

# 1

## SQL Server 2014 Architecture

### WHAT'S IN THIS CHAPTER?

---

- ▶ Getting to know the new important features in SQL Server 2014
- ▶ Understanding how new features relate to data professionals based on their roles
- ▶ Understanding SQL Server architecture
- ▶ Getting to know the editions of SQL Server and how they affect the data professional

SQL Server 2012 offered a fresh look at how organizations and their developers, information workers, and executives use and integrate data. SQL Server 2014 continues that story and pushes forward into the future with the integration of cloud and in-memory technologies. Exciting new features and improvements provide breakthrough performance, availability, and manageability for mission-critical workloads on premises or in the cloud. This chapter provides enough information to give you an understanding of how SQL Server operates.

### SQL SERVER 2014 ECOSYSTEM

The SQL Server data platform has become quite large over the past few releases. This first section provides a review of the overall SQL Server ecosystem, which is now referred to as less of a product and more of a data platform, because many interactions with other products and features drive increased performance, scale, and usability. SQL Server 2014 focuses on the following areas:

- ▶ **Mission-critical performance**—New In-Memory Online Transaction Processing (OLTP) features that allow faster performance with no application changes, updatable

columnstore indexes, buffer pool extensions using solid-state drives (SSDs), and AlwaysOn enhancements (including support of up to eight replicas) make this the most powerful, available release of SQL Server.

- ▶ **Faster insights**—With new Office-based Business Intelligence (BI) tools such as Power Query and Power Map, as well as enhancements on Power View and PowerPivot, getting insights on your data has never been easier. Also, enterprise options such as Parallel Data Warehouse with Polybase allow an organization to easily explore its Big Data and gain never-before-seen insights on its data by leveraging the full power of the Microsoft BI stack.
- ▶ **Platform for hybrid cloud**—Whether your environment is pure on-premises, virtualized, or all in the cloud, SQL Server 2014 has options for you. New features such as the Microsoft SQL Server Backup to Windows Azure Tool enables backup to Windows Azure Blob Storage, and encrypts and compresses SQL Server backups stored locally or in the cloud. You can also now choose a Windows Azure virtual machine (VM) as a replica in an AlwaysOn Availability Group configuration.

## NEW IMPORTANT FEATURES IN 2014

You will be excited about a number of new things, depending on your role and how you use SQL Server. This section touches on the features you should be checking out and getting your hands on. Many of these features are quick to get up and running, which is exciting for those readers who want to begin delivering impact right away.

## Production DBA

*Production database administrators (DBAs)* are a company's insurance policy that the production database won't go down. If the database does go down, the company cashes in its insurance policy in exchange for a recovered database. The Production DBA also ensures that the server performs optimally, and promotes database changes from development to quality assurance (QA) to production. New features include the following:

- ▶ **In-Memory OLTP**—This release offers a new memory-optimized OLTP database engine for SQL Server. Experience significant performance and scalability gains to your application with little to no code changes.
- ▶ **AlwaysOn Availability Groups**—Availability functionality includes Availability Groups and the ability to fail over databases in groups that mimic applications. Though not a new feature in 2014, enhancements in this release include support of up to eight secondaries (up from four in the 2012 release), as well as performance and management improvements.
- ▶ **SQL Server Backup to Windows Azure tool**—This free tool enables backing up to Windows Azure Blob Storage. It encrypts and compresses backups stored locally or in the cloud. This tool supports SQL Server 2005 and later.
- ▶ **SQL Server Backup to URL**—Originally released in SQL Server 2012 SP1 CU2, this feature is now fully integrated into the Management Studio interface so that you can back up or restore from the Windows Azure Blob Storage service.

- **SQL Server Managed Backup to Windows Azure**—This new feature fully automates SQL Server backups (full and transaction log) to the Windows Azure Blob Storage service based on retention period and transaction workload on the database.
- **Columnstore indexes**—The 2014 release includes updatable clustered columnstore indexes for optimizing large data volumes.
- **Encrypted backups**—Support is now included for encrypted backups, including algorithm support for Advanced Encryption Standard (AES) 128, AES 192, AES 256, and Triple Data Encryption Standard (DES). You must use either a certificate or an asymmetric key to perform encryption during backup.
- **Delayed durability**—This includes the ability to reduce latency by designating some or all transactions as delayed durable. This asynchronous process reports a `COMMIT` as successful before log records are written to disk. Delayed durable transactions become durable when the transaction log entries are flushed to disk in chunks. This feature is available for In-Memory OLTP only.
- **Compression and partitioning**—Improved compression and partitioning include the capability to rebuild single partitions.
- **Resource Governor**—This now enables you to specify limits on physical I/O within a resource pool.

## Development DBA

Since the release of SQL Server 2000, there has been a trend away from full-time Production DBAs, and the role has merged with that of the *Development DBA*. The trend may have slowed, though, with laws such as Sarbanes-Oxley, which require a separation of power between the person developing the change and the person implementing the change. In a large organization, a Production DBA may fall into the operations department, which consists of the network of administrators and Windows-support administrators. In other instances, a Production DBA may be placed in a development group. This removes the separation of power that is sometimes needed for regulatory reasons.

