



SQL Anywhere 12

New Features Summary

WHITE PAPER



Contents:

Introduction.....	2
Out-of-Box Performance	3
Automatic Tuning of Server Threads	3
Column Statistics Management.....	3
Improved Remote Data Access Performance	3
Enhanced Connection Pooling for Clients and HTTP Server	3
Other Performance Enhancements	4
Scalability and Monitoring	5
Read-Only Scale-Out Configurations	5
New Design of the Monitor User Interface	5
Developer Productivity.....	7
Spatial Functionality	7
New Spatial Viewer and Spatial Preview Tab in Interactive SQL.....	7
New Spatial Data Types, Methods, Constructors, and Functions.....	8
Spatial Wizards	8
Sequence Generator.....	8
UltraLite Improvements	8
Other Developer Productivity Enhancements	9
Advanced Data Synchronization	10
Performance	10
Single Administration Point	10
New Graphical Design for MobiLink 12 Administration	10
Server Farms Enhancements	10
Security	11
Additional Resources.....	12

Introduction

SQL Anywhere is the leading data management and enterprise synchronization solution for applications that operate outside the traditional data center. From its inception, SQL Anywhere was designed to deliver enterprise caliber features, high performance out of the box, and robust synchronization enabling the implementation of mission-critical database solutions for networked, embedded, and mobile environments.

Under the tenets of ease-of-use, self-management, and embeddability, SQL Anywhere 12, the latest release, continues to advance those offerings by providing major new features in the areas of out-of-box performance, scalability and monitoring, developer productivity, and advanced data synchronization. This document summarizes the new features of SQL Anywhere 12 under these four areas of focus, as well as new features added to MobiLink and UltraLite technologies. Obtain more information about SQL Anywhere and the new features and functionality of SQL Anywhere 12 by visiting www.sybase.com/sqlanywhere.

Out-of-Box Performance

The following added features to SQL Anywhere 12 enhance database performance, leading to faster queries and efficient use of resources.

Automatic Tuning of Server Threads

SQL Anywhere 12 is able to use multiple threads to execute a number of active tasks at a given time. An active task is one that is currently being executed by a thread (or fiber) in the database server. For example, an active task may be executing an access plan operator, or performing some other useful work. An active task may also be blocked, waiting for a resource (such as an I/O operation, or a lock on a row). An unscheduled task is one that is ready to execute, but is waiting for an available thread or fiber. The number of active tasks that can execute simultaneously depends on the number of database server threads and the number of logical processors in use on the computer.

The maximum number of threads, referred to as the multiprogramming level (MPL), can greatly affect query performance. Unfortunately, it can be difficult to determine the optimum MPL without experimentation of the application's workload. The database server thread auto-tuning feature of SQL Anywhere 12 simplifies this process by dynamically adding or removing worker tasks as required, thus improving its throughput and intelligently adapting to changes in the workload environment.

Column Statistics Management

To increase query execution and enhance performance, the database server is now able to gather column statistics to detect any inaccuracies and automatically fix them. This is accomplished by the optimizer and is based on column statistics stored in the database and on heuristics to choose the best optimization strategy.

For each access plan considered by the optimizer, an estimated result size (number of rows) is computed. For example, for each join method or index access based on the selectivity estimations of the predicates used in the query, an estimated result size is calculated. The estimated result sizes are used to compute the estimated disk access and CPU cost for each operator such as a join method, a group by method, or a sequential scan, used in the plan.

Column statistics are the primary data used by the optimizer to compute selectivity estimation of predicates. Therefore, they are vital to estimating correctly the cost of an access plan. Accurately analysing column statistics helps maintain efficient database performance.

Improved Remote Data Access Performance

A proxy table is a local table containing metadata used to access a table on a remote database server as if it were a local table. Proxy tables allow the user to import data directly from a variety of relational and non-relational data sources, as well as querying data across databases.

SQL Anywhere 12 improves remote data access operations by presenting a number of enhancements to proxy table performance, most notably an increase in the buffer size to hold more data in memory, as well as improved conversion of data types to make better utilization of computing resources.

Enhanced Connection Pooling for Clients and HTTP Server

Connection pooling can improve the performance of applications that make multiple, brief connections to the database server. When an application first attempts to connect to the database, it checks a pool for an existing

connection that uses the same connection parameters the user has specified. If a matching connection is found, that connection is used. Otherwise, a new connection is used. When you disconnect, the connection is returned to the pool so that it can be reused when that application reconnects.

