



Adaptive Server[®] Anywhere Database Administration Guide

Last modified: October 2002
Part Number: 38123-01-0802-01

Copyright © 1989–2002 Sybase, Inc. Portions copyright © 2001–2002 iAnywhere Solutions, Inc. All rights reserved.

No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of iAnywhere Solutions, Inc. iAnywhere Solutions, Inc. is a subsidiary of Sybase, Inc.

Sybase, SYBASE (logo), AccelaTrade, ADA Workbench, Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, Adaptive Server Enterprise Replication, Adaptive Server Everywhere, Adaptive Server IQ, Adaptive Warehouse, AnswerBase, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-FORMS, APT-Library, APT-Translator, ASEP, Backup Server, BayCam, Bit-Wise, BizTracker, Certified PowerBuilder Developer, Certified SYBASE Professional, Certified SYBASE Professional (logo), ClearConnect, Client Services, Client-Library, CodeBank, Column Design, ComponentPack, Connection Manager, Convoy/DM, Copernicus, CSP, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, Dynamo, e-ADK, E-Anywhere, e-Biz Integrator, E-Whatever, EC-GATEWAY, ECMAP, ECRT, eFulfillment Accelerator, Electronic Case Management, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, eProcurement Accelerator, eremote, Everything Works Better When Everything Works Together, EWA, Financial Fusion, Financial Fusion Server, First Impression, Formula One, Gateway Manager, GeoPoint, iAnywhere, iAnywhere Solutions, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InstaHelp, Intellidex, InternetBuilder, iremote, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, Logical Memory Manager, MainframeConnect, Maintenance Express, MAP, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, MethodSet, ML Query, MobiCATS, MySupport, Net-Gateway, Net-Library, New Era of Networks, Next Generation Learning, Next Generation Learning Studio, O DEVICE, OASiS, OASiS (logo), ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Business Interchange, Open Client, Open Client/Server, Open Client/Server Interfaces, Open ClientConnect, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, Partnerships that Work, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, PhysicalArchitect, Pocket PowerBuilder, PocketBuilder, Power Through Knowledge, Power++, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, Powering the New Economy, PowerJ, PowerScript, PowerSite, PowerSocket, Powersoft, Powersoft Portfolio, Powersoft Professional, PowerStage, PowerStudio, PowerTips, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Rapport, Relational Beans, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Report Workbench, Report-Execute, Resource Manager, RW-DisplayLib, RW-Library, S Designer, S-Designer, S.W.I.F.T. Message Format Libraries, SAFE, SAFE/PRO, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILS, smart.partners, smart.parts, smart.script, SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL Server SNMP SubAgent, SQL Server/CFT, SQL Server/DBM, SQL SMART, SQL Station, SQL Toolset, SQLJ, Stage III Engineering, Startup.Com, STEP, SupportNow, Sybase Central, Sybase Client/Server Interfaces, Sybase Development Framework, Sybase Financial Server, Sybase Gateways, Sybase Learning Connection, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase Synergy Program, Sybase User Workbench, Sybase Virtual Server Architecture, SybaseWare, Syber Financial, SyberAssist, SybMD, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, The Enterprise Client/Server Company, The Extensible Software Platform, The Future Is Wide Open, The Learning Connection, The Model For Client/Server Solutions, The Online Information Center, The Power of One, TradeForce, Transact-SQL, Translation Toolkit, Turning Imagination Into Reality, UltraLite, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viewer, Visual Components, VisualSpeller, VisualWriter, VQL, Warehouse Control Center, Warehouse Studio, Warehouse WORKS, WarehouseArchitect, Watcom, Watcom SQL, Watcom SQL Server, Web Deployment Kit, Web.PB, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server, and XP Server are trademarks of Sybase, Inc. or its subsidiaries.

All other trademarks are property of their respective owners.

Last modified October 2002. Part number 38123-01-0802-01.