Development DBAs play a traditional role in an organization. They wear more of a developer's hat, and serve as the development staff's database experts and representatives. This administrator ensures that all stored procedures are optimally written and that the database is modeled correctly, both physically and logically. The Development DBA also may be the person who writes the migration processes to upgrade the database from one release to the next. The Development DBA typically does not receive calls at 2:00 A.M like the Production DBA might for failed backups or similar problems. Following are some of the things Development DBAs should be excited about in this new release:

- Transact SQL (T-SQL) enhancements include features such as the inline specification of indexes for disk-based tables. The `SELECT...INTO` statement can now operate in parallel with databases, in compatibility mode set to at least 110.
- Updates for SQL Server Data Tools for Business Intelligence (SSDT-BI) include support for creating Power View reports against multidimensional models. Additional improvements to SSDT-BI include support to create projects for older versions (2005+) of Analysis Services and Reporting Services. Backward compatibility for Integration Services projects is not yet supported.

The Development DBA typically reports to the development group, and receives requests from a business analyst or another developer. In a traditional sense, Development DBAs should never have modification access to a production database. However, they should have read-only access to the production database to debug in a time of escalation.

## Business Intelligence DBA and Developer

The *BI DBA* is a role that has evolved as a result of the increased capabilities of SQL Server. In SQL Server 2012, BI grew to be an incredibly important feature set that many businesses could not live without. The BI DBA (or developer) is an expert at these features. SQL Server 2014 improves upon the end-user BI experience started in SQL Server 2012, with all new extensibility through the Power BI solution with Office365, as well as the existing BI offerings.

*Development BI DBAs* specialize in the best practices, optimization, and use of the BI toolset. In a small organization, a Development BI DBA may create your SQL Server Integration Services (SSIS) packages to perform extract, transform, and load (ETL) processes or reports for users. In a large organization, developers create the SSIS packages and SQL Server Reporting Services (SSRS) reports. The Development BI DBA is consulted regarding the physical implementation of the SSIS packages and SQL Server Analysis Services (SSAS) cubes, or SSAS tabular models.

Development BI DBAs may be responsible for the following types of functions:

- Modeling/consulting regarding Analysis Services cubes/tabular models and solutions
- Creating reports using Reporting Services
- Creating/consulting around ETL using SSIS
- Developing deployment packages to be sent to the Production DBA

These responsibilities, coupled with the following new features, make for an exciting time for the BI-oriented folks:

- Rapid data discovery with Power View and PowerPivot
- Lightning-fast query response time with use of columnstore indexes
- Users can easily discover, access, and combine various data sources themselves using Power Query
- Users can create powerful geospatial visualizations using Power Map
- Managed Self-Service BI with SharePoint and BI Semantic Model
- Credible, consistent data with Data Quality Services and Master Data Management capabilities
- Robust data warehousing solutions with Parallel Data Warehouse and Reference Architectures

## SQL SERVER ARCHITECTURE

SQL Server 2014 brings many enhancements to your data platform solution. These not only give you a whole new level of breakthrough performance and scalability, but also provide exciting Self-Service BI opportunities that enable end users to explore and analyze their data faster and easier than ever before. New features such as in-memory OLTP provide massive performance gains to your applications with little to no schema changes for the application. Other features such as AlwaysOn Availability Groups allow you to quickly and easily scale out your database application.

To know how to best leverage all of these exciting features and functionality, it's important to understand the basic architecture of SQL Server. This section covers the primary file types in SQL Server 2014, file management, SQL Client, and system databases. It also covers an overview of schemas, synonyms, and Dynamic Management Objects. Finally, it also goes into the data types, editions, and licensing options.

### Database Files and Transaction Log

The architecture of database and transaction log files remains relatively unchanged from prior releases. Database files serve two primary purposes, depending on their type. *Data files* hold the data, indexes, and other data support structure within the database. *Log files* hold the data from committed transactions to ensure consistency in the database.

#### Database Files

A database may consist of multiple *filegroups*. Each filegroup must contain one or more physical data files. Filegroups ease administrative tasks for a collection of files. Data files are divided into 8KB data pages, which are part of 64KB extents. You can specify how full each data page should be with the fill factor option of the `create/alter index` T-SQL command. In SQL Server 2014, you continue to have the capability to bring your database partially online if a single file is corrupted. In this instance, the DBA can bring the remaining files online for reading and writing, and users receive an error if they try to access the other parts of the database that are offline.

In SQL 2000 and before, the largest row you could write was 8,060 bytes. The exceptions to this limit are `text`, `ntext`, `image`, `varchar(max)`, `varbinary(max)`, and `nvarchar(max)` columns, which may each be up to 2GB, and are managed separately. Beginning with SQL 2005, the 8KB limit applies only to those columns of fixed length. The sum of fixed-length columns and pointers for other column types must still be less than 8,060 bytes per row. However, each variable-length column may be up to 8KB in size, allowing for a total row size of well more than 8,060 bytes. If your actual row size exceeds 8,060 bytes, you may experience some performance degradation because the logical row must now be split across multiple physical 8,060-byte rows.

#### Transaction Log

The purpose of the *transaction log* is to ensure that all committed transactions are persisted in the database and can be recovered, either through rollback or point-in-time recovery. The transaction log is a *write-ahead log*. As you make changes to a database in SQL Server, the data is written to the

log, and then the pages that need to be changed are loaded into memory. The pages are then “dirtyed” by having the changes written to them. Upon checkpoint, the dirty pages are written to disk, making them now clean pages that no longer need to be part of the write buffer. This is why you may see your transaction log grow significantly in the middle of a long-running transaction even if your recovery model is set to simple. (Chapter 17, “Backup and Recovery,” covers this in much more detail.)

