

## Chapter 1

### Column Names

Keep the names simple and intuitive. For more information see Chapter 11.

### Data Types

The general rule on data types is to use the smallest one you can. This conserves memory usage and disk space. Also keep in mind that SQL Server processes numbers much more efficiently than characters, so use numbers whenever practical. I have heard the argument that numbers should only be used if you plan on performing mathematical operations on the columns that contain them, but that just doesn't wash. Numbers are preferred over string data for sorting and comparison as well as mathematical computations. The exception to this rule is if the string of numbers you want to use starts with a zero. Take the social security number, for example. Other than the unfortunate fact that some social security numbers (like my daughter's) begin with a zero, the social security number would be a perfect candidate for using an integer instead of a character string. However, if you tried to store the integer 012345678 you would end up with 12345678. These two values may be numeric equivalents but the government doesn't see it that way. They are strings of numerical characters and therefore must be stored as characters rather than numbers.

When designing tables and choosing a data type for each column, try to be conservative and use the smallest, most efficient type possible. But, at the same time, carefully consider the exception, however rare, and make sure that the chosen type will always meet these requirements.





The data types available for columns in SQL Server 2000 and 2005 are specified in the following table.

Data Type	Storage	Description
Bigint	8 bytes	An 8-byte signed integer. Valid values are -9223372036854775808 through +9223372036854775807.
Int	4 bytes	A 4-byte signed integer. Valid values are -2,147,483,648 through +2,147,483,647.
SmallInt	2 bytes	A double-byte signed integer. Valid values are -32,768 through +32,767.
TinyInt	1 byte	A single-byte unsigned integer. Valid values are from 0 through 255.
Bit	1 bit	Integer data with either a 1 or 0 value.



*Table continued on following page*



Data Type	Storage	Description
 DateTime	8 bytes	Datetime is used to store dates from January 1, 1753 through December 31, 9999 (which could cause a huge Y10K disaster). The accuracy of the datetime data type is 3.33 milliseconds.
SmallDatetime	4 bytes	Smalldatetime stores dates from January 1, 1900 through June 6, 2079 with an accuracy of 1 minute.
 Char	1 byte per character. Maximum 8000 characters	The char data type is a fixed-length data type used to store character data. The number of possible characters is between 1 and 8000. The possible combinations of characters in a char data type are 256. The characters that are represented depend on what language, or collation, is defined. English, for example, is actually defined with a Latin collation. The Latin collation provides support for all English and western European characters.
 VarChar	1 byte per character. Maximum 8000 characters	The varchar data type is identical to the char data type with the exception of it being a variable length type. If a column is defined as char(8) it will consume 8 bytes of storage even if only three characters are placed in it. A varchar column only consumes the space it needs. Typically, char data types are more efficient when it comes to processing and varchar data types are more efficient for storage. The rule of thumb is: use char if the data will always be close to the defined length. Use varchar if it will vary widely. For example, a city name would be stored with varchar(167) if you wanted to allow for the longest city name in the world, which is Krung thep mahanakhon bovoratatanakosin mahintharayutthaya mahadilok popnopratratchathani burirom udomratchanivetmahasathan amornpiman avatarnsathit sakkathat-tiyavisnukarmprasit (the poetic name of Bangkok, Thailand). Use char for data that is always the same. For example, you could use char(12) to store a domestic phone number in the United States: (123)456-7890.
 Text	1 byte per character. Maximum 2,147,483,648 characters (2GB)	The text data type is similar to the varchar data type in that it is a variable-length character data type. The significant difference is the maximum length of about 2 billion characters (including spaces) and where the data is physically stored. With a varchar data type on a table column, the data is stored physically in the row with the rest of the data. With a text data type, the data is stored separately from the actual row and a pointer is stored in the row so SQL Server can find the text.