SQL Server Database Coding Standards and Guidelines
http://www.SQLAuthority.com

Naming

Tables: Rules: Pascal notation; end with an 's'
                     Examples: Products, Customers
                     Group related table names[1]

Stored Procs: Rules: sp<App Name>_[<Group Name >_]<Action><table/logical instance>
                     Examples: spOrders_GetNewOrders, spProducts_UpdateProduct

Triggers: Rules: TR_<TableName>_<action>
                     Examples: TR_Orders_UpdateProducts
                     Notes: The use of triggers is discouraged

Indexes: Rules: IX_<TableName>_<columns separated by _>
                     Examples: IX_Products_ProductID

Primary Keys: Rules: PK_<TableName>
                     Examples: PK_Products

Foreign Keys: Rules: FK_<TableName1>_<TableName2>
                     Example: FK_Products_Orderss

Defaults: Rules: DF_<TableName>_<ColumnName>
                     Example: DF_Products_Quantity

Columns: If a column references another table's column, name it <table name>ID
                     Example: The Customers table has an ID column
                             The Orders table should have a CustomerID column

General Rules: Do not use spaces in the name of database objects
                     Do not use SQL keywords as the name of database objects
                             In cases where this is necessary, surround the
                             object name with brackets, such as [Year]
                     Do not prefix stored procedures with 'sp_'[2]
                     Prefix table names with the owner name[3]

Structure

- Each table must have a primary key
  - In most cases it should be an IDENTITY column named ID
- Normalize data to third normal form
  - Do not compromise on performance to reach third normal form. Sometimes, a little de-normalization results in better performance.
- Do not use TEXT as a data type; use the maximum allowed characters of VARCHAR instead
- In VARCHAR data columns, do not default to NULL; use an empty string instead
- Columns with default values should not allow NULLs
- As much as possible, create stored procedures on the same database as the main tables they will be accessing

Formatting

- Use upper case for all SQL keywords
  - SELECT, INSERT, UPDATE, WHERE, AND, OR, LIKE, etc.
- Indent code to improve readability
- Comment code blocks that are not easily understandable
  - Use single-line comment markers(--)
  - Reserve multi-line comments (/*.. ..*/) for blocking out sections of code
- Use single quote characters to delimit strings.
  - Nest single quotes to express a single quote or apostrophe within a string
    - For example, SET @sExample = 'SQL''s Authority'
- Use parentheses to increase readability
  - WHERE (color='red' AND (size = 1 OR size = 2))
- Use BEGIN..END blocks only when multiple statements are present within a conditional code segment.
- Use one blank line to separate code sections.
- Use spaces so that expressions read like sentences.
  - fillfactor = 25, not fillfactor=25
- Format JOIN operations using indents
  - Also, use ANSI Joins instead of old style joins[4]
- Place SET statements before any executing code in the procedure.

Coding

- Optimize queries using the tools provided by SQL Server[5]
- Do not use SELECT *
- Return multiple result sets from one stored procedure to avoid trips from the application server to SQL server
- Avoid unnecessary use of temporary tables
  - Use 'Derived tables' or CTE (Common Table Expressions) wherever possible, as they perform better[6]
- Avoid using <> as a comparison operator
  - Use ID IN(1,3,4,5) instead of ID <> 2
- Use SET NOCOUNT ON at the beginning of stored procedures[7]
- Do not use cursors or application loops to do inserts[8]
  - Instead, use INSERT INTO
- Fully qualify tables and column names in JOINs
- Fully qualify all stored procedure and table references in stored procedures.
- Do not define default values for parameters.
  - If a default is needed, the front end will supply the value.
- Do not use the RECOMPILE option for stored procedures.
- Place all DECLARE statements before any other code in the procedure.
- Do not use column numbers in the ORDER BY clause.
- Do not use GOTO.
- Check the global variable @@ERROR immediately after executing a data manipulation statement (like INSERT/UPDATE/DELETE), so that you can rollback the transaction if an error occurs
  - Or use TRY/CATCH
- Do basic validations in the front-end itself during data entry
- Off-load tasks, like string manipulations, concatenations, row numbering, case conversions, type conversions etc., to the front-end applications if these operations are going to consume more CPU cycles on the database server
- Always use a column list in your INSERT statements.
  - This helps avoid problems when the table structure changes (like adding or dropping a column).