## SQL Server Native Client

The SQL Server Native Client (SNAC) is a data-access method that shipped with SQL Server 2005 and was enhanced in 2012. It is used by both an online object linking and embedding (OLE) database and Open Database Connectivity (ODBC) for accessing SQL Server. The SQL Server Native Client simplifies access to SQL Server by combining the OLE DB and ODBC libraries into a single access method. The access type exposes the following features in SQL Server:

- Database mirroring
- AlwaysOn readable secondary routing
- Multiple Active Result sets (MARS)
- Native client support for LocalDB
- FILESTREAM support
- Snapshot isolation
- Query notification
- XML data type support
- User-defined data types (UDTs)
- Encryption
- Performing asynchronous operations
- Using large value types
- Performing bulk copy operations
- Table-value parameters
- Large Common Language Runtime (CLR) user-defined types
- Password expiration
- Service Principal Name (SPN) support in client connections

With these features, you can use the feature in other data layers such as Microsoft Data Access Components (MDAC), but it takes more work. MDAC still exists, and you can use it if you don’t need some of the new functionality of SQL Server 2008/2012. If you develop a COM-based application, you should use SQL Native Client; and if you develop a managed code application (such as in C#), you should consider using the .NET Framework Data Provider for SQL Server, which is robust and includes the SQL Server 2008/2012 features as well.

SQL Server 2014 actually installs SQL Server 2012 Native Client because there is no SQL Server 2014 Native Client. Other major changes include the deprecation of the ODBC driver in SQL Server Native Client. Because of this deprecation, there will be no more updates to the ODBC driver in SQL Server Native Client. The successor of the ODBC driver in SQL Server Native Client is called the Microsoft ODBC Driver 11 for SQL Server on Windows, which is installed with SQL Server 2014 by default.

## Standard System Databases

The system databases in SQL Server are crucial, and you should leave them alone most of the time. The only exceptions to that rule are the `model` database (which enables you to deploy a change such as a stored procedure to any new database created) and `tempdb` (which will probably need to be altered to help with scaling your workload). The following sections go through the standard system databases in detail.

**WARNING** *If certain system databases are tampered with or corrupted, you run the risk that SQL Server will not start. The master database contains all the stored procedures and tables needed for SQL Server to remain online.*

## The resource Database

SQL Server 2005 added the `resource` database. This database contains all the read-only critical system tables, metadata, and stored procedures that SQL Server needs to run. It does not contain any information about your instance or your databases because it is written to only during an installation of a new service pack. The `Resource` database contains all the physical tables and stored procedures referenced logically by other databases. You can find the database in a default installation at `C:\Program Files\Microsoft SQL Server\MSSQL14.MSSQLSERVER\MSSQL\Binn`, and there is only one `Resource` database per instance.

**NOTE** *The use of drive C: in the path assumes a standard setup. If your machine is set up differently, you may need to change the path to match your setup. In addition, the `.MSSQLSERVER` is the instance name. If your instance name is different, use your instance name in the path.*

In SQL Server 2000, when you upgraded to a new service pack, you needed to run many long scripts to drop and re-create system objects. This process took a long time to run, and created an environment that couldn't be rolled back to the previous release after the service pack. Since SQL Server 2012, when you upgrade to a new service pack or hot fix, a copy of the `resource` database overwrites the old database. This enables you to both quickly upgrade your SQL Server catalog and roll back a release.

The `resource` database cannot be seen through Management Studio and should never be altered unless you're under instruction to do so by Microsoft Product Support Services (PSS). You can connect to the database under certain single-user mode conditions by typing the command `USE MSSQLSystemResource`. Typically, a DBA runs simple queries against it while connected to any database, instead of having to connect to the `resource` database directly. Microsoft provides some functions that enable this access. For example, if you were to run this query while connected to any database, it would return your `resource` database's version and the last time it was upgraded:

```
SELECT serverproperty('resourceversion') ResourceDBVersion,  
serverproperty('resourcelastupdatedatetime') LastUpdateDate;
```

**NOTE** *Do not place the resource database on an encrypted or compressed drive. Doing this may cause upgrade or performance issues.*

## The master Database

The `master` database contains the metadata about your databases (database configuration and file location), logins, and configuration information about the instance. You can see some of the metadata stored in `master` by running the following query, which returns information about the databases that exist on the server:

```
SELECT * FROM sys.databases;
```

The main difference between the `resource` and `master` databases is that the `master` database holds data specific to your instance, whereas the `resource` database just holds the schema and stored procedures needed to run your instance but does not contain any data specific to your instance.

**WARNING** *You should rarely create objects in the master database. If you create objects here, you may need to make more frequent master database backups.*

## The tempdb Database

The `tempdb` database is similar to the operating system paging file. It's used to hold temporary objects created by users, temporary objects needed by the database engine, and for row-version information. The `tempdb` database is created each time you restart SQL Server. The database will be re-created to its original database size when the SQL Server is started. Because the database is re-created each time, you cannot back it up. Data changes made to objects in the `tempdb` database benefit from reduced logging. You must have enough space allocated to your `tempdb` database because many operations that you use in your database applications use `tempdb`. Generally speaking, you should set `tempdb` to autogrow as it needs space. Typically, your `tempdb` size varies, but you should understand how much temp space your application will use at peak, and ensure that there is enough



space with 15–20 percent overhead for growth. If there is not enough space, the user may receive one of the following errors:

- 1101 or 1105—The session connecting to SQL Server must allocate space in `tempdb`.
- 3959—The version store is full.
- 3967—The version store must shrink because `tempdb` is full.