Contents

About This Manual.....	ix
SQL Anywhere Studio documentation.....	x
Documentation conventions.....	xiii
The Adaptive Server Anywhere sample database.....	xvi
Finding out more and providing feedback.....	xvii

PART ONE

Starting and Connecting to Your Database	1
---	----------

1	Running the Database Server	3
	Introduction	4
	Starting the server.....	8
	Some common options	10
	Stopping the database server.....	18
	Starting and stopping databases	20
	Running the server outside the current session	21
	Troubleshooting server startup	34

2	Connecting to a Database.....	37
	Introduction to connections	38
	Connecting from Sybase Central or Interactive SQL	42
	Simple connection examples	45
	Working with ODBC data sources	53
	Connecting from desktop applications to a Windows CE database.....	65
	Connecting to a database using OLE DB	68
	Connection parameters.....	70
	Troubleshooting connections	73
	Using integrated logins	83

3	Client/Server Communications.....	91
	Supported network protocols	92
	Using the TCP/IP protocol	93

	Using the SPX protocol.....	96
	Using Named Pipes	97
	Adjusting communication compression settings to improve performance	98
	Troubleshooting network communications	101
4	Adaptive Server Anywhere as an Open Server	105
	Open Clients, Open Servers, and TDS.....	106
	Setting up Adaptive Server Anywhere as an Open Server.....	108
	Configuring Open Servers	110
	Characteristics of Open Client and jConnect connections	116
5	The Database Server	119
	The database server	120
6	Connection and Communication Parameters	163
	Connection parameters.....	164
	Network communications parameters	189
 PART TWO		
	Working with Database Files	201
7	File Locations and Installation Settings.....	203
	Installation directory structure	204
	How Adaptive Server Anywhere locates files	206
	Environment variables	208
	Registry and INI files.....	213
8	Working with Database Files	217
	Overview of database files.....	218
	Using additional dbspaces	220
	Working with write files	224
	Using the utility database	226
9	Automating Tasks Using Schedules and Events	231
	Introduction	232
	Understanding schedules	234
	Understanding events	236

	Understanding event handlers	240
	Schedule and event internals.....	242
	Scheduling and event handling tasks	244
10	International Languages and Character Sets	249
	Introduction to international languages and character sets.....	250
	Understanding character sets in software	253
	Understanding locales	262
	Understanding collations.....	269
	Understanding character set translation	278
	Collation internals.....	282
	International language and character set tasks	287
11	Backup and Data Recovery.....	299
	Introduction to backup and recovery.....	300
	Understanding backups	305
	Designing backup procedures	308
	Configuring your database for data protection.....	317
	Backup and recovery internals.....	321
	Backup and recovery tasks.....	329
 PART THREE		
	Permissions, Security, and Replication	349
12	Managing User IDs and Permissions	351
	Database permissions overview	352
	Setting user and group options.....	356
	Managing individual user IDs and permissions	357
	Managing connected users.....	368
	Managing groups	369
	Database object names and prefixes.....	376
	Using views and procedures for extra security	378
	Changing Ownership on Nested Objects.....	381
	How user permissions are assessed	383
	Managing the resources connections use	384
	Users and permissions in the system tables	385
13	Keeping Your Data Secure.....	387
	Security features overview.....	388
	Controlling database access.....	390
	Auditing database activity	393

Running the database server in a secure fashion	397
Encrypting a database	399
Security tips.....	403

14	Replicating Data with Replication Server	405
	Introduction to replication.....	406
	Tutorial: Replicate data using Replication Server.....	409
	Configuring databases for Replication Server	419
	Using the LTM.....	422

PART FOUR

Utilities, Options, and Properties	433
---	------------

15	Database Administration Utilities.....	435
	Administration utilities overview	437
	The Backup utility.....	438
	The Collation utility.....	442
	The File Hiding utility.....	446
	The Compression utility	448
	The Console utility	450
	The Data Source utility.....	451
	The Debugger utility	456
	The Erase utility	458
	The Histogram utility	461
	The Information utility.....	463
	The Initialization utility.....	465
	The Interactive SQL utility.....	472
	The Language utility.....	477
	The License utility	479
	The Log Transfer Manager	481
	The Log Translation utility.....	488
	The Ping utility.....	494
	The Rebuild utility.....	497
	The Server Location utility	498
	The Service Creation utility	499
	The Spawn utility.....	503
	The Stop utility	505
	The Transaction Log utility.....	507
	The Uncompression utility.....	511
	The Unload utility	513
	The Upgrade utility	521
	The Validation utility.....	526
	The Write File utility	530

16	Database Options	535
	Introduction to database options.....	536
	Database options	541
	Compatibility options	544
	Replication options.....	547
	Interactive SQL options.....	548
	Alphabetical list of options	550
17	Database Performance and Connection Properties	609
	Database performance statistics.....	610
	Database properties.....	618
18	Physical Limitations	635
	Size and number limitations	636
	Index	639

About This Manual

Subject	This book covers material related to running, managing, and configuring databases. It describes database connections, the database server, database files, security, backup procedures, and replication with the Replication Server, as well as administration utilities and options.
Audience	This manual is for all users of Adaptive Server Anywhere. It is to be used in conjunction with other manuals in the documentation set.

