

Part I

Laying the Foundations

IN THIS PART

Chapter 1

The World of SQL Server

Chapter 2

Data Architecture

Chapter 3

Installing SQL Server

Chapter 4

Client Connectivity

Chapter 5

SQL Server Management and Development Tools

The World of SQL Server

IN THIS CHAPTER

Understanding SQL Server History and Overview

Understanding SQL Server Components and Tools

Understanding Notable Features in SQL 2012

What's New with SQL Server 2012?

SQL Server 2012 represents another tremendous accomplishment for the Microsoft data platform organization. A number of new features in this release drive performance and scalability to new heights. A large focus is on speed of data access, ease and flexibility of integration, and capability of visualization. These are all strategic areas in which Microsoft has focused on to add value since SQL Server 2005.

SQL Server History

SQL Server has grown considerably over the past two decades from its early roots with Sybase.

In 1989, Microsoft, Sybase, and Ashton-Tate jointly released **SQL Server 1.0**. The product was based on Sybase SQL Server 3.0 for UNIX and VMS.

SQL Server 4.2.1 for Windows NT released in 1993. Microsoft began making changes to the code.

SQL Server 6.0 (code named SQL 95) released in 1995. In 1996, the 6.5 upgrade (Hydra) was released in 1996. It included the first version of Enterprise Manager (StarFighter I) and SQL Server Agent (StarFighter II.)

SQL Server 7.0 (Sphinx), released in 1999 and was a full rewrite of the database engine by Microsoft. From a code sense, this was the first Microsoft SQL Server. SQL Server 7 also included English Query (Argo), OLAP Services (Plato), Replication, Database Design and Query tools (DaVinci), and Full-Text Search (aptly code named Babylon). Data Transformation Services (DTS) was introduced.

Part I: Laying the Foundations

SQL Server 2000 (Shiloh) 32-bit, version 8, introduced SQL Server to the enterprise with clustering, better performance, and OLAP. It supported XML through three different XML add-on packs. It added user-defined functions, indexed views, clustering support, OLAP, Distributed Partition Views, and improved Replication. SQL Server 2000 64-bit version for Intel Itanium (Liberty) released in 2003, along with the first version of Reporting Services (Rosetta) and Data Mining tools (Aurum). DTS becomes powerful and gained in popularity. Northwind joined Pubs as the sample database.

SQL Server 2005 (Yukon), version 9, was another rewrite of the database engine and pushed SQL Server further into the enterprise space. In 2005, a ton of new features and technologies were added including Service Broker, Notification Services, CLR, XQuery and XML data types, and SQLOS. T-SQL gained try-catch, and the system tables were replaced with Dynamic Management Views. Management Studio replaced Enterprise Manager and Query Analyzer. DTS was replaced by Integration Services. English Query was removed, and stored procedure debugging was moved from the DBA interface to Visual Studio. AdventureWorks and AdventureWorksDW replaced Northwind and Pubs as the sample database. SQL Server 2005 supported 32-bit, 64x, and Itanium CPUs. Steve Ballmer publically vowed to never again make customers wait 5 years between releases and to return to a 2-to-3-year release cycle.

SQL Server 2008 (Katmai), version 10, is a natural evolution of SQL Server adding Policy-Based Management, Data Compression, Resource Governor, and new beyond relational data types. Notification Services went the way of English Query. T-SQL finally has date and time data types, table-valued parameters, the debugger returns, and Management Studio gets IntelliSense.

SQL Server 2008R2, version 10.5, is a release mostly focused on new business intelligence features and SharePoint 2010 supportability. The list of major new work and code in the SQL Server 2005 and 2008/R2 releases have been fully covered in previous editions, but the high points would be SQLCLR (this was the integration of another long-term strategy project); XML support; Service Broker; and Integration Services, which is all ground up code. Microsoft formed a new team built on the original members of the DTS team, adding in some C++, hardware, AS and COM+ folks, and Report Builder. Additional features to support SharePoint 2010 functionality and other major releases are also critically important. Now you have SQL 2012; so look at where this new release can carry you forward.

SQL Server in the Database Market

SQL Server's position in the database market has consistently grown over time. This section discusses some of the primary competition to SQL Server, and what makes SQL a strong choice for data management, business intelligence, and cloud computing along with the strength of the SQL Server community.