## The model Database

`model` is a system database that serves as a template when SQL Server creates a new database. As each database is created, SQL Server copies the `model` database as the new database. The only time this does not apply is when you restore or attach a database from a different server.

If a table, stored procedure, or database option should be included in each new database that you create on a server, you may simplify the process by creating the object in `model`. When the new database is created, `model` is copied as the new database, including the special objects or database settings you have added to the `model` database. If you add your own objects to `model`, it should be included in your backups, or you should maintain a script that includes the changes.

## The msdb Database

`msdb` is a system database that contains information used by SQL Server Agent, log shipping, SSIS, and the backup-and-restore system for the relational database engine. The database stores all the information about jobs, operators, alerts, policies, and job history. Because it contains this important system-level data, you should back up this database regularly.

## Schemas

Schemas enable you to group database objects together. You may want to do this for ease of administration because you can apply security to all objects within a schema. Another reason to use schemas is to organize objects so that the users may find the objects they need easily. For example, you may create a schema called `HumanResource` and place all your employee tables and stored procedures into it. You could then apply security policies on the schema to allow appropriate access to the objects contained within it.

When you refer to an object, you should always use the two-part name. The `dbo` schema is the default schema for a database. An `Employee` table in the `dbo` schema is referred to as `dbo.Employee`. Table names must be unique within a schema. You could create another table called `Employee` in the `HumanResources` schema. It would be referred to as `HumanResources.Employee`. This table actually exists in the `AdventureWorks` sample database for SQL Server 2014. (All SQL Server 2014 samples must be downloaded and installed separately from [wrox.com](http://wrox.com).) A sample query using the two-part name follows:

```
USE AdventureWorks2014
GO
SELECT BusinessEntityID, JobTitle
FROM HumanResources.Employee;
```

Prior to SQL 2005, the first part of the two-part name was the username of the object owner. The problem with that implementation was related to maintenance. If a user who owned objects were to leave the company, you could not remove that user login from SQL Server until you ensured that all the objects owned by the user were changed to a different owner. All the code that referred to the objects had to be changed to refer to the new owner. By separating ownership from the schema name, SQL 2005 and later versions remove this maintenance problem.

## Synonyms

A *synonym* is an alias, or alternative name, for an object. This creates an abstraction layer between the database object and the user. This abstraction layer enables you to change some of the physical implementation and isolate those changes from the consumer. The following example is related to the use of linked servers. You may have tables on a different server that need to be joined to tables on a local server. You refer to objects on another server using the four-part name, as shown in the following code:

```
SELECT Column1, Column2
FROM LinkedServerName.DatabaseName.SchemaName.TableName;
```

For example, you might create a synonym for `LinkedServerName.DatabaseName.SchemaName.TableName` called `SchemaName.SynonymName`. Data consumers would refer to the object using the following query:

```
SELECT Column1, Column2
FROM SchemaName.SynonymName;
```

This abstraction layer now enables you to change the location of the table to another server, using a different linked server name, or even to replicate the data to the local server for better performance without requiring any changes to the code that refers to the table.

**NOTE** *A synonym cannot reference another synonym. The `object_id` function returns the `id` of the synonym, not the `id` of the related base object. If you need column-level abstraction, use a view instead.*

## Dynamic Management Objects

Dynamic Management Objects (DMOs) and functions return information about your SQL Server instance and the operating system. DMOs are grouped into two different classes: *Dynamic Management Views (DMVs)* and *Dynamic Management Functions (DMFs)*. DMVs and DMFs simplify access to data and expose new information that was not available in versions of SQL Server prior to 2005. DMOs can provide you with various types of information, from data about the I/O subsystem and RAM to information about Service Broker.

Whenever you start an instance, SQL Server begins saving server-state and diagnostic information in memory, which DMVs and DMFs can access. When you stop and start the instance, the information is flushed from the views and fresh data begins to be collected. You can query the views just like

any other table in SQL Server with the two-part qualifier. For example, the following query uses the `sys.dm_exec_sessions` DMV to retrieve the number of sessions connected to the instance, grouped by login name:

```
SELECT login_name, COUNT(session_id) as NumberSessions
FROM sys.dm_exec_sessions GROUP BY login_name;
```

Some DMFs are functions that accept parameters. For example, the following code uses the `sys.dm_io_virtual_file_stats` dynamic management function to retrieve the I/O statistics for the AdventureWorks data file:

```
USE AdventureWorks
GO
SELECT * FROM
sys.dm_io_virtual_file_stats(DB_ID('AdventureWorks2014'),
FILE_ID('AdventureWorks_Data'));
```

Many new DMVs and DMFs were introduced in SQL Server 2012. These views focus on improved insight into new and existing areas of functionality and include the following:

- AlwaysOn Availability Groups
- Change Data Capture–Related
- Change Tracking–Related
- Common Language Runtime–Related
- Database Mirroring–Related
- Database-Related
- Execution-Related
- SQL Server Extended Events
- FileStream and FileTable
- Full-Text Search and Semantic Search
- Index-Related
- I/O-Related
- Object-Related
- Query Notifications–Related
- Replication-Related
- Resource Governor–Related
- Security-Related
- Server-Related
- Service Broker–Related
- Spatial Data–Related

- SQL Server Operating System–Related
- Transaction–Related

SQL Server 2014 introduced the following new DMVs and DMFs:

