



PUBLIC

SQL Anywhere

Document Version: 17.01.0 – 2021-10-15

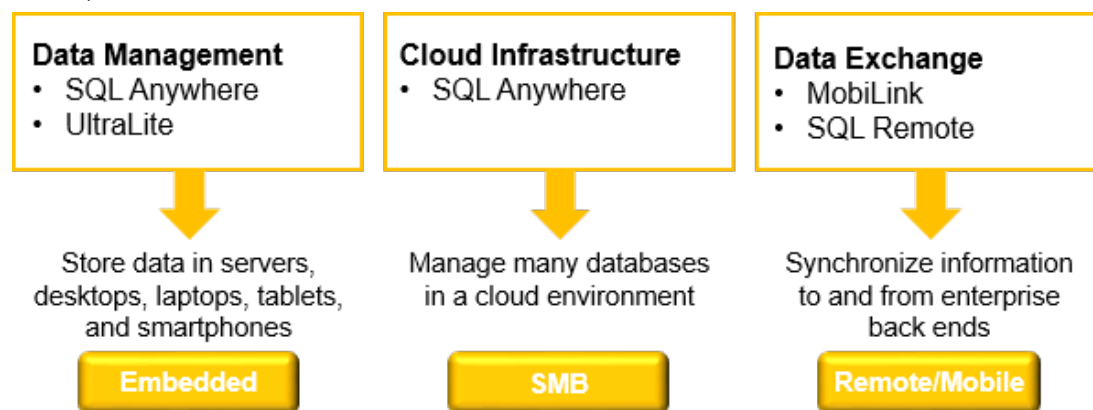
SQL Anywhere Introduction

Content

- 1 SQL Anywhere - Introduction. 3**
- 1.1 SQL Anywhere 17 Outside the Data Center. 5
- 1.2 Database Scenarios. 6
 - Desktop Applications and Embedded Databases. 6
 - Client/Server Applications. 7
 - Web Applications. 8
 - Remote Data Synchronization. 8
- 1.3 Editions and Licensing. 9
- 1.4 Separately Licensed Components. 10
- 1.5 Supported Platforms. 10
- 1.6 Accessibility. 11
- 1.7 Data Management Technologies. 17
 - SQL Anywhere and UltraLite Target Platforms. 18
 - Multi-Tier Computing Architecture. 18
 - ETL Features. 19
- 1.8 Overview of Data Exchange Technologies. 20
 - Challenges for Synchronization Technologies. 22
 - How to Choose a Synchronization Technology. 23
 - Consolidated and Remote Databases. 25
 - Hierarchical Database Configurations. 26
- 1.9 Design and Management Tools. 27
- 1.10 Sample Databases. 28
 - The SQL Anywhere Sample Database (demo.db). 29
 - The CustDB Sample Database Application. 32
- 1.11 Using the Documentation. 33
 - SAP Help Portal. 33
 - DocCommentXchange. 34
 - HTML Help. 34
 - Documentation Conventions. 35
 - Online Resources and Additional Help. 38

1 SQL Anywhere - Introduction

This book describes *SQL Anywhere 17*, a comprehensive package that provides data management and data exchange technologies, that enables the rapid development of database-powered applications for server, desktop, mobile, and remote office environments.



SQL Anywhere offers:

Embellability

SQL Anywhere can be easily embedded inside other applications. It combines high performance with a very small memory footprint. SQL Anywhere contains a range of features to enable self-management and maintenance in frontline environments, including features that enable optimization of computer resources, self-tuning for improved performance, and simplification of remote installation and support.

Interoperability

SQL Anywhere is available on many platforms, including Windows, Linux, Oracle Solaris, HP-UX, IBM AIX, and Apple macOS. Its database files can be copied between platforms. In addition, SQL Anywhere provides support for many smartphones using its UltraLite database technology for small devices. SQL Anywhere includes support for many common database interfaces, including ODBC, JDBC, ADO.NET, JavaScript, PHP, and Perl. Many popular application development tools can be used, including: Microsoft Visual Studio, SAP PowerBuilder, Eclipse, and various web tools. Stored procedures can be written in C/C++, Java, .NET, or Perl.

Performance out of the box

SQL Anywhere is designed to deliver outstanding performance without ongoing tuning and administration. Features such as dynamic cache sizing, automatic generation of statistics, a sophisticated query optimizer, parallel query processing, and materialized views mean that SQL Anywhere is ideal for environments that demand high performance but have no on-site database administrator. By offering, On-Line Analytical Processing (OLAP) SQL Anywhere offers the ability to perform complex data analysis within a single SQL statement, increasing the value of the results, while improving performance by decreasing the number of database queries.

Mobility

SQL Anywhere provides enterprise-caliber databases that operate on frontline systems and devices whether connectivity with enterprise systems is available or not. Its synchronization technologies ensure

data can be exchanged efficiently over wireless and wired networks with backend databases, application servers, and messaging systems.

Data encryption and security

SQL Anywhere provides full end-to-end data security with 128-bit strong encryption of database tables, files, and communications streams between the application and the database, and the MobiLink synchronization stream. SQL Anywhere can audit data access, offers built-in user authentication, and can integrate with third-party authentication systems. SQL Anywhere also offers a FIPS-certified version of encryption via a separately licensed component.

In this section:

[SQL Anywhere 17 Outside the Data Center \[page 5\]](#)

SQL Anywhere technologies are used in several types of applications.

[Database Scenarios \[page 6\]](#)

There are several scenarios you can choose from for building database applications and systems.

[Editions and Licensing \[page 9\]](#)

SQL Anywhere offers various editions that include separately licensed components, and have some restrictions on the number of CPUs used by the database server.

[Separately Licensed Components \[page 10\]](#)

There are components that are licensed separately and may need to be ordered separately if not included in your edition of SQL Anywhere.

[Supported Platforms \[page 10\]](#)

The *SQL Anywhere Supported Platforms and Engineering Support Status* web page lists the supported operating system platforms broken down by version of SQL Anywhere.

[Accessibility \[page 11\]](#)

SQL Anywhere supports Section 508 of the United States Federal Rehabilitation Act of 1973 (29 U.S.C. 794d) and meets its accessibility requirements when the product is run on the Windows x86 and x64 operating system platforms.

[Data Management Technologies \[page 17\]](#)

SQL Anywhere offers two relational databases: SQL Anywhere Server and UltraLite.

[Overview of Data Exchange Technologies \[page 20\]](#)

SQL Anywhere offers a wide range of options for exchanging data with existing enterprise systems and mobile devices.

[Design and Management Tools \[page 27\]](#)

There are many design and management tools for administering, monitoring, and troubleshooting databases.

[Sample Databases \[page 28\]](#)

Two sample databases are provided with the software.

[Using the Documentation \[page 33\]](#)

This Help assists you with your day-to-day tasks and includes detailed product information and procedures.

Related Information

[SQL Anywhere - Changes and Upgrading](#)

[Overview of Data Exchange Technologies \[page 20\]](#)

[Design and Management Tools \[page 27\]](#)

[Installing SQL Anywhere \(UNIX/Linux\)](#)

[Installing SQL Anywhere \(Microsoft Windows\)](#)

[Tutorial: Connecting to the Sample Database](#)

1.1 SQL Anywhere 17 Outside the Data Center

SQL Anywhere technologies are used in several types of applications.

Client-server applications

Whether it is 5, 50, 500 users or more, SQL Anywhere is a powerful database solution for server applications that provides high performance out of the box, with low maintenance and cost.

SQL Anywhere easily scales to support hundreds of active users, hundreds of gigabytes of data, and hundreds of millions of rows. Yet many ease-of-use and administration features ensure that costs stay down as performance scales up.

This deployment model works best when the majority of users are connected to the network.

Desktop applications

SQL Anywhere delivers enterprise-caliber features, without the bulky characteristics of an enterprise database. Its robust reliability and performance, along with highly efficient usage of memory and system resources, ensure that the database can be hidden from laptop and desktop users.

Organizations embed SQL Anywhere databases in their applications because SQL Anywhere databases are built for use in widely deployed, minimally administered environments, and require minimal memory and disk space.

