| Stored procedures and triggers | Yes | Yes | Yes | Yes |
| New enhancements to T-SQL | Yes | Yes | Yes | Yes |
| Integration of the Common Language Runtime and .NET support | Yes | Yes | Yes | Yes |
| User-defined types | Yes | Yes | Yes | Yes |
| XML datatype | Yes | Yes | Yes | Yes |
| XQuery support | Yes | Yes | Yes | Yes |

I introduce stored procedures in Chapter 9 and triggers in Chapter 14. I introduce the Common Language Runtime in Chapter 8, and I describe use of the XML datatype in Chapter 7.

Transact-SQL

SQL Server 2005 supports the *Structured Query Language* (SQL). In fact SQL is the main language that SQL Server uses. Like many other databases, SQL Server satisfies some standard SQL syntax and adds its own extensions to SQL, allowing you to easily write code.

The Microsoft flavor of SQL is called *Transact-SQL* (T-SQL). T-SQL allows you to add, modify, or query relational or XML data held in SQL Server 2005. In Chapter 5, I introduce you to using T-SQL to retrieve desired data. I introduce the manipulation of XML in Chapter 7.

SQL Server Management Studio

Transact-SQL programming is commonly done in the query pane of SQL Server Management Studio. I cover SQL Server Management Studio in Chapter 3.

You can also use the SQLCMD utility to issue T-SQL commands interactively or to run T-SQL script files. In applications you create, you can use T-SQL to retrieve or manipulate data. I don't describe these uses of T-SQL in this book.

Business Intelligence Development Studio

You can use Business Intelligence Development Studio (BIDS), which is based on Visual Studio components, to create business intelligence applications. BIDS is used to create SQL Server Integration Services, Analysis Services, and Reporting Services projects.

I cover business intelligence (BI) in Part VI. In Chapter 20, I show you how to use BIDS to create Integration Services projects. In Chapter 21, I introduce Analysis Services. And in Chapter 22, I introduce you to using BIDS with Reporting Services.

A Scalable Database

A database management system, such as SQL Server 2005, must grow as your business grows. Table 1-3 summarizes some key features of the various editions of SQL Server 2005 and lists the limitations of each. I do not cover in detail in this book how to make decisions on the most appropriate scalability decisions for your business.

| Table 1-3 SQL Server 2005 Editions and Their Limitations | | | | |
|---|------------------------|--------------------------|-------------------------|---------------------------|
| <i>Feature</i> | <i>Express Edition</i> | <i>Workgroup Edition</i> | <i>Standard Edition</i> | <i>Enterprise Edition</i> |
| CPU | 1 | 2 | 4 | No limit |
| RAM | 1GB | 3GB | No limit | No limit |
| Database size | 4GB | No limit | No limit | No limit |
| Partitioning | No | No | No | Yes |



Commercial decisions regarding which functionality is in which edition are subject to change up to the time of product release. Therefore, functionality may differ slightly from the information I give in this chapter.

An Available Database

If your business uses the World Wide Web to sell goods or provide information to customers, your customers may be located around the world. In that case, your SQL Server 2005-based applications need to be available 24 hours a day, 7 days a week. This need for continuous availability has stimulated many availability features in SQL Server 2005. Table 1-4 summarizes availability features in different editions of SQL Server 2005.

| Table 1-4 The Availability Features of SQL Server 2005 Editions | | | | |
|--|-------------------------------|---------------------------------|--------------------------------|----------------------------------|
| <i>Feature</i> | <i>Express Edition</i> | <i>Workgroup Edition</i> | <i>Standard Edition</i> | <i>Enterprise Edition</i> |
| Database mirroring | No | No | Partial | Yes |
| Failover clustering | No | No | Limited to 2 nodes | Yes |
| Backup log shipping | No | Yes | Yes | Yes |
| Online system changes | Yes | Yes | Yes | Yes |
| Online indexing | No | No | No | Yes |
| Online page and file restore | No | No | No | Yes |

If you cluster multiple server machines, they can operate together to give availability that would either be unavailable or very expensive with a single server. Each machine in a cluster is called a *node*. If a node fails, then another node in the cluster picks up the work of the failed machine. This significantly reduces downtime, at a cost of increased hardware.

Miscellaneous changes

You can add memory to the SQL Server machine if you have the relevant hardware. This allows you to add memory to a machine while it is available to users, so reducing downtime.

The Dedicated Administrator Connection (DAC) allows a database administrator to take control of the server even if an operation is using virtually all the CPU cycles. If a server process fails to complete, it can use virtually 100

percent of CPU time, making a server almost unusable or very slow. The DAC allows a server administrator to connect to the server and stop the runaway process without having to restart the server.

Online indexing

Online indexing is available only in the Enterprise Edition. Online indexing improves the availability of a database by avoiding the need to take a table or database offline while, for example, an index is rebuilt. I do not cover online indexing further in this book.