- Cluster Shared Volume (CSV)–Related
- Buffer Pool Extension–Related
- In-Memory OLTP–Related

## SQL Server 2014 Data Types

Data types are the foundation of table creation in SQL Server. As you create a table, you must assign a data type for each column. This section covers some of the more commonly used data types in SQL Server. Even if you create a custom data type, it must be based on a standard SQL Server data type. For example, you can create a custom data type (`Address`) by using the following syntax, but notice that it is based on the SQL Server standard `varchar` data type:

```
CREATE TYPE Address
FROM varchar(35) NOT NULL;
```

If you change the data type of a column in a large table in SQL Server Management Studio's table designer interface, the operation may take a long time. You can observe the reason for this by scripting the change from the Management Studio interface. Management Studio creates a secondary temporary table with a name such as `tmpTableName` and then copies the data into the table. Finally, the interface deletes the old table and renames the new table with the new data type. Other steps along the way, of course, handle indexes and any relationships in the table.

If you have a large table with millions of records, this process can take more than 10 minutes, and, in some cases, more than 1 hour. To avoid this, you can use a simple one-line T-SQL statement in the query window to change the column's data type. For example, to change the data type of the `JobTitle` column in the `Employees` table to a `varchar(70)`, you could use the following syntax:

```
ALTER TABLE HumanResources.Employee ALTER COLUMN JobTitle Varchar(70);
```

**WARNING** *When you convert to a data type that may be incompatible with your data, you may lose important data. For example, if you convert from a numeric data type that has data, such as 15.415, to an integer, the number 15.415 would be rounded to a whole number.*

You may want to write a report against your SQL Server tables that displays the data type of each column inside the table. You can do this dozens of ways, but a popular method shown in the following example joins the `sys.objects` table with the `sys.columns` table. You may not be familiar with two functions in the following code. The `TYPE_NAME()` function translates the data type ID into its proper name. To go the opposite direction, you could use the `TYPE_ID()` function. The other function of note is `SCHEMA_ID()`, which is used to return the identity value for the schema. This is useful primarily when you want to write reports against the SQL Server metadata.

```

USE AdventureWorks2014;
GO
SELECT o.name AS ObjectName,
       c.name AS ColumnName,
       TYPE_NAME(c.user_type_id) as DataType
FROM sys.objects o
JOIN sys.columns c
ON o.object_id = c.object_id
WHERE o.name = 'Department'
and o.Schema_ID = SCHEMA_ID('HumanResources');

```

This code returns the following results (the *Name* data type is a user-defined type):

ObjectName	ColumnName	DataType
-----		
Department	DepartmentID	smallint
Department	Name	Name
Department	GroupName	Name
Department	ModifiedDate	datetime

## Character Data Types

Character data types include `varchar`, `char`, and `text`. This set of data types stores character data. The primary difference between the `varchar` and `char` types is data padding. If you have a column called `FirstName` that is a `varchar(20)` data type and you store the value of “Brian” in the column, only 5 bytes are physically stored, plus a little overhead. If you store the same value in a `char(20)` data type, all 20 bytes would be used. SQL inserts trailing spaces to fill the 20 characters.

**NOTE** *You might ask yourself, if you want to conserve space, why would you ever use a char data type? There is a slight overhead to using a varchar data type. For example, if you store a two-letter state abbreviation, you’re better off using a char(2) column. Although some DBAs have opinions about this that border on religious conviction, generally speaking, it’s good to find a threshold in your organization, and specify that anything below this size becomes a char versus a varchar. In general, a good guideline is that any column less than or equal to 5 bytes should be stored as a char data type instead of a varchar data type, depending on your application needs. Beyond that point, the benefit of using a varchar begins to outweigh the cost of the overhead.*

The `nvarchar` and `nchar` data types operate the same way as their `varchar` and `char` counterparts, except these data types can handle international Unicode characters. This comes at a cost, though. Data stored as Unicode consumes 2 bytes per character. If you were to store the value of “Brian” in an `nvarchar` column, it would use 10 bytes, and storing it as an `nchar(20)` would use 40 bytes. Because of this overhead and added space, you should not use Unicode columns unless you have a business or language need for them. Think ahead to the future and consider instances in which you might need them. If there is no future business need, avoid them.

Table 1-1 shows the data types with short descriptions and the amount of storage required.

**TABLE 1-1: SQL Server Data Types**

DATA TYPE	DESCRIPTION	STORAGE SPACE
Char (n)	N between 1 and 8,000 characters	n bytes
Nchar (n)	N between 1 and 4,000 Unicode characters	(2 × n bytes)
Nvarchar (max)	Up to ((2 to the 30th power) – 1) (1,073,741,823) Unicode characters	2 × characters stored + 2 bytes overhead
Text	Up to ((2 to the 31st power) – 1) (2,147,483,647) characters	1 byte per character stored + 2 bytes overhead
Varchar (n)	N between 1 and 8,000 characters	1 byte per character stored + 2 bytes overhead
Varchar (max)	Up to ((2 to the 31st power) – 1) (2,147,483,647) characters	1 byte per character stored + 2 bytes overhead

## Exact Numeric Data Types

Numeric data types consist of `bit`, `tinyint`, `smallint`, `int`, `bigint`, `numeric`, `decimal`, `smallmoney`, and `money`. Each of these data types stores different types of numeric values. The first data type, `bit`, stores only a null, 0, or a 1, which in most applications translates into true or false. Using the `bit` data type is perfect for on and off flags, and it occupies only a single byte of space. Table 1-2 shows other common numeric data types.