SQL Anywhere Studio documentation

This book is part of the SQL Anywhere documentation set. This section describes the books in the documentation set and how you can use them.

The SQL Anywhere Studio documentation set

The SQL Anywhere Studio documentation set consists of the following books:

- ◆ **Introducing SQL Anywhere Studio** This book provides an overview of the SQL Anywhere Studio database management and synchronization technologies. It includes tutorials to introduce you to each of the pieces that make up SQL Anywhere Studio.
- ◆ **What's New in SQL Anywhere Studio** This book is for users of previous versions of the software. It lists new features in this and previous releases of the product and describes upgrade procedures.
- ◆ **Adaptive Server Anywhere Getting Started** This book is for people new to relational databases or new to Adaptive Server Anywhere. It provides a quick start to using the Adaptive Server Anywhere database-management system and introductory material on designing, building, and working with databases.
- ◆ **Adaptive Server Anywhere Database Administration Guide** This book covers material related to running, managing, and configuring databases.
- ◆ **Adaptive Server Anywhere SQL User's Guide** This book describes how to design and create databases; how to import, export, and modify data; how to retrieve data; and how to build stored procedures and triggers.
- ◆ **Adaptive Server Anywhere SQL Reference Manual** This book provides a complete reference for the SQL language used by Adaptive Server Anywhere. It also describes the Adaptive Server Anywhere system tables and procedures.
- ◆ **Adaptive Server Anywhere Programming Guide** This book describes how to build and deploy database applications using the C, C++, and Java programming languages. Users of tools such as Visual Basic and PowerBuilder can use the programming interfaces provided by those tools.

-
- ◆ **Adaptive Server Anywhere Error Messages** This book provides a complete listing of Adaptive Server Anywhere error messages together with diagnostic information.
 - ◆ **Adaptive Server Anywhere C2 Security Supplement** Adaptive Server Anywhere 7.0 was awarded a TCSEC (Trusted Computer System Evaluation Criteria) C2 security rating from the U.S. Government. This book may be of interest to those who wish to run the current version of Adaptive Server Anywhere in a manner equivalent to the C2-certified environment. The book does *not* include the security features added to the product since certification.
 - ◆ **MobiLink Synchronization User's Guide** This book describes all aspects of the MobiLink data synchronization system for mobile computing, which enables sharing of data between a single Oracle, Sybase, Microsoft or IBM database and many Adaptive Server Anywhere or UltraLite databases.
 - ◆ **SQL Remote User's Guide** This book describes all aspects of the SQL Remote data replication system for mobile computing, which enables sharing of data between a single Adaptive Server Anywhere or Adaptive Server Enterprise database and many Adaptive Server Anywhere databases using an indirect link such as e-mail or file transfer.
 - ◆ **UltraLite User's Guide** This book describes how to build database applications for small devices such as handheld organizers using the UltraLite deployment technology for Adaptive Server Anywhere databases.
 - ◆ **UltraLite User's Guide for PenRight! MobileBuilder** This book is for users of the PenRight! MobileBuilder development tool. It describes how to use UltraLite technology in the MobileBuilder programming environment.
 - ◆ **SQL Anywhere Studio Help** This book is provided online only. It includes the context-sensitive help for Sybase Central, Interactive SQL, and other graphical tools.

In addition to this documentation set, SQL Modeler and InfoMaker include their own online documentation.

Documentation formats

SQL Anywhere Studio provides documentation in the following formats:

-
- ◆ **Online books** The online books include the complete SQL Anywhere Studio documentation, including both the printed books and the context-sensitive help for SQL Anywhere tools. The online books are updated with each maintenance release of the product, and are the most complete and up-to-date source of documentation.

To access the online books on Windows operating systems, choose Start►Programs►Sybase SQL Anywhere 8►Online Books. You can navigate the online books using the HTML Help table of contents, index, and search facility in the left pane, and using the links and menus in the right pane.

To access the online books on UNIX operating systems, run the following command at a command prompt:

```
dbbooks
```

- ◆ **Printable books** The SQL Anywhere books are provided as a set of PDF files, viewable with Adobe Acrobat Reader.

The PDF files are available on the CD ROM in the *pdf_docs* directory. You can choose to install them when running the setup program.

- ◆ **Printed books** The following books are included in the SQL Anywhere Studio box:
 - ◆ *Introducing SQL Anywhere Studio.*
 - ◆ *Adaptive Server Anywhere Getting Started.*
 - ◆ *SQL Anywhere Studio Quick Reference.* This book is available only in printed form.