Remote office applications

SQL Anywhere data exchange architectures address the challenges of managing and sending data within and between offices and workers that are dispersed geographically.

Companies choose SQL Anywhere database and data exchange technologies to provide remote workers with the data they require to run their operations effectively, while providing the central office with the critical information that gives the pulse of the business.

Mobile and wireless applications

Recognized as the industry's leading mobile database, SQL Anywhere gives mobile workers the ability to access their data and corporate applications.

Companies depend on SQL Anywhere for reliable management of data and mobile applications, which run on laptops, handheld devices, and smartphones.

Related Information

[Client/Server Applications \[page 7\]](#)

[Desktop Applications and Embedded Databases \[page 6\]](#)

[Consolidated and Remote Databases \[page 25\]](#)

[UltraLite Overview](#)

1.2 Database Scenarios

There are several scenarios you can choose from for building database applications and systems.

In this section:

[Desktop Applications and Embedded Databases \[page 6\]](#)

You can build a complete application and database on a single computer.

[Client/Server Applications \[page 7\]](#)

You can build an installation with many applications running on different computers, connected over a network to a single network database server running on a separate computer.

[Web Applications \[page 8\]](#)

You can build an installation that uses a SQL Anywhere web server or a third party web server to send data to clients inside or outside the corporate network.

[Remote Data Synchronization \[page 8\]](#)

You can build an installation that can scale to tens of thousands of clients with bi-directional data replication to heterogeneous database systems.

1.2.1 Desktop Applications and Embedded Databases

You can build a complete application and database on a single computer.

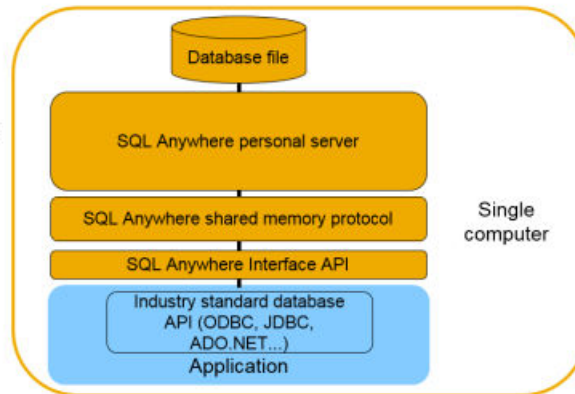
In the simplest arrangement, this is a **standalone application** or **personal application**: it is self-contained with no connection to other database. The database server and the database can be started by the client application, and it is common to refer to the database as an **embedded database**. As far as the end user is concerned, the database is a part of the application.

Many relational database management systems require experienced staff for administration. A characteristic of SQL Anywhere databases is the ability to run entirely without administration.

The SQL Anywhere personal database server is generally used for embedded applications. Embedded applications have the following architecture, with a client application connecting through a programming interface to a database server running on the same computer:

Desktop/laptop embedded database

- Personal database server running on the client



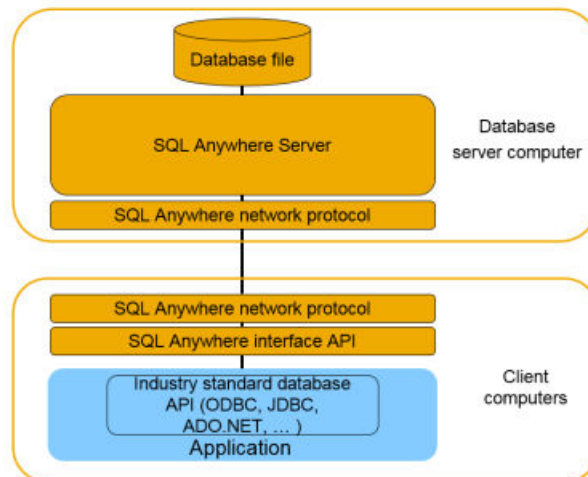
1.2.2 Client/Server Applications

You can build an installation with many applications running on different computers, connected over a network to a single network database server running on a separate computer.

This is a **client/server** or **multi-user database** environment, and has the following architecture. The interface library is located on each client computer.

Server database

- Serve data to clients inside or outside the corporate network
- Thick client application



In this case, the database server is the SQL Anywhere network database server, which supports network communications over TCP/IP.

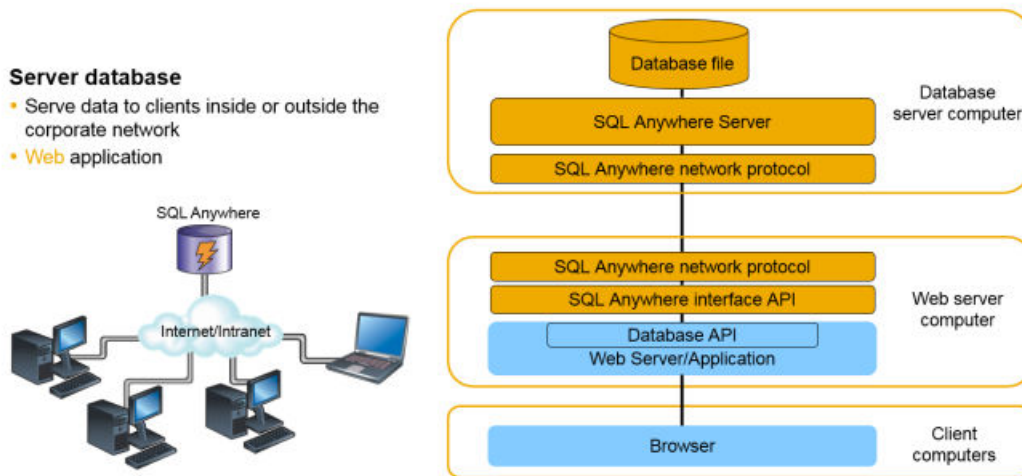
For a client application to work in a client/server environment, specify additional connection parameters, typically the HOST connection parameter and optionally the ServerName connection parameter.

Related Information

[Types of Deployment](#)
[Database Connections](#)

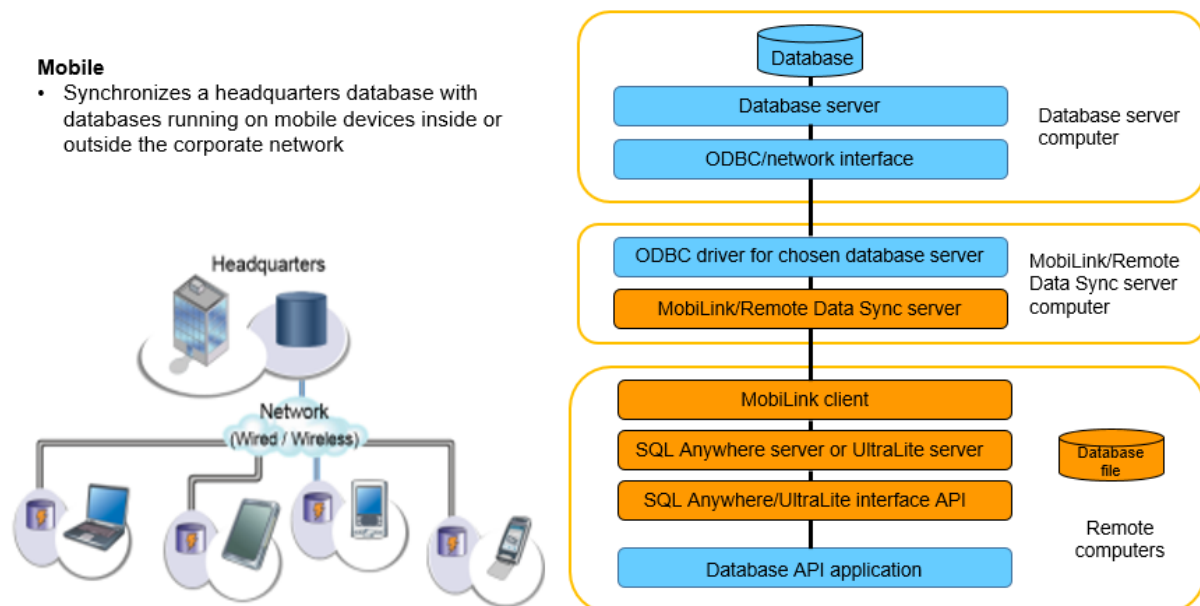
1.2.3 Web Applications

You can build an installation that uses a SQL Anywhere web server or a third party web server to send data to clients inside or outside the corporate network.