**TABLE 1-2: Exact Numeric Data Types**

DATA TYPE	DESCRIPTION	STORAGE SPACE
Bit	0, 1, or null	1 byte for each 8 columns of this data type
tinyint	Whole numbers from 0 to 255	1 byte
smallint	Whole numbers from –32,768 to 32,767	2 bytes
Int	Whole numbers from –2,147,483,648 to 2,147,483,647	4 bytes
bigint	Whole numbers from –9,223,372,036,854,775,808 to 9,223,372,036,854,775,807	8 bytes

DATA TYPE	DESCRIPTION	STORAGE SPACE
numeric(p,s) or decimal(p,s)	Numbers from $-1,038 + 1$ through $1,038 - 1$	Up to 17 bytes
money	$-922,337,203,685,477.5808$ to $922,337,203,685,477.5807$	8 bytes
smallmoney	$-214,748.3648$ to $214,748.3647$	4 bytes

Numeric data types, such as `decimal` and `numeric`, can store a variable number of digits to the right and left of the decimal place. *Scale* refers to the number of digits to the right of the decimal. *Precision* defines the total number of digits, including the digits to the right of the decimal place. For example, 14.88531 would be a `numeric(7,5)` or `decimal(7,5)`. If you were to insert 14.25 into a `numeric(5,1)` column, it would be rounded to 14.3.

## Approximate Numeric Data Types

The data types `float` and `real` are included in this group. They should be used when floating-point data must be represented. However, because they are approximate, not all values can be represented exactly.

The `n` in the `float(n)` is the number of bits used to store the mantissa of the number. SQL Server uses only two values for this field. If you specify between 1 and 24, SQL uses 24. If you specify between 25 and 53, SQL uses 53. The default is 53 when you specify `float()`, with nothing in parentheses.

Table 1-3 shows the approximate numeric data types with a short description and the amount of storage required.

**TABLE 1-3:** Approximate Numeric Data Types

DATA TYPE	DESCRIPTION	STORAGE SPACE
<code>float[(n)]</code>	$-1.79E + 308$ to $-2.23E-308,0$ , $2.23E-308$ to $1.79E + 308$	$N \leq 24 - 4$ bytes $N > 24 - 8$ bytes
<code>real()</code>	$-3.40E + 38$ to $-1.18E-38,0$ , $1.18E-38$ to $3.40E + 38$	4 bytes

**NOTE** *The synonym for real is float(24).*

## Binary Data Types

Binary data types such as `varbinary`, `binary`, and `varbinary(max)` store binary data such as graphic files, Word documents, or MP3 files. The values are hexadecimal 0x0 to 0xf. The image

data type stores up to 2GB outside the data page. The preferred alternative to an `image` data type is `varbinary(max)`, which can hold more than 8KB of binary data and generally performs slightly better than an `image` data type. Introduced in SQL Server 2012 was the capability to store `varbinary(max)` objects in operating system files via FileStream storage options. This option stores the data as files and is not subject to the 2GB size limit of `varbinary(max)`.

Table 1-4 shows the binary data types with a short description and the amount of storage required.

**TABLE 1-4:** Binary Data Types

DATA TYPE	DESCRIPTION	STORAGE SPACE
<code>Binary(n)</code>	N between 1 and 8,000 hex digits	n bytes
<code>Varbinary(n)</code>	N between 1 and 8,000 hex digits	1 byte per character stored + 2 bytes overhead
<code>Varbinary(max)</code>	Up to $2^{31} - 1$ (2,147,483,647) characters	1 byte per character stored + 2 bytes overhead

## Date and Time Data Types

The `datetime` and `smalldatetime` types both store date and time data. The `smalldatetime` is 4 bytes and stores from January 1, 1900, through June 6, 2079, and is accurate to the nearest minute.

The `datetime` data type is 8 bytes and stores from January 1, 1753, through December 31, 9999, to the nearest 3.33 millisecond.

SQL Server 2012 introduced four new date-related data types: `datetime2`, `datetimeoffset`, `date`, and `time`. You can find examples using these data types in SQL Server Books Online.

The `datetime2` data type is an extension of the `datetime` data type, with a wider range of dates. Time is always stored with hours, minutes, and seconds. You can define the `datetime2` data type with a variable parameter at the end—for example, `datetime2(3)`. The 3 in the preceding expression means to store fractions of seconds to three digits of precision, or 0.999. Valid values are between 0 and 7, with a default of 3.

The `datetimeoffset` data type is just like the `datetime2` data type, with the addition of the time offset. The time offset is + or – up to 14 hours, and contains the UTC offset so that you can rationalize times captured in different time zones.

The `date` data type stores the date only, a long-requested piece of functionality. Alternatively, the `time` data type stores the time only. The `time` data type also supports the `time(n)` declaration, so you can control granularity of the fractional seconds. As with `datetime2` and `datetimeoffset`, n can be between 0 and 7.

Table 1-5 shows the date/time data types with a short description and the amount of storage required.