The complete set of books is available as the SQL Anywhere Documentation set from Sybase sales or from e-Shop, the Sybase online store, at <http://e-shop.sybase.com/cgi-bin/eshop.storefront/>.

Documentation conventions

This section lists the typographic and graphical conventions used in this documentation.

Syntax conventions

The following conventions are used in the SQL syntax descriptions:

- ◆ **Keywords** All SQL keywords are shown like the words ALTER TABLE in the following example:

ALTER TABLE [*owner*.]*table-name*

- ◆ **Placeholders** Items that must be replaced with appropriate identifiers or expressions are shown like the words *owner* and *table-name* in the following example.

ALTER TABLE [*owner*.]*table-name*

- ◆ **Repeating items** Lists of repeating items are shown with an element of the list followed by an ellipsis (three dots), like *column-constraint* in the following example:

ADD *column-definition* [*column-constraint*, ...]

One or more list elements are allowed. If more than one is specified, they must be separated by commas.

- ◆ **Optional portions** Optional portions of a statement are enclosed by square brackets.

RELEASE SAVEPOINT [*savepoint-name*]

These square brackets indicate that the *savepoint-name* is optional. The square brackets should not be typed.

- ◆ **Options** When none or only one of a list of items can be chosen, vertical bars separate the items and the list is enclosed in square brackets.

[**ASC** | **DESC**]

For example, you can choose one of ASC, DESC, or neither. The square brackets should not be typed.

- ◆ **Alternatives** When precisely one of the options must be chosen, the alternatives are enclosed in curly braces.

[**QUOTES** { **ON** | **OFF** }]

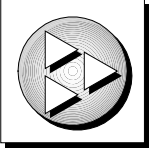
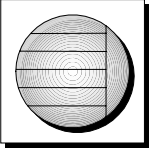
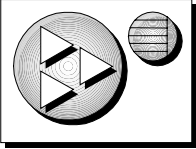
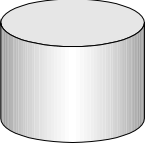

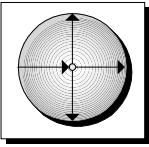

If the QUOTES option is chosen, one of ON or OFF must be provided.
The brackets and braces should not be typed.

- ◆ **One or more options** If you choose more than one, separate your choices with commas.

{ **CONNECT, DBA, RESOURCE** }

Graphic icons

The following icons are used in this documentation:

Icon	Meaning
	A client application.
	A database server, such as Sybase Adaptive Server Anywhere or Adaptive Server Enterprise.
	An UltraLite application and database server. In UltraLite, the database server and the application are part of the same process.
	A database. In some high-level diagrams, the icon may be used to represent both the database and the database server that manages it.
	Replication or synchronization middleware. These assist in sharing data among databases. Examples are the MobiLink Synchronization Server, SQL Remote Message Agent, and the Replication Agent (Log Transfer Manager) for use with Replication Server.
	A Sybase Replication Server.
	A programming interface.

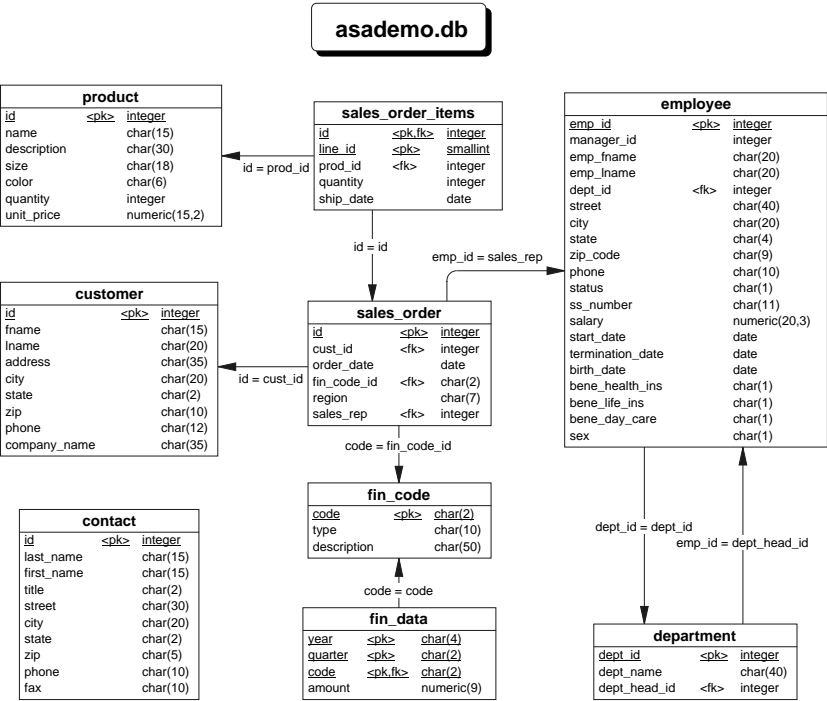
The Adaptive Server Anywhere sample database

Many of the examples throughout the documentation use the Adaptive Server Anywhere sample database.

