

SimbaEngine SDK

Simba SQL Engine Specification



An ODBC SQL Parsing and Execution Engine for Database Access

Simba SQL Engine is a component of SimbaEngine SDK. It provides SQL processing in data drivers and data providers for non-SQL data stores built with SimbaEngine SDK. Simba SQL Engine can also be used in custom solutions where SQL parsing or execution is required.

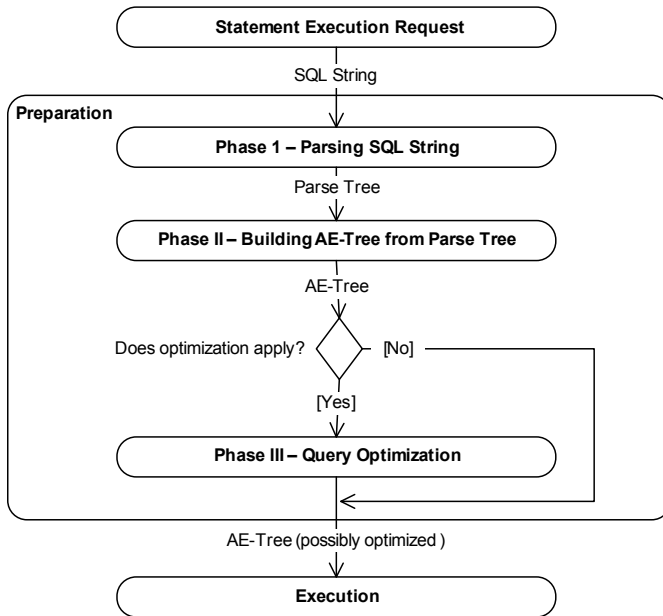
Highlights

- Build an ODBC or JDBC data driver, or an ADO.NET or OLE DB data provider for your non-SQL data.
- Quickstart sample driver has you retrieving data into Microsoft Excel in a few days.
- Make object-oriented, hierarchical or other non-tabular data available to common desktop tools such as Excel.
- Embed SQL parsing or execution into your product.
- SQL-92 compliant – targeted for the SQL generated by popular reporting and BI tools.
- Easy to use, complete, high-performance, compatible with desktop tools, portable.
- Query optimizing.
- Simba's Collaborative Query Execution feature leverages available data store query processing.
- Standard Simba Data Store Interface (DSI) API – works seamlessly with SimbaODBC and custom DSI implementations.
- SQL parsing only, or parsing and execution.
- Create your own data driver, or Simba can do the work for you.

Description

Simba SQL Engine is an easy-to-use, complete, high-performance SQL parser and execution engine for building ODBC, JDBC, ADO.NET or OLE DB data access solutions for non-SQL data stores. You can use it with other SimbaEngine SDK components to create either local or remote data driver access to non-SQL data sources – such as web services, object-oriented, network, real time (e.g. SCADA) and ISAM databases – or you can use it in a custom solution and call it from your own code. Simba SQL Engine's clean and logical interfaces make it easy to work with, and it can either parse, or parse and execute a SQL statement and retrieve the result set. Included with the SimbaEngine SDK is the Quickstart ODBC driver that is a working example of an ODBC driver built using SQL Engine that you can modify to demonstrate reading your own data in just a few days.

The Simba SQL Engine parser conforms to the SQL-92 specification, but it can be extended to implement any SQL function or feature. When Simba SQL Engine prepares a SQL statement, it validates the tables and columns referenced in the statement against the metadata from the underlying data store. Then it creates an intermediate version of the query called the Algebraic Expression Tree (AE-Tree). The AE-Tree is eventually transformed into an Execution Tree, which can be executed. By carefully moving or modifying nodes, the AE-Tree can be updated to change the way the query is executed without changing the result. The underlying data store can modify the AE-Tree, so that when the query is executed, Simba SQL Engine delegates some or all of the processing to the underlying data store. We call this Collaborative Query



Block Diagram: Simba SQL Engine Process Flow

Execution. The advantage is that any high-performance processing available in the data store is available for use with Simba SQL Engine. If the entire query can be delegated, Simba SQL Engine will only retrieve the result set. This is particularly important for analytical databases with built-in, high-performance aggregation processing that end users expect to see even through ODBC.

Simba SQL Engine is a complete, compatible, high-performance and easy to use solution for SQL processing on non-SQL data stores. Everything you need to prepare and execute SQL queries is included with a clean and logical interface. With compliance to the SQL-92 standard, it is carefully designed to support common reporting and BI tools, so end users can get on with their work. Collaborative Query Execution allows you to use the high-performance processing built into your own data store to boost the performance of Simba SQL Engine. In a few days of coding, you can begin to evaluate how well Simba SQL Engine will work for you, and in a few weeks you can have a commercial database driver ready to ship. You can use Simba SQL Engine with the other components of SimbaEngine SDK to quickly deploy local or remote data access, or you can use it alone, called from your own code. However you choose to deploy it, Simba SQL Engine is the easiest and most flexible SQL preparation and execution engine available today.

Specifications

- SQL-92-compliant parser with comprehensive support for ODBC 3.52-compliant data types, as well as scalar and aggregate functions.
- ANSI and Unicode UTF-8, -16 and -32 support for both data and metadata.
- Available in 32-bit and 64-bit versions for Windows, Linux and UNIX.
- Maximum number of simultaneous connections limited only by memory.
- Maximum result set length limited only by memory and swap table space.
- Maximum 65,356 columns supported in a select list.
- Maximum 65,356 columns supported in a table.
- Maximum 2GB in a column-value for 32-bit versions, 4 GB for 64-bit versions.
- Maximum number of characters in a SQL statement is configurable and is only limited by memory.
- Maximum number of joined tables limited only by memory and swap table space.