1.2.4 Remote Data Synchronization

You can build an installation that can scale to tens of thousands of clients with bi-directional data replication to heterogeneous database systems.



1.3 Editions and Licensing

SQL Anywhere offers various editions that include separately licensed components, and have some restrictions on the number of CPUs used by the database server.

Licensing and CPUs

Per-seat licensing

A per-seat license restricts the number of client connections to the database server. With per-seat licensing, the network database server can use all CPUs available on your computer, but not more than the maximum allowed by the SQL Anywhere edition you are running. If a CPU contains multiple cores or threads, all cores and threads of the CPU may be used.

The personal database server is limited to four cores on one CPU. If a core contains multiple threads, all threads of that core may be used. The personal server is limited to 10 connections.

Core-based licensing

With core-based licensing, the network database server uses up to the number of cores specified in your license, but not more than the maximum allowed by the SQL Anywhere edition you are running.

When you have a core license, there are no restrictions on the number of client connections to the network database server.

The following options further restrict the processors that the database server uses:

- -gt database server option
- -gta database server option
- -gtc database server option
- sa_server_option system procedure with the ProcessorAffinity option

Related Information

[Separately Licensed Components \[page 10\]](#)

[SAP SQL Anywhere Editions](#) 

[Server Licensing Utility \(dblic\)](#)

[-gt Database Server Option](#)

[-gta Database Server Option](#)

[-gtc Database Server Option](#)

[sa_server_option System Procedure](#)

1.4 Separately Licensed Components

There are components that are licensed separately and may need to be ordered separately if not included in your edition of SQL Anywhere.

For a list of separately licensed components, see [SAP SQL Anywhere Editions](#).

1.5 Supported Platforms

The [SQL Anywhere Supported Platforms and Engineering Support Status](#) web page lists the supported operating system platforms broken down by version of SQL Anywhere.

Web Page Name	URL	Description
SQL Anywhere Components by Platforms	http://scn.sap.com/docs/DOC-63186	Provides a list of components that are available across each platform supported by SQL Anywhere.
SAP SQL Anywhere	http://scn.sap.com/docs/DOC-35654	Provides the engineering support status for each SQL Anywhere version.
SQL Anywhere Supported Linux Platforms	http://scn.sap.com/docs/DOC-35851	Provides information about the Linux platform support, including support exceptions, key components, and tested Linux distributions.
SQL Anywhere Supported Kerberos Clients	http://scn.sap.com/docs/DOC-63443	Provides a list of Kerberos clients/runtimes that have been tested with SQL Anywhere.
ODBC Drivers for MobiLink	http://scn.sap.com/docs/DOC-63337	Provides a list of the recommended ODBC drivers for different versions of the MobiLink server.
SQL Anywhere Client Interfaces	http://scn.sap.com/docs/DOC-52993	Provides a list of supported client interfaces for SQL Anywhere.

Support for Virtualized Environments

SQL Anywhere is supported on virtual platforms such as VMware, Xen, and Microsoft Virtual Server/Windows Server Virtualization with some restrictions:

1. The virtual machine environments must be officially certified and approved by the corresponding operating system platform vendor.
2. The operating system running in the virtual machine is already officially certified by SAP.

If an issue cannot be reproduced, then the customer may be asked to reproduce the problem running natively on the certified operating system, without the use of such virtualization systems.

3. The extent of support does not cover performance. SAP does not guarantee the performance of SQL Anywhere products running on virtual machines.

Hardware Requirements

For information about hardware requirements for SQL Anywhere, see <http://scn.sap.com/docs/DOC-35653>



Related Information

[Software Updates](#)

[Physical Limitations on Size and Number of Databases](#)

[UltraLite Database Limitations](#)

1.6 Accessibility

SQL Anywhere supports Section 508 of the United States Federal Rehabilitation Act of 1973 (29 U.S.C. 794d) and meets its accessibility requirements when the product is run on the Windows x86 and x64 operating system platforms.

The tables below identify how SQL Anywhere meets the accessibility requirements identified in Section 508 of the US Federal Rehabilitation Act. An accessibility enablement component provides software that enables the use of accessibility tools.

For the purposes of this document, the following are components of SQL Anywhere:

- SQL Anywhere Server
- SQL Central
- Interactive SQL
- Online documentation

Third-Party Tools

Testing environment: SQL Anywhere 17.0.11 (64-bit Windows) using both the Microsoft Accessibility technologies and the JAWS Screen Reader from Freedom Scientific, Inc.

SQL Anywhere 17.0.11 Voluntary Product Accessibility Template Summary

Criteria	Supporting Features	Remarks and Explanations
Section 1194.21 Software Applications and Operating Systems	Supports all features	Features and capabilities may also be complemented with Microsoft Windows Accessibility Features and Aids.
Section 1194.22 Web-based Internet Information and Applications	Not Applicable	
Section 1194.23 Telecommunications Products	Not Applicable	
Section 1194.24 Video and Multi-media Products	Not Applicable	
Section 1194.25 Self-Contained, Closed Products	Not Applicable	
Section 1194.26 Desktop and Portable Computers	Not Applicable	
Section 1194.31 Functional Performance Criteria	Supports with exceptions	Supports third-party accessibility products and technologies supported on Microsoft Windows. Some third-party products may not be fully compatible.
Section 1194.41 Information, Documentation and Support	Supports	

Section 1194.21 Software Applications and Operating Systems

Criteria	Supporting Features	Remarks and Explanations
(a) When software is designed to run on a system that has a keyboard, product functions shall be executable from a keyboard where the function itself or the result of performing a function can be discerned textually.	Supports Keyboard access is provided throughout SQL Anywhere. Keyboard shortcuts, shortcut keys, and menu commands are readily available in SQL Anywhere.	SQL Anywhere fully supports Microsoft Windows accessibility features.

Criteria	Supporting Features	Remarks and Explanations
<p>(b) Applications shall not disrupt or disable activated features of other products that are identified as accessibility features, where those features are developed and documented according to industry standards. Applications also shall not disrupt or disable activated features of any operating system that are identified as accessibility features where the application programming interface for those accessibility features has been documented by the manufacturer of the operating system and is available to the product developer.</p>	<p>Supports</p> <p>SQL Anywhere does not disrupt or disable any accessibility features of the operating system.</p> <p>SQL Anywhere supports large fonts in the title bar, toolbars, as well as operating system accessibility features including StickyKeys, FilterKeys, MouseKeys, and High Contrast.</p>	<p>Microsoft Windows accessibility aids can be used with the software. Screen readers can be used with the software; however, Java-based components require screen readers that support Java Assistive Technology.</p> <p>The Microsoft Windows magnifier application is available to magnify window text.</p> <p>Windows Screen Magnifier may not always follow focus when using Tab or Shortcut Keys.</p>
<p>(c) A well-defined on-screen indication of the current focus shall be provided that moves among interactive interface elements as the input focus changes. The focus shall be programmatically exposed so that Assistive Technology can track focus and focus changes.</p>	<p>Supports</p> <p>Visual on-screen focus, highlighting and tracking is provided throughout SQL Anywhere.</p>	<p>For the Java-based components in SQL Anywhere, third-party screen reader used must support Java assistive technology.</p> <p>Windows Screen Magnifier may not always follow focus for Java-based components when using Tab or Shortcut Keys.</p>
<p>(d) Sufficient information about a user interface element including the identity, operation and state of the element shall be available to Assistive Technology. When an image represents a program element, the information conveyed by the image must also be available in text.</p>	<p>Supports</p> <p>SQL Anywhere supports both the Microsoft Active Accessibility and Java Accessibility programming interface which enable its components to be compatible with assistive technology.</p>	<p>For the Java-based components in SQL Anywhere, third-party screen reader used must support Java assistive technology.</p>
<p>(e) When bitmap images are used to identify controls, status indicators, or other programmatic elements, the meaning assigned to those images shall be consistent throughout an application's performance.</p>	<p>Supports</p> <p>SQL Anywhere uses standard and consistent images throughout the SQL Anywhere suite of applications.</p>	
<p>(f) Textual information shall be provided through operating system functions for displaying text. The minimum information that shall be made available is text content, text input caret location, and text attributes.</p>	<p>Supports</p> <p>SQL Anywhere uses standard system functions for displaying text.</p>	