The sample database is held in a file named *asademo.db*, and is located in your SQL Anywhere directory.

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

The following figure shows the tables in the sample database and how they relate to each other.



Finding out more and providing feedback

We would like to receive your opinions, suggestions, and feedback on this documentation.

You can provide feedback on this documentation and on the software through newsgroups set up to discuss SQL Anywhere technologies. These newsgroups can be found on the *forums.sybase.com* news server.

The newsgroups include the following:

- ◆ sybase.public.sqlanywhere.general.
- ◆ sybase.public.sqlanywhere.linux.
- ◆ sybase.public.sqlanywhere.mobilink.
- ◆ sybase.public.sqlanywhere.product_futures_discussion.
- ◆ sybase.public.sqlanywhere.replication.
- ◆ sybase.public.sqlanywhere.ultralite.

Newsgroup disclaimer

iAnywhere Solutions has no obligation to provide solutions, information or ideas on its newsgroups, nor is iAnywhere Solutions obliged to provide anything other than a systems operator to monitor the service and insure its operation and availability.

iAnywhere Solutions Technical Advisors as well as other staff assist on the newsgroup service when they have time available. They offer their help on a volunteer basis and may not be available on a regular basis to provide solutions and information. Their ability to help is based on their workload.

PART ONE

Starting and Connecting to Your Database

This part describes how to start the Adaptive Server Anywhere database server, and how to connect to your database from a client application.

C H A P T E R 1

Running the Database Server

About this chapter This chapter describes how to start and stop the Adaptive Server Anywhere database server, and the options open to you on startup under different operating systems.

Contents	Topic	Page
	Introduction	4
	Starting the server	8
	Some common options	10
	Stopping the database server	18
	Starting and stopping databases	20
	Running the server outside the current session	21
	Troubleshooting server startup	34

Introduction

Adaptive Server Anywhere provides two versions of the database server:

- ◆ **The personal database server** This executable does not support client/server communications across a network. Although the personal database server is provided for single-user, same-machine use—for example, as an embedded database engine—it is also useful for development work.

On Windows operating systems except Windows CE, the name of the personal server executable is *dbeng8.exe*. On UNIX operating systems its name is *dbeng8*.

- ◆ **The network database server** Intended for multi-user use, this executable supports client/server communications across a network.

On Windows operating systems, including Windows CE, the name of the network server executable is *dbsrv8.exe*. On Novell NetWare the name is *dbsrv8.nlm*, and on UNIX operating systems it is *dbsrv8*.

Server differences

The request-processing engine is identical in the two servers. Each supports exactly the same SQL, and exactly the same database features. The main differences include:

- ◆ **Network protocol support** Only the network server supports communications across a network.
- ◆ **Number of connections** The personal server has a limit of ten simultaneous connections. The limit for the network server depends on your license.
- ◆ **Number of CPUs** The personal database server uses a maximum of two CPUs for request processing. The network database server has no set limit.
- ◆ **Default number of internal threads** You can configure the number of requests the server can process at one time using the *-gn* option. The network database server has a default of 20 threads and no set limit, while the personal database server has a default and limit of 10 threads.

☞ For more information about database server options, see "The database server" on page 120.

- ◆ **Startup defaults** To reflect their use as a personal server and a server for many users, the startup defaults are slightly different for each.

First steps

You can start a personal server running on a single database very simply. For example, you can start both a personal server and a database called *test.db* by typing the following command in the directory where *test.db* is located:

```
dbeng8 test
```

Where to specify commands

You can specify commands in several ways, depending on your operating system. For example, you can:

- ◆ type the command at the command prompt.
- ◆ place the command in a shortcut or desktop icon.
- ◆ run the command in a batch file.
- ◆ include the command as a **StartLine (START)** connection parameter in a connection string.

🔗 For more information, see "StartLine connection parameter" on page 187.

There are slight variations in the basic command from platform to platform. These are described in the following section.