SQL Server's Competition

SQL Server competes primarily with two other major database platforms, Oracle and IBM's DB2. Both of these products have existed for longer than SQL Server, but the last four releases of SQL Server have brought them closer together. They are adding features that SQL has had for years and vice versa. Many of the scalability improvements added since SQL 2005 have been directly focused on overtaking the performance and other qualities of these products. Microsoft has succeeded in these releases in besting benchmarks set by many other products both in the relational database platforms as well as in data integration, analytics, and reporting. These improvements, along with the strongest integrated ecosystem, including cloud (Windows Azure SQL Database), portal (SharePoint 2010), and business intelligence make SQL Server the market leader.

Strength of Community

SQL Server has one of the strongest communities of any technology platform. There are many websites, blogs, and community contributors that make up a great ecosystem of support. Some great avenues to get involved with include the following:

- PASS (Professional Association of SQL Server) SQLPASS.org
- SQL Saturday events — SQLSaturday.com
- SQLServerCentral.com
- BIDN.com
- MSSQLTips.com
- SQLServerPedia.com
- Twitter.com — #SQLHelp

Many of these are started and operated by Microsoft SQL Server MVPs and companies focused on SQL Server, education, and mentoring.

SQL Server Components

SQL Server is composed of the database engine, services, business intelligence tools, and other items including cloud functionality. This section outlines the major components and tools you need to become familiar with as you begin to explore this platform.

Database Engine

The SQL Server *Database Engine*, sometimes called the *relational engine*, is the core of SQL Server. It is the component that handles all the relational database work. SQL is a descriptive language, meaning that SQL describes only the question to the engine; the engine takes over from there.

Within the relational engine are several key processes and components, including the following:

- The *Algebrizer* checks the syntax and transforms a query to an internal representation used by the following components.
 - SQL Server's *Query Optimizer* determines how to best process the query based on the costs of different types of query-execution operations. The estimated and actual query-execution plans may be viewed graphically, or in XML, using Management Studio or SQL Profiler.
 - The *Query Engine*, or *Query Processor* executes the queries according to the plan generated by the Query Optimizer.
 - The *Storage Engine* works for the Query Engine and handles the actual reading and writing to and from the disk.
 - The *Buffer Manager* analyzes the data pages used and prefetches data from the data file(s) into memory, thus reducing the dependency on disk I/O performance.
 - The *Checkpoint* process writes dirty data pages (modified pages) from memory to the data file.
 - The *Resource Monitor* optimizes the query plan cache by responding to memory pressure and intelligently removing older query plans from the cache.
 - The *Lock Manager* dynamically manages the scope of locks to balance the number of required locks with the size of the lock.
 - SQL Server eats resources for lunch, and for this reason it needs direct control of the available resources (memory, threads, I/O request, and so on). Simply leaving the resource management to Windows isn't sophisticated enough for SQL Server. SQL Server includes its own OS layer, called *SQLOS*, which manages all its internal resources.

SQL Server 2012 supports installation of many instances of the relational engine on a physical server. Although they share some components, each instance functions as a complete separate installation of SQL Server.

Services

The following components are client processes for SQL Server used to control, or communicate with, SQL Server.

SQL Server Agent

The Server Agent is an optional process that, when running, executes SQL jobs and handles other automated tasks. It can be configured to automatically run when the system boots or may be started from the SQL Server Configuration Manager or the Management Studio's Object Explorer.

Database Mail

The Database Mail component enables SQL Server to send mail to an external mailbox through SMTP. Mail may be generated from multiple sources within SQL Server, including T-SQL code, jobs, alerts, Integration Services, and maintenance plans.

Microsoft Distributed Transaction Coordinator (MSDTC)

The Distributed Transaction Coordinator is a process that handles dual-phase commits for transactions that span multiple SQL Servers. DTC can be started from within Windows' Computer Administration/Services. If the application regularly uses distributed transactions, you should start DTC when the operating system starts.

Business Intelligence

Business intelligence (BI) is the name given to the discipline and tools that enable the management of data for the purpose of analysis, exploration, reporting, mining, and visualization. Although aspects of BI appear in many applications, the BI approach and toolset provide a rich and robust environment to understand data and trends.

SQL Server provides a great toolset to build BI applications, which explains Microsoft's continued gains in the growing BI market. SQL Server includes three services designed for BI: Integration Services (IS, sometimes called SSIS for SQL Server Integration Services), Reporting Services (RS), and Analysis Services (AS). Development for all three services can be done using the new SQL Server Data Tools, which is the new combining of Business Intelligence Development Studio and database development into a new environment in Visual Studio.