Online page and file restore

This feature is available only in Enterprise Edition. If you have to restore data from backups, the database becomes available for use more quickly than was possible in SQL Server 2000. I do not cover this feature further in this book.

A Reliable Database

Disasters, small and large, can happen. If the hard drive crashes on your SQL Server machine, your business doesn't need to crash with it.

Backing up data

You can back up your data with SQL Server Management Studio (see Chapter 3). You need to back up all your data regularly, especially the following system databases, with the exception of `tempdb`:

- ✓ **Master:** The `master` database contains system level information for a SQL Server 2005 system.
- ✓ **Model:** The `model` database is the template that is used when you create a new database.
- ✓ **Msdb:** The `msdb` database is used by SQL Server Agent to record information for scheduling alerts and jobs.
- ✓ **Resource:** The `resource` database is new in SQL Server 2005 and contains the system objects for SQL Server 2005. When you update SQL Server 2005 — for example a service pack — the new version replaces

the `resource` database. Be careful not to restore an out-of-date version after applying a service pack.

- ✓ **Distribution:** The `distribution` database exists only if the SQL Server machine is a distributor for replication. The database contains metadata about replication.
- ✓ **Temp:** The `temp` database is deleted when you close down SQL Server 2005. Although `tempdb` is a system database, you cannot back it up.

Replication

SQL Server 2005 supports replication of data by using a publish/subscribe metaphor. I describe replication in more detail in Chapter 19.

A Manageable Database

You need to manage many aspects of a SQL Server 2005 installation. Table 1-5 summarizes some features of manageability in SQL Server 2005. I cover SQL Server Management Studio in Chapter 3. I don't describe the Database Engine Tuning Advisor or Full Text Search in detail in this book.

| Table 1-5 SQL Server 2005 Manageability Features | | | | |
|--|------------------------|--------------------------|-------------------------|---------------------------|
| <i>Feature</i> | <i>Express Edition</i> | <i>Workgroup Edition</i> | <i>Standard Edition</i> | <i>Enterprise Edition</i> |
| Automatic performance tuning | Yes | Yes | Yes | Yes |
| SQL Server Management Studio | No | Yes | Yes | Yes |
| Database Engine Tuning Advisor | No | Yes | Yes | Yes |
| Full text search | No | Yes | Yes | Yes |
| SQL Agent job scheduling | No | Yes | Yes | Yes |

Command-line tools

In SQL Server 2005, the main command-line tool is SQLCMD. The SQLCMD utility allows you to manage a SQL Server 2005 installation by using the T-SQL language. You can use T-SQL interactively from the command line or can use SQLCMD to run T-SQL scripts. I do not cover the SQLCMD in detail in this book.

Graphical tools

The main graphical tool for administering SQL Server 2005 is SQL Server Management Studio. Management Studio allows you to administer many SQL Server 2005 database engine instances, Analysis Services instances, Integration Services instances, and Reporting Services instances from a single interface. I describe SQL Server Management Studio in more detail in Chapter 3.

SQL Server Agent

Any database administrator has tasks that need to be carried out repeatedly and, often, these tasks take place at set times. SQL Server Agent (or simply *SQL Agent*) is the software component that allows you to carry out such tasks automatically. For example, you may need to back up data at 2 a.m. every day. It is much more convenient for you to be at home asleep and let SQL Server Agent take the strain. See Chapter 16 for more about SQL Server Agent.

Performance tools

SQL Server Profiler allows you to monitor and analyze performance of a SQL Server instance. New in SQL Server 2005 is the ability to monitor and analyze the performance of Analysis Services. I do not cover SQL Server Profiler in detail in this book.

A Database That Supports Business Intelligence

SQL Server 2005 supports many pieces of business intelligence functionality, grouped under the headings of Integration Services, Analysis Services, and

Reporting Services. Table 1-6 summarizes the availability of business intelligence functionality by SQL Server 2005 edition.

| Table 1-6 | | Business Intelligence Functionality | | |
|--|--------------------------------|--|--------------------------------|----------------------------------|
| <i>Feature</i> | <i>Express Edition</i> | <i>Workgroup Edition</i> | <i>Standard Edition</i> | <i>Enterprise Edition</i> |
| Data can be used by Report Server | Yes | Yes | Yes | Yes |
| Report Builder | No | Yes | Yes | Yes |
| Scale out of Report Servers | No | No | No | Yes |
| Data warehousing | No | No | Yes | Yes |
| Business Intelligence Development Studio | Not included but is compatible | Not included but is compatible | Yes | Yes |
| Analysis Services | No | No | Yes | Yes |
| Advanced analytic functions | No | No | Yes | Yes |
| Data mining | No | No | Yes | Yes |
| Integration Services | No | No | No | Yes |

Chapter 20 describes SQL Server Integration Services. Chapter 21 has more on Analysis Services. And Chapter 22 includes information on Reporting Services.