The new ConnectionPool (CPOOL) connection parameter controls the behavior of client connection pooling. Once an application makes a specified number of connections with the same connection string, the connection is pooled. By default, an application must make five connections with the same connection string before a connection is cached. The connection name can be unique each time, but all other connection parameters must be identical for a cached connection to be reused. The user can set the maximum number of cached connections from each application. If the application process connects again and there are cached connections available for the same connection string, the cached connection is reused. Connections remain in the cached state for the time specified by the CPOOL connection parameter (60 seconds by default).

Enhancements to the HTTP server's automatic connection pooling allows it to handle a much larger amount of HTTP traffic and serve web services with a higher level of concurrency. Several new options are provided that allow users to specify the normal threshold size of database connections, specify the maximum duration that an unused connection can be retained in the connection pool, and intelligently remove any unused database connections.

Other Performance Enhancements

SQL Anywhere 12 includes several more performance enhancements. Immediate materialized views now support outer joins, enabling faster query execution for applications that analyze large amounts of data. There's also a number of new database options to improve performance, for example the ability to prevent a low priority task from blocking other connections beyond a specified time. Further, SQL Anywhere 12 highlights improvements in index performance, validation of large databases, and request prioritization. Developers using object-relational mapping (ORM) technologies, such as the Microsoft ADO.NET Entity Framework, will experience better performance as SQL Anywhere 12 enhances the optimization of queries generated through ORMs.



Scalability and Monitoring

Significant enhancements made to SQL Anywhere 12 improve data availability for systems dealing with large volumes of information. Additional capabilities to monitor the health and performance of SQL Anywhere databases and MobiLink servers are also new to this release.

Read-Only Scale-Out Configurations

Read-only scale-out is a configuration that allows offloading of reporting or other operations that require read-only access to the database. SQL Anywhere 12 takes advantage of read-only scale-out capabilities leading to applications gaining scalability as the load is split across multiple database servers.

A scale-out tree consists of the root node and copy nodes. The root node of a scale-out system has the only copy of the database that is writable. The root node can be either a single database, or a database mirroring system. Below the root node, there can be one or many branches of copy nodes. A copy node is a database server that runs a copy of a database that supports read-only access, and a copy node can be the parent of other copy nodes.

The root node is the only database server that accepts both read and write requests. Once you start additional copy nodes, the root database server starts sending transaction log pages to the copy nodes in the tree. The root database server also sends the pages to any other nodes that are defined as its children if they are connected and ready to receive log pages. The pages are normally sent without waiting for a response; however, the root database server occasionally requests an acknowledgement to ensure that the copy node does not receive more asynchronous requests than it can handle.

When the copy node receives pages, it writes them to disk and then sends them to its children (if it has any). The parent detects if a child node becomes unavailable, and if this happens, the parent stops pushing transaction log pages to the child. If the child is restarted, it requests the transaction log pages that it does not have, and then the parent resumes pushing transaction log pages to the child.

Database Mirroring Implementations

SQL Anywhere 12 allows access to a database running on a mirror server using a read-only connection. Database mirroring with read-only scale-out ensures the availability of the root node. With this method, the root node of the scale-out system consists logically of the primary, mirror, and arbiter servers instead of a single database server. Once you start the database servers in the mirroring system, additional database servers can access the participating mirror servers to maintain read-only copies of the database.

When copy nodes are used in a database mirroring system, they can be the child of the current primary server, the current mirror server, or another copy node. The copy node gets its transaction log pages from its parent. Unlike mirror nodes, copy nodes do not have a state information file because their state does not influence which database server has ownership of the database.

Applying read-only scale-out configurations to database mirroring enhances new or existing high availability environments.

New Design of the Monitor User Interface

The Monitor collects metrics and performance data from SQL Anywhere databases and MobiLink servers running on other computers, while a separate computer accesses the Monitor via a web browser. This tool allows administrators to monitor a large number of servers and act on alerts.

The following is a list of additions to the SQL Anywhere Monitor introduced in version 12.0.0:

- **New user interface-** The Monitor user interface is dashboard-based. The dashboards contain widgets to display metrics, alerts, and resource information. A user can create new widgets, maximize and minimize them as well as move them around the dashboard.
- **Monitor MobiLink server farms and Relay Server farms-** A user can use the Monitor to monitor MobiLink server farms and Relay Server Farms as well as SQL Anywhere databases and MobiLink servers.
- **Close connections from within the Monitor-** From within the Monitor, a user can close a resource database's connections.
- **Import SQL Anywhere resources to be monitored-** Previously, a user could only add one SQL Anywhere resource at a time to the Monitor. Now you can create a list of SQL Anywhere resources in a .csv file and import this list into the Monitor.
- **Perform on-demand maintenance-** Administrators can perform unscheduled maintenance on the Monitor.
- **New backup alert for SQL Anywhere resources-** Administrators can configure the Monitor to issue an alert when a SQL Anywhere resource has not been successfully backed up for a given number of days.
- **Time metrics-** The times reported in the Monitor are always local to the browser you are using.
- **Export metrics-** Export metrics that have a graph or table associated with them, to an xml file. For example, most of the metrics in the **Key Performance Metrics** widget can be exported.
- **Troubleshooting features-** Administrators can monitor the Monitor via the **Message Log**, and **Exception Reports**.