Criteria	Supporting Features	Remarks and Explanations
(g) Applications shall not override user selected contrast and color selections and other individual display attributes.	<p>Supports</p> <p>SQL Anywhere supports system-wide colors settings as identified in the operating system display properties.</p>	<p>Customizations made in SQL Anywhere utilities for visual representation are specific to the application, and does not change system-wide settings.</p>
(h) When animation is displayed, the information shall be displayable in at least one non-animated presentation mode at the option of the user.	<p>Supports</p> <p>SQL Anywhere uses animation sparingly.</p>	<p>Where animation is used, the information conveyed is not essential, or is supplementary.</p>
(i) Color coding shall not be used as the only means of conveying information, indicating an action, prompting a response, or distinguishing a visual element.	<p>Supports</p> <p>SQL Anywhere uses color coding in association with other visual indicators when conveying information, indicating an action, prompting a response, or distinguishing a visual element.</p>	<p>The SQL Central database statistics graphing application provides a user-selectable range of colors, while providing text/numeric output of the information as well.</p>
(j) When a product permits a user to adjust color and contrast settings, a variety of color selections capable of producing a range of contrast levels shall be provided.	<p>Supports</p> <p>Where color is used to convey information, SQL Anywhere allows a user to select from a wide range of colors.</p>	<p>SQL Anywhere supports operating system settings for high contrast color schemes that make the screen easier to view for some users by heightening screen contrast with alternate color combinations.</p>
(k) Software shall not use flashing or blinking text, objects, or other elements having a flash or blink frequency greater than 2 Hz and lower than 55 Hz.	<p>Supports</p> <p>SQL Anywhere does not use flashing or blinking objects or text in any application user interface that falls within this frequency range.</p>	
(l) When electronic forms are used, the form shall allow people using Assistive Technology to access the information, field elements, and functionality required for completion and submission of the form, including all directions and cues.	<p>Supports</p> <p>SQL Anywhere windows support assistive technologies.</p>	<p>For the Java-based components in SQL Anywhere, third-party screen reader used must support Java assistive technology.</p>

Section 1194.31 Functional Performance Criteria

Criteria	Supporting features	Remarks and explanations
(a) At least one mode of operation and information retrieval that does not require user vision shall be provided, or support for Assistive Technology used by people who are blind or visually impaired shall be provided.	Supports	SQL Anywhere supports 3rd party technologies that make computer programs more accessible to people who use assistive technology. SQL Anywhere also supports the use of screen readers to access user interface information. 3rd party products supporting Java Assistive Technologies are required for Java-Based tools.
(b) At least one mode of operation and information retrieval that does not require visual acuity greater than 20/70 shall be provided in audio and enlarged print output working together or independently, or support for Assistive Technology used by people who are visually impaired shall be provided.	Supports with exceptions	SQL Anywhere supports the use of screen readers to access user interface information. SQL Anywhere, with few exceptions, also supports system large font settings.
(c) At least one mode of operation and information retrieval that does not require user hearing shall be provided, or support for Assistive Technology used by people who are deaf or hard of hearing shall be provided.	Supports SQL Anywhere does not require user hearing for access to any SQL Anywhere component functionality.	In all instances where SQL Anywhere provides an audio cue, it provides a visual cue as well.
(d) Where audio information is important for the use of a product, at least one mode of operation and information retrieval shall be provided in an enhanced auditory fashion, or support for assistive hearing devices shall be provided.	Supports SQL Anywhere does not require user hearing for access to any SQL Anywhere component functionality.	In all instances where SQL Anywhere provides an audio cue, it provides a visual cue as well.
(e) At least one mode of operation and information retrieval that does not require user speech shall be provided, or support for Assistive Technology used by people with disabilities shall be provided.	Supports SQL Anywhere components do not require speech recognition.	

Criteria	Supporting features	Remarks and explanations
(f) At least one mode of operation and information retrieval that does not require fine motor control or simultaneous actions and that is operable with limited reach and strength shall be provided.	Supports SQL Anywhere supports Microsoft Windows operating system tools such as MouseKeys, StickyKeys and FilterKeys, as well as other assistive technologies.	

Section 1194.41 Information, Documentation and Support

Criteria	Supporting Features	Remarks and Explanations
(a) Product support documentation provided to end-users shall be made available in alternate formats upon request, at no additional charge.	Fully supports	Documentation is provided in digital format for customers on the web and available for no-charge download.
(b) End-users shall have access to a description of the accessibility and compatibility features of products in alternate formats or alternate methods upon request, at no additional charge.	Fully supports	SQL Anywhere includes Help File documentation listing accessibility features. An individual may access and view Help File topics online or may select a topic for print.
(c) Support services for products shall accommodate the communication needs of end-users with disabilities.	Fully supports SQL Anywhere does not require user hearing for access to any SQL Anywhere component functionality.	The technical support group offers a variety of services including Technical Documents, White Papers, Product Manuals, Newsgroups, online Software Problem Case Management, and Bug Fixes and Patches. For information about these and other support services, visit the technical support web site. See SAP Support Portal .

Accessibility Using the Java Access Bridge

The Java Access Bridge module, which is loaded whenever you use SQL Central or Interactive SQL, supports assistive technologies. Third-party software such as screen readers use this module to provide access to software features.

1.7 Data Management Technologies

SQL Anywhere offers two relational databases: SQL Anywhere Server and UltraLite.

SQL Anywhere Server

SQL Anywhere Server provides enterprise-caliber functionality, including full transaction processing, referential integrity, materialized views, snapshot isolation, high availability via database mirroring and server clustering, SQL and Java stored procedures, triggers, row-level locking, automatic event scheduling, automatic backup and recovery, full-text searching, and support for spatial data. SQL Anywhere Server easily scales to hundreds of concurrent users and hundreds of gigabytes of data. Yet its small footprint and its many features that automate administration make it an ideal database to embed into server and desktop applications that are widely deployed in customer and remote sites.

UltraLite

For environments that demand smaller data-driven applications, the UltraLite database is ideal. UltraLite is a full relational database management system designed specifically to minimize memory and system requirements for deployment to handhelds and other mobile devices. It provides full transaction processing, a choice of development models, and a built-in synchronization client for exchanging data with other databases.

In this section:

[SQL Anywhere and UltraLite Target Platforms \[page 18\]](#)

When deciding whether to use a SQL Anywhere database or an UltraLite database, consider your target platform.

[Multi-Tier Computing Architecture \[page 18\]](#)

In multi-tier computing, application logic is held in an application server, such as SAP NetWeaver, WebLogic, or WebSphere, which sits between the database server and the client applications.

[ETL Features \[page 19\]](#)

Extract, Transform, and Load (ETL) is the process by which large amounts of data are extracted from disparate data sources and consolidated into a single database.

1.7.1 SQL Anywhere and UltraLite Target Platforms

When deciding whether to use a SQL Anywhere database or an UltraLite database, consider your target platform.

Database	Target Platform	Considerations
SQL Anywhere	Unix, Apple macOS, Microsoft Windows, or Linux	SQL Anywhere is often preferred because it provides a fuller feature set and its additional memory requirements are rarely an issue.
UltraLite	Embedded Linux, Google Android, Apple iOS, Microsoft Windows, or Linux	If the target platform is on an electronic device, then consider memory constraints and possibly the tasks your application needs to perform.

Related Information

[UltraLite and SQL Anywhere Feature Comparisons](#)

1.7.2 Multi-Tier Computing Architecture

In multi-tier computing, application logic is held in an application server, such as SAP NetWeaver, WebLogic, or WebSphere, which sits between the database server and the client applications.