SSIS

Integration Services moves data among nearly any types of data sources and is SQL Server's Extract-Transform-Load (ETL) tool. IS uses a graphical tool to define how data can be moved from one connection to another connection. IS packages have the flexibility to either copy data column for column or perform complex transformations, lookups, and exception handling during the data move. IS is extremely useful during data conversions, collecting data from many dissimilar data sources, or gathering for data warehousing data that can be analyzed using Analysis Services.

IS has many advantages over using custom programming or T-SQL to move and transform data; chief among these are speed and traceability. If you have experience with other databases but are new to SQL Server, this is one of the tools that will impress you. If any other company were marketing SSIS, it would be the flagship product, but instead it's bundled inside SQL Server without much fanfare and at no extra charge. Be sure to find the time to explore IS.

SSAS

The Analysis Services service hosts two key components of the BI toolset: Online Analytical Processing (OLAP) hosts multidimensional databases where data is stored in cubes, whereas Data Mining provides methods to analyze datasets for nonobvious patterns in the data.

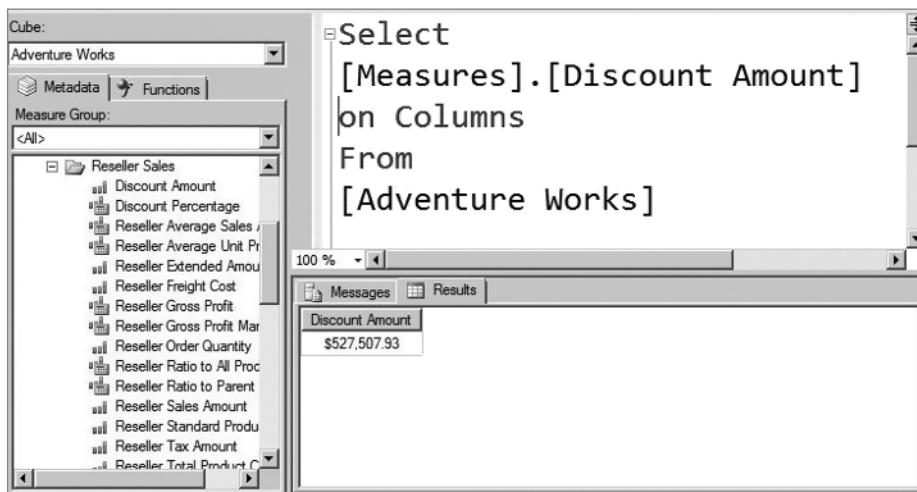
OLAP

Building cubes in a multidimensional database provides a fast, pre-interpreted, flexible analysis environment. Robust calculations can be included in a cube for later query and reporting, going a long way toward the “one version of the truth” that is so elusive in many organizations. Results can be used as the basis for reports, but the most powerful uses involve the interactive data exploration using tools such as Excel pivot tables or similar query and analysis applications. Tables and charts that summarize billions of rows can be generated in seconds, allowing users to understand the data in ways they never thought possible.

Although relational databases in SQL Server are queried using T-SQL, cubes are queried using the Multidimensional Expressions (MDX), a set-based query language tailored to retrieving multidimensional data. (See Figure 1-1.) This enables relatively easy custom application development in addition to standard analysis and reporting tools.

FIGURE 1-1

Example of MDX query in Analysis Services.



Data Mining

Viewing data from cubes or even relational queries can reveal the obvious trends and correlations in a dataset, but data mining can expose the nonobvious ones. The robust set of mining algorithms enables tasks such as finding associations, forecasting, and classifying cases into groups. When a model is trained on an existing set of data, it can predict new cases that occur, for example, predicting the most profitable customers to spend scarce advertising dollars on or estimating expected component failure rates based on its characteristics.

SSRS

Reporting Services (RS) for SQL Server 2012 is a full-featured, web-based, managed reporting solution. RS reports can be exported to PDF, Excel, or other formats with a single click and are easy to build and customize.

Reports are defined graphically or programmatically and stored as .rdl files in the RS databases in SQL Server. They can be scheduled to be pre-created and cached for users, e-mailed to users, or generated by users on-the-fly with parameters. Reporting Services is bundled with SQL Server so there are no end-user licensing issues. It's essentially free; although most DBAs place it on its own dedicated server for better performance. There is new functionality in SSRS 2012 with the addition of Power View. This is a SharePoint integrated feature that provides for rich drag and drop visualization and data exploration. It is one of the hottest new features in SQL 2012.