By offering insight into the health of mission-critical servers, the server monitor greatly enhances administrative tasks.



Developer Productivity

A major new feature of SQL Anywhere 12 is support for spatial data for both SQL Anywhere and UltraLite databases, as well as MobiLink synchronization technology. Other developer productivity new features include support for sequences in SQL Anywhere and enhancements to UltraLite databases.

Spatial Functionality

Spatial data is data that describes the position, shape, and orientation of objects in a defined space. SQL Anywhere 12 allows storage and querying of spatial data. Spatial data in SQL Anywhere is represented as 2D geometries in the form of points, curves (line strings and strings of circular arcs), and polygons.

Spatial support in SQL Anywhere lets application developers associate spatial information with their data. For example, a table representing stores could store the location of the store as a point, or store the delivery area for the store as a polygon. This can be represented in SQL as:

```
CREATE TABLE Locations(  
    ID INT,  
    ManagerName CHAR(16),  
    StoreName CHAR(16),  
    Address ST_Point,  
    DeliveryArea ST_Polygon )
```

The spatial type `ST_Point` in the example represents a single point, and `ST_Polygon` represents an arbitrary polygon. With this schema, the application can show all store locations on a map, or find out if a store delivers to a particular address using a query similar to the following:

```
CREATE VARIABLE @pt ST_Point;  
SET @pt = ST_Geometry::ST_GeomFromText( 'POINT(1 1)' );  
  
SELECT * FROM Locations  
WHERE DeliveryArea.ST_Contains( @pt ) = 1
```

Spatial data is extremely useful in many situations, such as efficient workforce dispatching and location-based inventory tracking. Application developers will be able to take advantage of the database server's new spatial data capabilities to implement state-of-the-art geospatial solutions or enhance existing applications by adding support for geographic information.

New Spatial Viewer and Spatial Preview Tab in Interactive SQL

A new viewer tool, the **Spatial Viewer**, has been added to Interactive SQL to allow the user to view spatial geometries. Spatial data can be queried in the top portion of the viewer, and then results are represented as an image in the lower portion of the viewer.

When viewing a result row in Interactive SQL, the user can now preview a geometry as a Scalable Vector Graphic (SVG) using the new **Spatial Preview** tab.

New Spatial Data Types, Methods, Constructors, and Functions

New types, methods, and constructors have been added to allow the user to access, model, and analyze spatial data. As well, many spatial compatibility functions have been provided to mimic regular SQL functions when accessing and manipulating spatial data. These functions have been provided for compatibility with other products, and make use of the spatial methods and constructors provided in SQL Anywhere.

Spatial Wizards

Sybase Central presents two new wizards in support of the spatial data feature. **The Create Spatial Reference System** wizard allows developers to create new spatial reference systems. **The Create Unit Of Measure** wizard allows developers to create new units of measure for use with spatial data.

Sequence Generator

SQL Anywhere now supports the generation of unique sequence values by a sequence generator. It is stored as a database object, and its value can be referred to anywhere that an expression can be used. Sequences can be used in applications to prevent concurrency and performance issues caused by the generation of unique key values outside the database. A sequence can generate values that are unique across multiple tables or that are different from a set of natural numbers.

UltraLite Improvements

UltraLite is a compact relational database with many of the features and functionality of SQL Anywhere. UltraLite can be installed as part of a SQL Anywhere solution for enterprise-wide mobile data management, or as part of a standalone embedded solution.

UltraLite provides built-in functionality to mobilize corporate data. When deployed as a MobiLink client, a user can implement a synchronization solution that delivers mission-critical information in a timely and reliable way. Users can record and access the data they require—even when they cannot directly access the corporate network.

UltraLite is now able to operate on the Apple iPhone, providing developers with a wider choice of devices to deploy their mobile solutions. UltraLite application development is fully supported for the Xcode suite of tools running on the Mac OS X operating system.

Just like SQL Anywhere, UltraLite also supports the storage and processing of spatial data. Coupled with MobiLink technology, you can synchronize spatial data from enterprise repositories to an UltraLite database running on a smartphone.