In many situations, a single application server can access multiple databases in addition to non-relational data stores. In the Internet case, client applications are browser-based, and the application server is generally a web server extension. Many modern multi-tier applications use a service-oriented architecture (SOA) based on web services.

SAP NetWeaver stores application logic in the form of components, and makes these components available to client applications. The components can be SAP PowerBuilder components, JavaBeans, or COM components.

Application servers can also provide transaction logic to their client applications, guaranteeing that sets of operations are executed atomically across multiple databases. SQL Anywhere is well suited to multi-tier computing, and can participate in distributed transactions coordinated by Microsoft Distributed Transaction Coordinator (MSDTC). Both SAP NetWeaver and Microsoft Transaction Server use DTC to provide transaction services to their client applications.

The built-in support for web services makes SQL Anywhere a good choice for many multi-tier or SOA applications.

Related Information

[Three-tier Computing and Distributed Transactions](#)

1.7.3 ETL Features

Extract, Transform, and Load (ETL) is the process by which large amounts of data are extracted from disparate data sources and consolidated into a single database.

In the extraction phase, data is parsed and evaluated for suitability. During transformation, data is manipulated to achieve the format required for storage. Some common transformations include the elimination of unnecessary columns, calculation of computed values, and translation of values such as dates into a common format so that the data can be consolidated. The data is then loaded into the database at a frequency and scope consistent with the organization's needs.

SQL Anywhere offers several features in support of ETL. For example:

OPENSTRING operation

Use the OPENSTRING operation in the FROM clause to transform and load data from client- and server-side data sources.

OPENXML operator

Use the OPENXML string operator to extract data from XML documents.

MERGE statement

Use the MERGE statement to merge data from different source objects.

Proxy tables

Use proxy tables to access objects such as tables, views, and materialized views in a remote database.

System procedure calls in the FROM clause

Use various system procedures in the FROM clause of a query to extract and transform data for loading.

Related Information

[Proxy Tables](#)

[System Procedures](#)

[FROM Clause](#)

[OPENXML Operator](#)

[MERGE Statement](#)

1.8 Overview of Data Exchange Technologies

SQL Anywhere offers a wide range of options for exchanging data with existing enterprise systems and mobile devices.

Data exchange technologies include synchronization, replication, messaging, and mobile web service technologies. There are three SQL Anywhere synchronization technologies: Relay Server, MobiLink, and SQL Remote. When an application modifies shared data at any one database, the changes are propagated to other databases in the synchronization system. Changes can be propagated by various means and through a variety of channels, allowing flexible application architecture while preserving data integrity.

All SQL Anywhere synchronization technologies provide two-way synchronization. Changes made at the central database are propagated to remote databases. Changes made at remote databases are propagated to the central database, and to other remote databases.

Relay Server

Relay Server enables secure, load-balanced communication between mobile devices and backend servers through a web server. Supported backend servers include MobiLink, SQL Anywhere, SAP Afaria, SAP Mobile Office, SAP Mobile Platform, and SAP Mobile Server.

SQL Anywhere offers two synchronization technologies:

MobiLink synchronization

MobiLink offers session-based, bi-directional synchronization. It is ideal for exchanging data between a consolidated database and many remote UltraLite or SQL Anywhere databases, or between a central, non-relational data source and many remote UltraLite or SQL Anywhere databases.

MobiLink allows upload-only, download-only, and two-way synchronization.

During a MobiLink synchronization, the remote database uploads changes that were made to it since the previous synchronization with the MobiLink server. On receiving this data, the MobiLink server updates the consolidated database and then downloads changes from the consolidated database to the remote database. It also ensures the transactional integrity of the databases in the event a connection between them is lost, and provides mechanisms for the resolution of data change conflicts.

MobiLink file transfer functionality lets you transfer files to remote applications on the same connection you use to synchronize data, which is useful when populating new remote databases or upgrading software.

In addition, MobiLink provides direct row handling for synchronizing remote data with any central data source. The data sources to which you can synchronize can include an application, a web server, a web service, an application server, a text file, a spreadsheet, a non-relational database, or an RDBMS that is not supported as a consolidated database.

SQL Remote replication

SQL Remote is a data-replication technology designed for two-way synchronization between a consolidated database and large numbers of remote databases, typically including many mobile databases.

SQL Remote supports SQL Anywhere databases. Each database in the system must have a similar schema.

SQL Remote preserves transactional integrity, making it ideal for many business applications, particularly those that operate in environments where connections are unreliable. Furthermore, memory and disk space requirements are minimal for all components of the replication system.

Messages are typically files or specially formatted email messages. A **message agent**, attached to each database, sends messages regarding changes to its own data. The same agent also receives messages from one or more other databases and modifies the database according to the contents of the received messages.

In message-based communications, each message carries its destination address and other control information so that no direct connection is necessary between applications exchanging information. For example, an email message contains the destination address; there is no direct connection between the sending server and the recipient.

Replication messages are delivered occasionally, in periods of seconds, minutes, hours, or days, and as such, a high transaction volume at each remote site can lead to a very large volume of messages. SQL Remote is best suited to systems with a relatively low volume of replicated data per remote database. However, at the consolidated site, SQL Remote can prepare messages efficiently by preparing messages for multiple sites simultaneously.

Building a synchronization system on top of a message system means that a message-based synchronization system, such as SQL Remote, does not need to implement a store-and-forward system to get messages to their destination. Just as session-based client/server applications do not implement their own protocol stacks to pass information between client and server, so SQL Remote uses existing message systems to pass the messages.

To work reliably, a message-based synchronization system must both guarantee that all messages reach their destination and that the messages are applied in the same order that they are sent. SQL Remote incorporates a protocol to guarantee application of synchronization updates in the correct order.

The following table summarizes the characteristics of MobiLink and SQL Remote.

Synchronization Technology	Number of Databases	Connection	Frequency	Consolidated Database Types
MobiLink	Large	Occasional	Medium	Many options
SQL Remote	Large	Occasional	Low	SQL Anywhere

In this section:

[Challenges for Synchronization Technologies \[page 22\]](#)

Any synchronization technology must address challenges that arise as a result of the increased flexibility permitted by synchronization.

[How to Choose a Synchronization Technology \[page 23\]](#)

There are several considerations to make when choosing your synchronization technology.

[Consolidated and Remote Databases \[page 25\]](#)

MobiLink and SQL Remote provide data synchronization between a central database and a set of remote databases.

[Hierarchical Database Configurations \[page 26\]](#)

For databases in a **hierarchical configuration**, every database has a single parent database, except the consolidated database, which has no parent.

Related Information

[SQL Remote Systems](#)
[Relay Server Architecture](#)

1.8.1 Challenges for Synchronization Technologies

Any synchronization technology must address challenges that arise as a result of the increased flexibility permitted by synchronization.

Transactional Integrity

One of the challenges of any synchronization system is to ensure that each database always retains transactional integrity.

SQL Remote replicates portions of the transaction log in such a way that transactions are maintained during synchronization: either a whole transaction is replicated, or none of it is replicated. This ensures transactional integrity at each database in the system.

In MobiLink, you can also choose to replicate each transaction, but by default MobiLink coalesces multiple transactions on the remote database and applies them in a single transaction. This generally results in more efficient uploads. In both cases, MobiLink maintains transactional integrity.

Data Consistency

Another challenge to synchronization systems is to maintain data consistency throughout the system. Synchronization systems maintain a **loose consistency** in the system as a whole: that is, all changes are replicated to each site over time in a consistent manner, but different sites may have different copies of data at any instant.

Related Information

[Synchronization Techniques](#)

1.8.2 How to Choose a Synchronization Technology

There are several considerations to make when choosing your synchronization technology.

Your Consolidated Database System

In a typical synchronization environment, a large database serves as a central repository for information. Sometimes you can choose a database system that suits your needs. Other times, a central database already exists and you must adapt the synchronization system to work with it.

MobiLink works with many popular database servers, including SQL Anywhere, SAP HANA, SAP Adaptive Server Enterprise, Oracle, Microsoft SQL Server, Microsoft Azure, MySQL, SAP IQ, and IBM DB2 UDB. Using the MobiLink server APIs for .NET and Java, you can synchronize with any data source, including application servers, web servers, text files, and other database products.

i Note

Support for IBM DB2 consolidated databases is deprecated.