Tools and Add-Ons

SQL Server 2012 retains most of the UI feel of SQL Server 2008, with a few significant enhancements.

SQL Server Management Studio

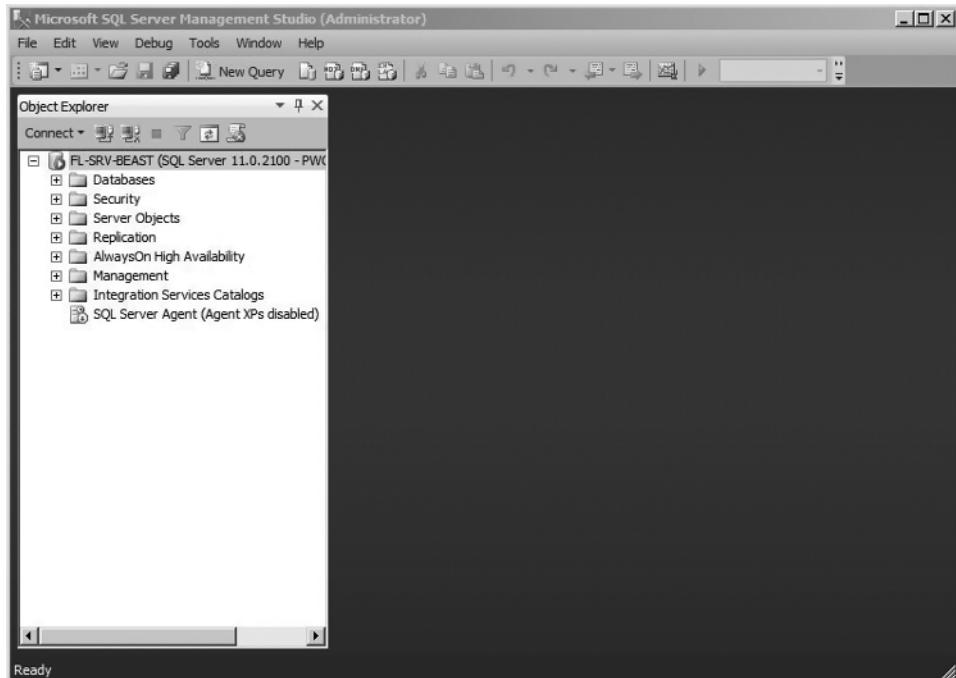
Management Studio is a Visual Studio-esque integrated environment that's used by database administrators and database developers. At its core is the visual Object Explorer complete with filters and the capability to browse all the SQL Server servers (database engine, Analysis Services, Reporting Services, and so on). Management Studio's Query Editor is an excellent way to work with raw T-SQL code, and it's integrated with the Solution Explorer to manage projects. Although the interface can look crowded (see Figure 1-2), the windows are easily configurable and can auto-hide.

SQL Server Configuration Manager

This tool is used to start and stop any server, set the start-up options, and configure the connectivity. It may be launched from the Start menu or from Management Studio. It can show you all the services and servers running on a particular server.

FIGURE 1-2

SQL Server Management Studio Query Interface.



SQL Profiler/Trace/Extended Events

SQL Server has the capability to expose a trace of selected events and data points. The server-side trace has nearly no load on the server. SQL Profiler is the UI for viewing traces in real time (with some performance cost) or viewing saved Trace files. Profiler is great for debugging an application or tuning the database. Profiler is being deprecated in favor of extended events. This will enable a deeper level of tracing with a decreased load on the server overall. This feature is continually enhanced and grown by support for other features such as Reporting services, Analysis Services, etc.

Performance Monitor

Although Profiler records large sets of details concerning SQL traffic and SQL Server events, Performance Monitor is a visual window into the current status of the selected performance counters. Performance Monitor is found within Windows's administrative tools. When SQL Server is installed, it adds a ton of useful performance counters to Performance Monitor. It's enough to make a network administrator jealous.

Database Engine Tuning Advisor

The Database Engine Tuning Advisor analyzes a batch of queries (from Profiler) and recommends index and partition modifications for performance. The scope of changes it can recommend is configurable, and the changes may be applied in part or in whole at the time of the analysis or later. The features of DBTA have been significantly enhanced in this newest version.