New security enhancements include support for FIPS 140-2 certified encryption on 64-bit Windows. Further UltraLite database encryption is now performed using 256-bit AES instead of 128-bit AES.

The ulinit database creation utility has been enhanced to allow you to create an UltraLite database based on information in a SQL Anywhere database even when the schema being extracted from the SQL Anywhere database contains elements that UltraLite does not support (such as column data types or default values for instance). This improvement allows for additional flexibility when designing and developing UltraLite applications.

A new version of the UltraLite C/C++ API has been added in this release. It offers incremental design and performance improvements, especially when using the UltraLite engine.



Other Developer Productivity Enhancements

As with previous versions, SQL Anywhere continues to improve its support for technologies traditionally used by application programmers. To ease migration-related issues for Oracle and MySQL developers, SQL Anywhere 12 offers support for sequences, as well as creating or replacing objects in the database with a single statement, thus reducing lines of code and improving development timeframes. SQL Anywhere 12 also provides support for the latest Microsoft Visual Studio and .NET Framework technologies. Application developers will be able to immediately implement SQL Anywhere database applications with Visual Studio 2010 and .NET Framework 4.0, including the ADO.NET Entity Framework.

Advanced Data Synchronization

MobiLink enhancements include improved performance, easier synchronization of local and remote databases, and better server farm support.

Performance

MobiLink now supports dynamic memory caching. With this feature, there is increased input to the memory cache so a larger cache may be needed to prevent swapping data to disk. The MobiLink server automatically grows its cache up to 60% of the available process address space and shrinks its cache if other processes on the system require more memory.

SQL Anywhere demonstrates enhanced behavior and robustness in low-memory conditions. The database engine can now drop the MobiLink client (dbmlsync) connection to the remote database (and rollback uncommitted operations dbmlsync has) if another connection is waiting for access to any database resource that dbmlsync has locked. This allows the other connection to go forward without waiting for the synchronization to complete.

Single Administration Point

SQL Anywhere now allows an administrator to manage remote databases involved in MobiLink synchronization. System administrators in charge of data synchronization are now able to manage the entire system using a single administration point, leading to simplification, automation and greater control of the data exchange environment.

Central administration of remote databases can be used to perform the following:

- Centrally control when an application database synchronizes with MobiLink.
- Push schema changes to application databases.
- Diagnose problems with specific application databases or with the synchronization system in general.
- Upload log files.

By controlling and maintaining local database schemas in the consolidated database, IT staff will experience performance gains in synchronizations. Additionally, the central point of administration enables administrators to quickly identify and resolve issues encountered in the synchronization environment.

New Graphical Design for MobiLink 12 Administration

The MobiLink plug-in has been redesigned in version 12 to support central administration of remote databases. The two MobiLink modes, Model and Admin, have been combined in the new plug-in. Now, the user can use the MobiLink plug-in to create synchronization projects that contain consolidated databases, groups, synchronization models, remote tasks, and network shares. Old synchronization models can be imported into synchronization projects. There is a new Relay Server plug-in for Sybase Central that enables you to configure backend farms and servers.

Server Farms Enhancements

The MobiLink arbiter ensures that only a single MobiLink server in a server farm is running as the primary server. New remote ID locking logic is used to prevent redundant synchronizations from the same remote ID, simplifying server farm configurations as the MobiLink server no longer requires to be explicitly started as running in a server farm.



Security

MobiLink server and clients now support 256-bit AES cipher suites for both RSA and ECC. Also, support has been added for the RFC 4492 version of the ECC cipher suites.

Additional Resources

For more information about SQL Anywhere 12, please refer to www.sybase.com/sqlanywhere.
The following resources are also available:

- Online documentation with comments:
http://dcx.sybase.com/index.html#SA12_bd_en/contents.htm
- SQL Anywhere Developer Community:
<http://www.sybase.com/developer/library/sql-anywhere-techcorner>
- SQL Anywhere Blog Center:
<http://www.sybase.com/sqlanyblogs>
- iAnywhere CodeXchange – Samples and Utilities:
<http://www.sybase.com/detail?id=1058600>



SYBASE, INC.
WORLDWIDE HEADQUARTERS
ONE SYBASE DRIVE
DUBLIN, CA 94568-7902 USA
Tel: 1 800 8 SYBASE

www.sybase.com

Copyright © 2010 Sybase, Inc. All rights reserved. Unpublished rights reserved under U.S. copyright laws. Sybase, and the Sybase logo are trademarks of Sybase, Inc. or its subsidiaries. All other trademarks are the property of their respective owners. ® indicates registration in the United States. Specifications are subject to change without notice. 3/09.