In a SQL Remote system, the central database must be SQL Anywhere.

Your Remote Database System

Synchronization technologies also differ in the types of remote databases that they can support.

MobiLink supports SQL Anywhere and UltraLite remote databases.

SQL Remote supports SQL Anywhere remote databases.

Network Characteristics

MobiLink and SQL Remote are both well suited to occasionally connected environments, where remote sites must operate for hours or days in isolation, although more frequent synchronization is possible whenever a network connection is available.

MobiLink is session based. A real-time connection is required during synchronization. If this connection is interrupted before synchronization is complete, the process does not complete until the next synchronization. In contrast, SQL Remote relays information via messages, which can be sent or received asynchronously. These messages may take the form of files on a hard disk, or email messages. These messages can be processed whenever they are received, allowing synchronization to occur incrementally.

Frequency of Synchronization

In some situations, it may be important that your information is replicated immediately. In others, synchronization once or twice a day may suffice. In fact, more frequent synchronization may be impossible when no network connection is available.

Both MobiLink and SQL Remote are primarily intended for situations where synchronization occurs infrequently, such as every few hours or days, but both can be used to synchronize as frequently as every few seconds.

The Number of Remote Sites

MobiLink and SQL Remote both work well with a very large number of remote users. MobiLink scalability is limited only by the scalability of the consolidated database management system. The SQL Remote message-based design allows a typical installation to handle thousands of remote users.

There is no hard limit on the maximum number of remote sites with any of these systems. The actual number depends on the amount of information replicated, the frequency of synchronization, and the design of your application.

Transaction Ordering

By default, MobiLink works by grouping the results of multiple transactions on the remote database into one set of changes to be applied to the consolidated database. Alternatively, you can choose to preserve the order of transactions and upload them separately. In both cases, synchronization always occurs at a transaction boundary, and so referential integrity is preserved. Uncommitted data is never synchronized, and so data integrity is preserved.

SQL Remote replicates data by scanning the transaction log and preparing messages, as appropriate, for each transaction. It orders these messages and sends them to the remote or consolidated site. When processing receives the messages, SQL Remote always processes them in the same order as they were applied to the other database. When necessary, it automatically delays processing a message until all earlier messages have been applied.

Achieving Data Consistency at a Particular Time

Immediately following each MobiLink synchronization session, the data in the two databases is consistent. The ability to guarantee the consistency of the data at a remote site at a particular point in time is an advantage of MobiLink session-based synchronization. For example, if it is important that the data at a remote site accurately reflect the data in the consolidated database at a particular time, such as 10 a.m., this objective can be achieved by synchronizing before this time. As long as the synchronization completes successfully, the currency of the data at the remote site is assured.

When changes to the data are replicated through an exchange of messages, it is difficult to guarantee that the data in a particular remote site is completely consistent with the data in the consolidated site at any particular

point in time. For example, sometimes a message is lost in transit. SQL Remote automatically recognizes this fault and resends the message, but such interruptions can cause unexpected delays.

1.8.3 Consolidated and Remote Databases

MobiLink and SQL Remote provide data synchronization between a central database and a set of remote databases.


In MobiLink, the **consolidated database** is one of several supported RDBMSs. The consolidated database, which typically resides on a corporate server, tracks synchronization information and optionally contains the data to be replicated. Other central data may be stored in any other format, such as a non-relational database, web service, or text file.

MobiLink also provides direct row handling, which enables data synchronization to consolidated data sources other than relational databases including enterprise resource planning (ERP) systems or application servers.

In SQL Remote, all data that is to be synchronized is contained in a SQL Anywhere consolidated database.

A **remote database** can run either at the same site as the consolidated database or at a physically distant site such as a handheld device. The remote database can share all or some of the data in the consolidated database.

The following figure displays a schematic illustration of a small synchronization system.



Remote Users

A typical synchronization system includes many remote databases. Each remote database contains a subset of the information in the central database. Each remote database is a physically separate database, usually on a separate computer or mobile device. All remote databases must stay consistent with the central database.

The entire synchronization system may be considered a single dispersed database, with the master copy of all shared data being kept at the central database.

Each remote site that synchronizes with the central database is considered a remote user of the central database. In the case that a remote site is a multi-user server, the entire site is considered to be a single remote user of the central database.


1.8.4 Hierarchical Database Configurations

For databases in a **hierarchical configuration**, every database has a single parent database, except the consolidated database, which has no parent.

SQL Remote supports hierarchical configurations of databases; it does not support peer-to-peer synchronization or other non-hierarchical configurations. MobiLink is also normally used with a hierarchical configuration, but it can also be used in other configurations.


For any two databases directly sharing data in a hierarchical configuration, one is always above or below the other in the hierarchy.

Hierarchical database configurations



Databases in a non-hierarchical configuration do not have a well-defined notion of above or below.

Non-hierarchical configurations



In a MobiLink or SQL Remote system, each database contains all or a subset of the data replicated by the database above it in the hierarchy.

Remote databases can contain tables that are not present at the consolidated database as long as they are not involved in synchronization. SQL Remote requires that the table and column names in the remote databases match the ones in the consolidated database. In contrast, MobiLink allows data to be stored in different columns and tables in the remote databases than in the consolidated database, allowing greater flexibility.

1.9 Design and Management Tools

There are many design and management tools for administering, monitoring, and troubleshooting databases.

i Note

Adobe will stop updating and distributing the Flash Player at the end of 2020. Because the SQL Anywhere Monitor is based on Flash, you cannot use it once Flash support ends. In many cases, tasks that were previously performed in the Monitor can be performed in the SQL Anywhere Cockpit. See [SQL Anywhere Monitor Non-GUI User Guide](#).

SQL Central: centralized control and administration

SQL Central is an integrated database administration and development tool that provides access to database settings, properties, and utilities in a graphical user interface. Via plug-ins, SQL Central manages SQL Anywhere Server, UltraLite, MobiLink, and other products.

In addition to helping with routine tasks, SQL Central also provides performance statistics, procedure profiling, stored procedure debugging, and the management of events and schedules, web services, and connection profiles. SQL Central helps administer any tasks that are performed by sending SQL statements to the database server, or performed by SQL Anywhere utilities.

A variety of SQL Central tools are available to help you analyze and monitor the current performance of your SQL Anywhere database. These tools include graphical plans, query executions, logging requests, timing utilities, and the performance monitor. In addition, SQL Central offers:

Support for spatial data

Spatial data describes the position, shape, and orientation of objects in a defined space. 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.

Index selection and optimization by using the Index Consultant

The Index Consultant analyzes workloads and provides recommendations on how to select indexes to optimize performance. The Index Consultant can be run from either SQL Central or Interactive SQL.

Interactive SQL: SQL query editor

Interactive SQL is a database utility designed to execute SQL statements and display database data. The built-in query editor and other tools, such as the graphical plan display, help you to analyze, troubleshoot, and optimize queries.

SQL Anywhere Monitor

The Monitor is a browser-based administration tool that provides you with information about the health and availability of SQL Anywhere databases, MobiLink servers, and MobiLink server farms. The Monitor provides constant data collection, email alert notifications, a browser-based interface, and the ability to monitor multiple databases, MobiLink servers, and MobiLink server farms. It also monitors the availability of web servers, proxy servers, and host computers.

SQL Anywhere Cockpit

Cockpit is a light-weight, low-impact monitoring application that makes it easy to view connected users, database and server properties, as well as the database server messages.

SQL Anywhere Profiler

Profiler is a development and troubleshooting tool for fine tuning SQL statements. Use it to:

- Profile stored procedures, functions, triggers, and events
- Receive recommendations to help improve the performance of your database application
- Capture database activity while your application is running

MobiLink Profiler: synchronization monitoring

The MobiLink Profiler is a graphical administration tool that provides details about the performance of MobiLink synchronizations. The MobiLink Profiler collects details and statistical summaries about all synchronizations that occur, including start and end times, data volume uploaded and downloaded, successful completions, conflicts, and more.