Command-Line Utilities

You can use various command-line utilities to execute SQL code (`sqlcmd`) or perform bulk copy program (`bcp`) from the DOS prompt or a command-line scheduler. Integration Services and SQL Server Agent have rendered these tools somewhat obsolete, but in the spirit of extreme flexibility, Microsoft still includes them.

Management Studio has a mode that enables you to use the Query Editor as if it were the command-line utility `sqlcmd`.

Online Resources

The SQL Server documentation team did an excellent job with Books Online (BOL) — SQL Server's mega help on steroids. The articles tend to be complete and include several examples. The indexing method provides a short list of applicable articles. BOL may be opened from Management Studio or directly from the Start menu.

BOL is well integrated with the primary interfaces. Selecting a keyword within Management Studio's Query Editor and pressing F1 launches BOL to the selected keyword. The Enterprise Manager help buttons can also launch the correct BOL topic.

Management Studio also includes a dynamic Help window that automatically tracks the cursor and presents help for the current keyword.

Searching returns both online and local MSDN articles. In addition, BOL searches the Codezone Community for relevant articles.

The Community Menu and Developer Center both launch web pages that enable users to ask a question or learn more about SQL Server.

CodePlex.com

If you haven't discovered [CodePlex.com](#), allow me to introduce it to you. [CodePlex.com](#) is Microsoft's site for open source code. That's where you can find AdventureWorks, the official sample database for SQL Server 2012, along with AdventureWorksLT (a smaller version for AdventureWorks) and AdventureWorksDW (the BI companion to AdventureWorks).

Editions of SQL Server 2012

The edition layout of SQL Server has changed again with this release to align closer with the way organizations use the product. Following are three main editions:

- **Enterprise:** This edition focused on mission critical applications and data warehousing.
- **Business intelligence:** This new edition has premium corporate features and self-service business intelligence features. If your environment is truly mission critical however, this may be missing some key features you might want. The key is to leverage this edition on your BI servers and use Enterprise where needed.
- **Standard:** This edition remains to support basic database capabilities including reporting and analytics.

You may wonder about the previous editions and how to move from what you have to the new plan. Following is a breakdown of deprecated editions and where the features now reside.

- **Datacenter:** Its features are now available in Enterprise Edition.
- **Workgroup:** Standard will become your edition for basic database needs.
- **Standard for small business:** Standard becomes your sole edition for basic database needs.

Notable SQL Server 2012 Enhancements

SQL 2012 has added many areas to its ecosystem. This includes new appliances, integration with “Big Data,” and connectors that leverage this technology as sources and destinations for analytics. Reference architectures have been improved and are released with improvements for SQL 2012. New features that add incredible performance boosts make these architectures a major weapon in return on investment (ROI) for many organizations.

Many of the important features that have been added to SQL Server 2012 fall into several categories, including the following:

- **Availability Enhancements**
 - AlwaysOn Failover Cluster instances
 - AlwaysOn Availability Groups
 - Online operations
- **Manageability Enhancements**
 - SQL Server Management Studio enhancements
 - Contained databases
 - Data-Tier Applications

- Windows PowerShell
- Database Tuning Advisor enhancements
- New Dynamic Management Views and Functions
- **Programmability Enhancements**
 - FileTables
 - Statistical Semantic Search functionality
 - Full-Text Search improvements
 - New and improved Spatial features
 - Metadata discovery and Execute Statement metadata support
 - Sequence Objects
 - THROW statement
 - 14 new T-SQL functions
 - Extended Events enhancements and more
- **Security Enhancements**
 - Enhanced Provisioning during setup
 - New permissions levels
 - New role management
 - Significant SQL Audit enhancements
 - Improved Hashing algorithms
- **Scalability and Performance Enhancements**
 - ColumnStore Indexes and Velocity
 - Online Index operation support for x(max) columns
 - Partition support increased to 15,000
- **Business Intelligence Features**
 - New Data Cleansing Components
 - Improved usability for SSIS and new deployment functionality
 - Master Data functionality has been significantly enhanced
 - New exciting features for Power Pivot
 - Power View data exploration and visualization
 - Tabular Models in SSAS
 - Expanded Extended Events throughout the BI ecosystem

These enhancements are discussed in detail through the upcoming chapters. More exciting details to come!

Summary

SQL Server 2012 has created many new opportunities for building some incredible scalable and high-performance applications and solutions. Many improvements have been added for availability, performance, configuration, intelligence and insight, and cloud functionality. This book covers all these new features, how to use them, and how to best leverage them for your organization.