- Minimize the use of NULLs, as they often confuse front-end applications, unless the applications are coded intelligently to eliminate NULLs or convert the NULLs into some other form.
    o Any expression that deals with NULL results in a NULL output.
    o The ISNULL and COALESCE functions are helpful in dealing with NULL values.
- Do not use the identitycol or rowguidcol.
- Avoid the use of cross joins, if possible.
- When executing an UPDATE or DELETE statement, use the primary key in the WHERE condition, if possible. This reduces error possibilities.
- Avoid using TEXT or NTEXT datatypes for storing large textual data.[9]
    o Use the maximum allowed characters of VARCHAR instead
- Avoid dynamic SQL statements as much as possible.[10]
- Access tables in the same order in your stored procedures and triggers consistently.[11]
- Do not call functions repeatedly within your stored procedures, triggers, functions and batches.[12]
- Default constraints must be defined at the column level.
- Avoid wild-card characters at the beginning of a word while searching using the LIKE keyword, as these results in an index scan, which defeats the purpose of an index.
- Define all constraints, other than defaults, at the table level.
- When a result set is not needed, use syntax that does not return a result set.[13]
- Avoid rules, database level defaults that must be bound or user-defined data types. While these are legitimate database constructs, opt for constraints and column defaults to hold the database consistent for development and conversion coding.
- Constraints that apply to more than one column must be defined at the table level.
- Use the CHAR data type for a column only when the column is non-nullable.[14]
- Do not use white space in identifiers.
- The RETURN statement is meant for returning the execution status only, but not data.

Reference:

1) Group related table names:

      Products_USA
      Products_India
      Products_Mexico

2) The prefix sp_ is reserved for system stored procedures that ship with SQL Server. Whenever SQL Server encounters a procedure name starting with sp_, it first tries to locate the procedure in the master database, then it looks for any qualifiers (database, owner) provided, then it tries dbo as the owner. Time spent locating the stored procedure can be saved by avoiding the "sp_" prefix.

3) This improves readability and avoids unnecessary confusion. Microsoft SQL Server Books Online states that qualifying table names with owner names helps in execution plan reuse, further boosting performance.

4)
False code:
SELECT *
FROM Table1, Table2
WHERE Table1.d = Table2.c

True code:
SELECT *
FROM Table1
INNER JOIN Table2 ON Table1.d = Table2.c

5) Use the graphical execution plan in Query Analyzer or SHOWPLAN_TEXT or SHOWPLAN_ALL commands to analyze your queries. Make sure your queries do an "Index seek" instead of an "Index scan" or a "Table scan." A table scan or an index scan is a highly undesirable and should be avoided where possible.

6) Consider the following query to find the second highest offer price from the Items table:

```
SELECT MAX(Price)
FROM Products
WHERE ID IN
(
SELECT TOP 2 ID
FROM Products
ORDER BY Price Desc
)
```

The same query can be re-written using a derived table, as shown below, and it performs generally twice as fast as the above query:

```
SELECT MAX(Price)
FROM
(
SELECT TOP 2 Price
FROM Products
ORDER BY Price DESC
)
```

7) This suppresses messages like '(1 row(s) affected)' after executing INSERT, UPDATE, DELETE and SELECT statements. Performance is improved due to the reduction of network traffic.

8) Try to avoid server side cursors as much as possible. Always stick to a 'set-based approach' instead of a 'procedural approach' for accessing and manipulating data. Cursors can often be avoided by using SELECT statements instead. If a cursor is unavoidable, use a WHILE loop instead. For a WHILE loop to replace a cursor, however, you need a column (primary key or unique key) to identify each row uniquely.

9) You cannot directly write or update text data using the INSERT or UPDATE statements. Instead, you have to use special statements like READTEXT, WRITETEXT and UPDATETEXT. So, if you don't have to store more than 8KB of text, use the CHAR(8000) or VARCHAR(8000) datatype instead.

10) Dynamic SQL tends to be slower than static SQL, as SQL Server must generate an execution plan at runtime. IF and CASE statements come in handy to avoid dynamic SQL.

11) This helps to avoid deadlocks. Other things to keep in mind to avoid deadlocks are:

- Keep transactions as short as possible.
- Touch the minimum amount of data possible during a transaction.
- Never wait for user input in the middle of a transaction.
- Do not use higher level locking hints or restrictive isolation levels unless they are absolutely needed.

12) You might need the length of a string variable in many places of your procedure, but don't call the LEN function whenever it's needed. Instead, call the LEN function once and store the result in a variable for later use.

13) IF EXISTS (SELECT 1 FROM Products WHERE ID = 50)
   Instead Of:
   IF EXISTS (SELECT COUNT(ID) FROM Products WHERE ID = 50)

14) CHAR(100), when NULL, will consume 100 bytes, resulting in space wastage. Preferably, use VARCHAR(100) in this situation. Variable-length columns have very little processing overhead compared with fixed-length columns.