TABLE 1-5: Date and Time Data Types

DATA TYPE	DESCRIPTION	STORAGE SPACE
Date	January 1, 1 to December 31, 9999	3 bytes
Datetime	January 1, 1753 to December 31, 9999, Accurate to nearest 3.33 millisecond	8 bytes
Datetime2 (n)	January 1, 1 to December 31, 9999 N between 0 and 7 specifies fractional seconds	6 to 8 bytes
Datetimeoffset (n)	January 1, 1 to December 31, 9999 N between 0 and 7 specifies fractional seconds +- offset	8 to 10 bytes
SmallDateTime	January 1, 1900 to June 6, 2079, Accurate to 1 minute	4 bytes
Time (n)	Hours:minutes:seconds. 9999999 N between 0 and 7 specifies fractional seconds	3 to 5 bytes

## Other System Data Types

Table 1-6 shows several other data types, which you have not yet seen.

TABLE 1-6: Other System Data Types

DATA TYPE	DESCRIPTION	STORAGE SPACE
Cursor	This contains a reference to a cursor, and may be used only as a variable or stored procedure parameter.	Not applicable
Hierarchyid	Contains a reference to a location in a hierarchy.	1 to 892 bytes + 2 bytes overhead
SQL_Variant	May contain the value of any system data type except text, ntext, image, timestamp, xml, varchar (max), nvarchar (max), varbinary (max), and user-defined data types. The maximum size allowed is 8,000 bytes of data + 16 bytes of metadata.	8,016 bytes
Table	Used to store a data set for further processing. The definition is like a <code>Create Table</code> . Primarily used to return the result set of a table-valued function; they can also be used in stored procedures and batches.	Dependent on table definition and number of rows stored
Timestamp or Rowversion	Unique per table, automatically stored value. Generally used for version stamping; the value is automatically changed on insert and with each update.	8 bytes

*continues*

TABLE 1-6 (continued)

DATA TYPE	DESCRIPTION	STORAGE SPACE
Uniqueidentifier	Can contain Globally Unique Identifier (GUID). guid values may be obtained from the Newsequentialid() function. This function returns values that are unique across all computers. Although stored as a binary 16, it is displayed as a char(36).	16 bytes
XML	This is Unicode by definition.	Up to 2GB

**NOTE** *A cursor data type may not be used in a Create Table statement.*

The XML data type stores an XML document or fragment. It is stored like an nvarchar(max) in size depending on the use of UTF-16 or UTF-8 in the document. The XML data type enables the use of special constructs for searching and indexing. (This is covered in more detail in Chapter 15, “Replication.”)

## CLR Integration

In SQL Server, you can also create your own data types, functions, and stored procedures using the portion of the Common Language Runtime (CLR) referred to as SQLCLR. This enables you to write more complex data types to meet your business needs in Visual Basic or C#, for example. These types are defined as a class structure in the base CLR language.

## EDITIONS OF SQL SERVER

SQL Server 2014 is available in several editions, and the features available to you in each edition vary widely. The editions you can install on your workstation or server also vary based on the operating system. The editions of SQL Server range from SQL Express on the lowest end to Enterprise Edition on the highest.

## Edition Overview

Three major editions of SQL Server 2014 are available. Additional, smaller editions are intended for select use cases, so they are not covered here. For more details on these, see [www.microsoft.com/sqlserver](http://www.microsoft.com/sqlserver). As shown in Table 1-7, SQL Server 2014 offers three principal SKUs:

- ▶ **Enterprise Edition**—This contains all the new SQL Server 2014 capabilities, including high availability and performance updates and features that make this a mission critical-ready database. This edition contains all the BI functionality as well.
- ▶ **Business Intelligence Edition**—The BI Server Edition offers the full suite of powerful BI capabilities in SQL Server 2014, including PowerPivot, Power View, Data Quality Services, and Master Data Services. One major focus of this edition is empowering end users with BI functionality.

This is ideal for projects that need advanced BI capabilities, but don't require the full OLTP performance and scalability of Enterprise Edition. The new BI Edition is inclusive of the standard edition and still contains the basic OLTP offerings.

- **Standard Edition**—This remains as it is today, designed for departmental use with limited scale. It has basic database functionality and basic BI functionality. New features are now included in the Standard Edition, such as compression.

**TABLE 1-7: SQL Server 2014 Editions**

ENTERPRISE EDITION (INCLUSIVE OF BUSINESS INTELLIGENCE EDITION)	BUSINESS INTELLIGENCE EDITION (INCLUSIVE OF STANDARD EDITION)	STANDARD EDITION
Mission-critical and Tier 1 applications	Corporate and scalable reporting of analytics	Departmental databases
Data warehousing	Power View and PowerPivot enable self-service analytics	Limited business intelligence project
Private cloud and highly virtualized environments		
Large, centralized, or external-facing business intelligence		

Note: Web Edition will only be available via Service Provider License Agreement (SPLA)

Table 1-8 provides you with an at-a-glance view of some of the key features in SQL Server 2014 by edition.

**TABLE 1-8: SQL Server 2014 Capabilities by Edition**

SQL SERVER CAPABILITIES	SQL SERVER EDITIONS		
	STANDARD	BUSINESS INTELLIGENCE	ENTERPRISE
Maximum number of cores	16 cores	16 cores (database; operating system maximum for business intelligence)	Operating system maximum
Supported Memory	128GB	128GB	Operating system maximum
Basic OLTP	X	X	X

*continues*

TABLE 1-8 (continued)

SQL SERVER CAPABILITIES	SQL SERVER EDITIONS		
	STANDARD	BUSINESS INTELLIGENCE	ENTERPRISE
In-memory OLTP			X
Basic reporting and analytics	X	X	X
Programmability and developer's tools	X	X	X
Manageability (SSMS, policy-based management)	X	X	X
Enterprise data management (data quality, master data services)		X	X
Self-service BI (Power View, PowerPivot for SPS)		X	X
Corporate BI (semantic model, advanced analytics)		X	X
Advanced security (advanced auditing, transparent data encryption)			X
Data warehousing (ColumnStore index, compression, partitioning)			X
Maximum scalability and performance			X
High availability	Limited	Basic	X
ColumnStore index			X
AlwaysOn			X
Buffer Pool Extension	X	X	X
Virtualization licensing	1 virtual machine	1 virtual machine	Unlimited

The basic functionality is kept in the Standard Edition. Key enterprise BI features are included in the Business Intelligence SKU. When you need to go to high-end data warehousing and enterprise-level high availability, the Enterprise Edition is the right edition for you.