Utilities

SQL Anywhere includes various utilities for performing administration tasks such as backing up a database and performing synchronizations. Utilities are useful for including in batch files for repeated use.

Related Information

[SQL Central](#)
[Diagnostic Tracing \(Deprecated\)](#)
[SQL Anywhere Profiler](#)
[Optimize Indexes to Improve Performance](#)
[Interactive SQL](#)
[SQL Anywhere Monitor](#)
[MobiLink Profiler](#)
[Database Administration Utilities](#)
[UltraLite Utilities](#)
[MobiLink Utilities](#)
[Database Creation](#)
[Cockpit Security](#)
[SQL Anywhere Profiler](#)
[SQL Remote Utilities and Options Reference](#)

1.10 Sample Databases

Two sample databases are provided with the software.

In this section:

[The SQL Anywhere Sample Database \(demo.db\) \[page 29\]](#)

For consistency and simplicity, many of the examples throughout the documentation use the SQL Anywhere sample database, `%SQLANYSAMP17%\demo.db`.

[The CustDB Sample Database Application \[page 32\]](#)

The CustDB sample application is a useful tool for learning how to develop UltraLite and MobiLink applications.

1.10.1 The SQL Anywhere Sample Database (demo.db)

For consistency and simplicity, many of the examples throughout the documentation use the SQL Anywhere sample database, `%SQLANYSAMPI7%\demo.db`.

By default this file is installed in the SQL Anywhere samples directory (for example, `C:\Users\Public\Documents\SQL Anywhere 17\Samples` on Windows 7).

The sample database uses the following default user ID and password:

- User ID = DBA
- Password = sql (passwords in SQL Anywhere are case sensitive)

Caution

The sample database always has the same user ID and password; change the DBA user ID and password to restrict access to the database.

The sample database uses the following ODBC data source: [SQL Anywhere 17 Demo](#).

- You do not need to enter a user ID because the SQL Anywhere 17 Demo data source already contains this information, but you must specify the password `sql`.
In a production environment, do not store passwords in an ODBC data source.
- The database server does not need to be running before connecting since the SQL Anywhere 17 Demo data source automatically starts the database server.

The sample database represents a small company that sells athletic clothing. It contains internal information about the company (employees, departments, and finances), product information, and sales information (sales orders, customers, and contacts). All data in the database is fictional.

Roles in the Sample Database

The sample database contains one user-extended role (GROUPO) and three user-defined roles (READ_ROLE, MODIFY_ROLE, and EXEC_ROLE).

GROUPO is a user that has been extended as a role. GROUPO is the owner of the tables described above. It is also the owner of a number of sample views, materialized views, and stored procedures that are defined in the sample database. Any new user that is granted the role GROUPO can reference these database objects without the owner qualifier. For example, they can refer to the Departments table without adding the owner prefix (GROUPO.Departments).

Any user that is granted the role READ_ROLE has read access to the rows in the tables described above. Users with this role can perform a SELECT on any of these tables.

Any user that is granted the role MODIFY_ROLE has write access to the rows in the tables described above. Users with this role can perform an INSERT, UPDATE, or DELETE on any of these tables.

Any user that is granted the role EXEC_ROLE can execute any of the sample stored procedures that are defined in the sample database with the exception of ManageContacts (ShowContacts, ShowCustomers, ShowCustomerProducts, ShowProductInfo, ShowSalesOrders, ShowSalesOrderDetail, ShowSalesOrderItems, and debugger_tutorial). The ManageContacts stored procedure is used to add, delete and update contacts in the Contacts table.

Any user with only roles READ_ROLE, MODIFY_ROLE, or EXEC_ROLE cannot perform schema changes on the database. For example, they cannot add or remove columns from existing tables nor can they create or delete tables. Additional privileges are required for this.

In a large enterprise, there would be many such roles. For example, a Sales department might own the sales-related and customer-related tables, a Manufacturing or Purchasing department might own the products-related tables, a Human Resources department might own the employees-related and departments-related tables, and so on. Each of these departments would have their own collection of roles set up with various access rights to the tables they own. However, the sample database reflects a small enterprise and the use of roles has been deliberately kept simple.

The user "DBA" has all privileges for the sample database and can alter the database schema, create new users, add roles, and so on. This user has been granted the GROUPO user-extended role so that it can reference the sample tables and stored procedures without qualifying them with an owner ID.

Two additional users have been created to demonstrate the roles described above.

The user "UPDATER" (password "update") has been granted the roles READ_ROLE, MODIFY_ROLE, and EXEC_ROLE. This user can perform a SELECT, INSERT, UPDATE, or DELETE on any of the sample tables. This user has EXECUTE privilege on all of the sample stored procedures. The GROUPO user-extended role is also granted so that this user can reference the sample tables and stored procedures without qualifying them with an owner ID.

The user "BROWSER" (password "browse") has been granted the roles READ_ROLE, and EXEC_ROLE. This user can perform a SELECT on any of the sample tables but cannot change the rows in the tables in any way. This user has EXECUTE privilege on all of the sample stored procedures with the exception of the ManageContacts stored procedure since this procedure is used to add, delete and update contacts in the Contacts table (the role MODIFY_ROLE is required to execute ManageContacts). The GROUPO user-extended role is also granted so that this user can reference the sample tables and stored procedures without qualifying them with an owner ID.

In this section:

[Recreating the Sample Database \(demo.db\) \[page 31\]](#)

Restore the sample database to its original state or preserve the sample database in its current state and recreate the sample database in its original state using a different name.

Related Information

[Changing a Password \(SQL Central\)](#)

[Tutorial: Connecting to the Sample Database](#)

1.10.1.1 Recreating the Sample Database (*demo.db*)

Restore the sample database to its original state or preserve the sample database in its current state and recreate the sample database in its original state using a different name.

Context

Testing features and completing the tutorials sometimes results in changes to the sample database that can prevent the successful completion of subsequent tutorials and tests.

Procedure

Choose one of the following options:

Option	Description
Recreate the sample database (Windows)	<p>Run the following command to erase <i>demo.db</i> and create a new copy of the sample database with objects and data:</p> <pre>newdemo "%SQLANYAMP17%\demo.db"</pre> <p>When you are prompted, choose to erase any existing files.</p>
Recreate the sample database (UNIX and Linux)	<p>Run the following command in the directory where the sample database is located to erase <i>demo.db</i> and create a new copy of the sample database with objects and data:</p> <pre>newdemo.sh demo.db</pre> <p>When you are prompted, choose to erase any existing files.</p>
Create a copy of the sample database with a different name (Windows)	<p>Run the following command to create a database called <i>mydemo.db</i> that contains objects and data.</p> <p>If you do not specify a path, the database and transaction log are created in the current directory.</p> <pre>newdemo path\mydemo.db</pre>
Create a copy of the sample database with a different name (UNIX and Linux)	<p>Run the following command to create a database called <i>mydemo.db</i> that contains objects and data.</p> <p>If you do not specify a path, the database and transaction log are created in the current directory.</p> <pre>newdemo.sh path/mydemo.db</pre>

Results

The database has been recreated to its original state or a copy of the sample database has been created with a different name.

1.10.2 The CustDB Sample Database Application

The CustDB sample application is a useful tool for learning how to develop UltraLite and MobiLink applications.

The sample database is a sales status database for a hardware supplier. It holds customer, product, and sales force information for the supplier.

There are two parts to the CustDB sample application:

UltraLite

For UltraLite, CustDB can be deployed on any device supported by UltraLite using any platform supported by UltraLite. You can see all the source code used to create the CustDB UltraLite application and run the sample. The CustDB sample application is set up for MobiLink synchronization.

You can find the UltraLite CustDB sample application in `%SQLANYAMP17%\UltraLite\CustDB\`.

MobiLink

If you are interested in exploring MobiLink features, there is a CustDB consolidated database that contains sample synchronization logic. You can use this consolidated database with the CustDB UltraLite sample remote database to run the sample synchronization system.

The MobiLink consolidated CustDB database is created by running script files against a supported relational database (SQL Anywhere, Adaptive Server Enterprise, Oracle, Microsoft SQL Server, or IBM DB2). These setup files are located in `%SQLANYAMP17%\MobiLink\CustDB\`.

i Note

Support for IBM DB2 consolidated databases is deprecated.