As you can see in Table 1-8, for the Standard Edition, there is a 16-core maximum compute capacity for the Database Engine. For the Business Intelligence Edition, there is a 16-core maximum for database use and up to the OS Maximum for BI features (Analysis Services, Reporting Services). The Enterprise Edition can be used up to the maximum number of cores in the operating system.

Both the Business Intelligence Edition and Enterprise Edition offer the full premium BI capabilities of SQL Server 2014, including Enterprise Data Management, Self-Service BI, and Corporate BI features. The Enterprise Edition adds mission-critical and Tier-1 database functionality with the maximum scalability, performance, and high availability. With the Enterprise Edition under Software Assurance, customers also get unlimited virtualization, with the ability to license unlimited virtual machines, as well as full mobility of those licenses.

## Licensing

Starting with SQL Server 2012, licensing has changed significantly. This new licensing scheme can impact your environment if you do not run Software Assurance on those licenses. This section is an overview of the changes, not an exhaustive licensing discussion. See your Microsoft account manager for more details on these changes.

SQL Server pricing and licensing better aligns to the way customers purchase database and BI products. The new pricing and licensing offers a range of benefits for customers, including the following:

- **SQL Server offers market leading Total Cost of Ownership (TCO):**
  - SQL Server continues to be the clear high-value leader among major vendors. This is exemplified through the pricing model, features included in non-Enterprise editions, and world-class support, no matter the edition.
  - Customers with Software Assurance get significant benefits and ways to help ease the transition to new licensing models. These benefits include access to upgrades and enhanced support options during your license term.
  - Customers with Enterprise Agreements have the easiest transition to the latest versions of SQL Server and realize the greatest cost-savings.
- **SQL Server 2014 is cloud-optimized:**
  - SQL Server 2014 is the most virtualization-friendly database platform with expanded virtualization licensing, the flexibility to license per VM, and excellent support for Hyper-V.
  - Hybrid on-premises and cloud-based solutions are fully supported with features such as support for SQL Server Data Files in Windows Azure, SQL Server Managed Backup to Windows Azure, ability to deploy a database from an on-premises instance directly to an instance in a Windows Azure Virtual Machine, support for backing up databases to Windows Azure BLOB Storage, and ability to add a replica to your AlwaysOn Availability Group in a Windows Azure Virtual Machine.
- **SQL Server pricing and licensing is designed to enable customers to pay as they grow:**
  - The new streamlined editions are aligned to database and BI needs.
  - For data center scenarios, licensing is better aligned to hardware capacity.
  - For BI, licensing is aligned to user-based access, which is the way most customers are accustomed to purchasing BI.

## CPU Core (Not Processor) Licenses

With the release of SQL Server 2012, Microsoft switched to a per-core processing model. The CPU Core licenses (available only for the Standard and Enterprise Edition) are sold in two core “packs.” So, a quad-core CPU needs two of these packs per socket. These license packs cost half of what a SQL Server 2008 R2 CPU license costs. The catch here is that you must purchase at least four cores per CPU.

Following are some examples:

- Two sockets with 2 cores each, you need four license “packs” (eight core licenses).
- Two sockets with 4 cores each, you need four license “packs” (eight core licenses).
- Two sockets with 6 cores each, you need six license “packs” (12 core licenses).
- Two sockets with 8 cores each, you need eight license “packs” (16 core licenses).

## Virtualized SQL Server and Host-Based Licensing

When you run a virtualized SQL Server, you must license at least four cores for the VM. If you have more than four virtual CPUs on the VM, you must have a CPU Core license for each virtual CPU that you have assigned to the VM.

SQL Server 2014 still includes host-based licensing for those customers with Software Assurance and an Enterprise Agreement. The host-based licensing works just like it did before: you purchase enough Enterprise Edition CPU Core licenses for the host, and you can run as many VMs running SQL Server as you want. This will likely be the preferred way for many of you. Those customers not running with Software Assurance or an Enterprise Agreement will want to contact their Microsoft Representative or reseller, because host-based licensing is not available for those customers.

As you can see, this is a lot of change. The pricing has changed, too, but is not discussed here because of the wide variations depending on your individual agreements with Microsoft. Microsoft has tried hard to ensure that the cost does not go up dramatically for many of you, so don't fret, but be judicious and check this out with your Microsoft Account teams.

## SUMMARY

The architecture for SQL Server 2014 has advancements under the covers that will improve performance, increase developer efficiency and system availability, and decrease overall operating cost. SQL Server 2014 also delivers exciting new features that allow your system to scale and offer greater performance than ever before. It is important to focus on your current and future roles, and understand the types of features and editions that will apply to your situation and organizational needs.

In Chapter 2, you learn about best practices for upgrading to SQL Server 2014.