The CustDB consolidated database uses the following ODBC data source: [SQL Anywhere 17 CustDB](#).

The SQL Anywhere 17 CustDB uses the following default user ID and password:

- User ID = ml_server
- Password = sql (passwords in SQL Anywhere are case sensitive)

⚠ Caution

The sample database always has the same user ID and password; change the DBA user ID and password to restrict access to the database.

Related Information

[CustDB Sample for MobiLink](#)

1.11 Using the Documentation

This Help assists you with your day-to-day tasks and includes detailed product information and procedures.

The complete documentation is available in the following formats:

- SAP Help Portal (https://help.sap.com/viewer/product/SAP_SQL_Anywhere/17.0/en-US)
The latest documentation is available on SAP Help Portal.
- DocCommentXchange (<http://dcx.sap.com/index.html>)
- HTML Help
- PDF

To download the documentation in other formats, see <http://scn.sap.com/docs/DOC-49456>.

In this section:

[SAP Help Portal \[page 33\]](#)

SAP Help Portal provides documentation for all SAP products.

[DocCommentXchange \[page 34\]](#)

DocCommentXchange (DCX) is a community for accessing and discussing the documentation.

[HTML Help \[page 34\]](#)

The navigation pane of this Help includes an integrated table of contents, index, and word search that allow you to find the information you need quickly and easily. It also allows you to bookmark frequently viewed topics for future reference.

[Documentation Conventions \[page 35\]](#)

There are several conventions used in this documentation.

[Online Resources and Additional Help \[page 38\]](#)

There are several online resources and sources of additional help available.

1.11.1 SAP Help Portal

SAP Help Portal provides documentation for all SAP products.

Use SAP Help Portal to:

- View documentation
- Provide suggestions and corrections to improve documentation for all users in future releases by clicking the [Send feedback](#) link on a topic.

1.11.2 DocCommentXchange

DocCommentXchange (DCX) is a community for accessing and discussing the documentation.

Use DocCommentXchange to:

- View documentation
- Check for clarifications users have made to sections of documentation
- Provide suggestions and corrections to improve documentation for all users in future releases

1.11.3 HTML Help

The navigation pane of this Help includes an integrated table of contents, index, and word search that allow you to find the information you need quickly and easily. It also allows you to bookmark frequently viewed topics for future reference.

There are four tabs that make up the navigation pane:

Tab	Description
Contents	Click Contents to reveal the Table of Contents. Click the book icons to reveal topics and sub-topics. As you move through the Help, the Table of Contents tracks the different topics that you access. This orients you within the Help, and lets you to see other relevant topics at a glance.
Index	Click Index to view an alphabetical list of the subjects covered in the Help. To move through the index, either type a word or scroll through the list. Double-click an entry to navigate to the related topic or topics.

Tab	Description
Search	<p>Click Search to search for occurrences of a word or phrase. In the Search field, type the words you want to search for and press Enter, or click List Topics. The topics that contain these words are listed. Double-click a topic name to open it. When searching, use the following commands:</p> <ul style="list-style-type: none"> • Use xx AND yy to search for the topics that contain both xx and yy. • Use xx OR yy to search for the topics that contain either xx and/or yy. • Use xx NOT yy to search for the topics that contain xx but not yy. • Use xx NEAR yy to search for topics that contain xx and yy near each other. • Type a phrase within quotation marks to search for topics that contain the exact phrase. For example, "xx yy". <p>Since the terms and, or, not, and near are interpreted as commands in Help, you must put quotes around them when you want to search for them in the documentation. For example, to search for near, you must enter near.</p> <p>When searching for help topics, use the wildcard characters ? and * in place of other characters:</p> <ul style="list-style-type: none"> • Use the question mark (?) wildcard to replace a single character. For example, typing serve? finds server. • Use the (*) asterisk wildcard to replace one or more characters. For example, typing searc* finds search, searching, and searched.
Favorites	<p>Click Favorites to bookmark frequently viewed topics for future reference.</p> <p>Click Add to add the currently displayed topic to the Topics list. To view the topic, double-click it, or select it and click Display.</p>

1.11.4 Documentation Conventions

There are several conventions used in this documentation.

Operating systems

SQL Anywhere runs on many platforms. Typically, the behavior of the software is the same on all platforms, but occasionally there are variations or limitations. These limitations are commonly based on the underlying operating system (Microsoft Windows, UNIX, or Linux), and seldom on the particular version or

variant of the operating system (for example, IBM AIX). To simplify references to operating systems, the documentation groups the supported operating systems as follows:

Microsoft Windows

Unless otherwise specified, when the documentation refers to Microsoft Windows, it refers to all supported Microsoft Windows-based platforms.

UNIX

Unless otherwise specified, when the documentation refers to UNIX, it refers to all supported UNIX platforms, including macOS.

Linux

Unless otherwise specified, when the documentation refers to Linux, it refers to all supported Linux platforms.

For the complete list of platforms supported by SQL Anywhere, see the [SAP SQL Anywhere Supported Platforms and Engineering Support Status](#).

Directory and file names

These conventions are used to simplify the documentation of directory and file names:

Uppercase and lowercase directory names

On Microsoft Windows, UNIX and Linux, directory and file names can contain uppercase and lowercase letters.

On Microsoft Windows, references to directories and files are *not* case sensitive. Mixed case directory and file names are common, but it is common to refer to them using all lowercase letters.

On UNIX and Linux, references to directories and files are case sensitive. Mixed case directory and file names are not common. Most use all lowercase letters.

Slashes separating directory and file names

The documentation uses backslashes (\) as the directory separator. On UNIX and Linux, replace the backslash with the forward slash (/).

Executable files

The documentation shows executable file names using Microsoft Windows conventions, with a suffix such as `.exe` or `.bat`. On UNIX and Linux, executable file names have no suffix.

32-bit and 64-bit versions

The documentation does not always distinguish between 32-bit and 64-bit versions of the software. In some cases, the documentation provides an example file path that includes `bin32` or `bin64`. The two are interchangeable, and depend on the version of the software that was installed. For example, the file paths `C:\Program Files\SQL Anywhere 17\bin32` and `C:\Program Files\SQL Anywhere 17\bin64` are equivalent.

Command prompts and command shell syntax

If these examples contain characters that the shell considers special, the command requires modification for the specific shell. The modifications are beyond the scope of this documentation, but generally, use quotes around the parameters containing those characters or use an escape character before the special characters.

Parentheses and curly braces

Some command line options require a parameter that accepts detailed value specifications in a list. The list is enclosed with parentheses or curly braces. The documentation uses parentheses. For example:

```
-x tcpip(host=127.0.0.1)
```

Where parentheses cause syntax problems, substitute curly braces:

```
-x tcpip{host=127.0.0.1}
```

If both forms result in syntax problems, enclose the entire parameter in quotes as required by the shell:

```
-x "tcpip(host=127.0.0.1)"
```

Semicolons

On UNIX and Linux, enclose semicolons in quotes.

Quotes

If you must specify quotes in a parameter value, the quotes can conflict with the traditional use of quotes to enclose the parameter. For example, to specify an encryption key whose value contains double quotes, enclose the key in quotes and then escape the embedded quote:

```
-ek "my \"secret\" key"
```

In many shells, the value of the key would be `my "secret" key`.

Related Information

[SQL Anywhere on Windows](#)

1.11.5 Online Resources and Additional Help

There are several online resources and sources of additional help available.

Table 1: Recommended Resources



Resources	Links
Technical papers	<ul style="list-style-type: none">• SQL Anywhere I/O Requirements for Windows and Linux• Capacity Planning with SQL Anywhere• Diagnosing Application Performance Issues with SQL Anywhere• Analytic Functions in SQL Anywhere• Optimizing Adaptive Server Anywhere Performance Over a WAN• Semantics and Compatibility of Transact-SQL Outer Joins
Software readme files	SQL Anywhere Readmes
SQL Anywhere forums	<ul style="list-style-type: none">• SAP SQL Anywhere Forum• SAP Community Network - SAP SQL Anywhere

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
 - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
 - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

© 2022 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.