🔗 You can also start a personal server using a database file name in a connection string. For more information, see "Connecting to an embedded database" on page 47.

Start the database server

The way you start the database server varies slightly depending on the operating system you use. This section describes how to specify commands for the simple case of running a single database with default settings, on each supported operating system.

Notes

- ◆ These commands start the personal server (**dbeng8**). To start a network server, simply replace **dbeng8** with **dbsrv8**.
- ◆ If the database file is in the starting directory for the command, you do not need to specify *path*.
- ◆ If you do not specify a file extension in *database-file*, the extension *.db* is assumed.

❖ To start the personal database server using default options (Windows):

- 1 Open a command prompt.
- 2 Type the following command:

```
start dbeng8 path\database-file
```

If you omit the database file, a window appears where you can locate a database file using a Browse button.

❖ **To start the personal database server using default options (UNIX):**

- 1 Open a command prompt.
- 2 Type the following command:

```
dbeng8 path/database-file
```

❖ **To start the personal database server using default options (NetWare):**

- ◆ The database server for NetWare is a NetWare Loadable Module (NLM) (*dbsrv8.nlm*). A NLM is a program that you can run on your NetWare server. Load a database server on your NetWare server as follows:

```
load dbsrv8.nlm path\database-file
```

The database file must be on a NetWare volume. A typical filename is of the form *DB:\database\sales.db*.

You can load the server from a client machine using the Novell remote console utility. See your Novell documentation for details.

You can put the command into your Novell *autoexec.ncf* file so that Adaptive Server Anywhere loads automatically each time you start the NetWare server.

There is no personal server for Novell NetWare, just a network server.

What else is there to it?

Although you can start a personal server in the simple way described above, there are many other aspects to running a database server in a production environment. For example,

- ◆ You can choose from many **options** to specify such features as how much memory to use as cache, how many CPUs to use (on multi-processor machines), and the network protocols to use (network server only). The options are one of the major ways of tuning Adaptive Server Anywhere behavior and performance.
- ◆ You can run the server as a Windows **service**. This allows it to continue running even when you log off the machine.

- ◆ You can start the personal server from an application and shut it down when the application has finished with it. This is typical when using the database server an **embedded database**.

The remainder of this chapter describes these options in more detail.

Starting the server

The general form for the server command is as follows:

```
executable [ server-options ] [ database-file [ database-options ], ...]
```

If you supply no options and no database file, then on Windows operating systems a dialog appears, allowing you to use a Browse button to locate your database file.

The elements of the database server command include the following:

- ◆ **Executable** This can be either the personal server or the network server.

✍ For the file names on different operating systems, see "Introduction" on page 4.

In this chapter, unless discussing network-specific options, we use the personal server in sample commands. The network server takes a very similar set of options.

- ◆ **Server options** These options control the behavior of the database server for all running databases.
- ◆ **Database file** You can specify zero, one, or more database file names. Each of these databases starts and remains available for applications.

Caution

The database file and the transaction log file must be located on the same physical machine as the database server. Database files and transaction log files located on a network drive can lead to poor performance and data corruption.

- ◆ **Database options** For each database file you start, you can provide database options that control certain aspects of its behavior.

In this section, we look at some of the more important and commonly-used options.

✍ For full reference information on each of these options, see "The database server" on page 120.

In examples throughout this chapter where there are several options, we display them on separate lines for clarity, as they could be written in a configuration file. If you type them directly at the command prompt, you must type them all on one line.

Case sensitivity

Database and server options are generally case sensitive. You should type all options in lower case.

Listing available options

❖ **To list the database server options:**

- 1 Open a command prompt.
- 2 Type the following command:
`dbeng8 -?`

Starting the database server on Windows CE

The database server supplied for Windows CE is the network database server (*dbsrv8.exe*). The network server supports communications over a TCP/IP network link.

The usual client/server arrangement has the database server running on a machine with more power and resources than the client applications. Clearly, this is not the case with Windows CE; instead, the less powerful machine is running the database server.

The advantage to supplying a network server on Windows CE is that you can run database applications on your desktop computer to carry out tasks on your Windows CE database. For example:

- ◆ You can use Sybase Central on your desktop PC to manage your database.
- ◆ You can use Interactive SQL on your desktop to load and unload data, and carry out queries.
- ◆ You can use InfoMaker to produce reports.

The Windows CE database server does not start the TCP/IP network link unless it is explicitly requested:

```
dbsrv8 -x tcpip ...
```

Some common options

This section describes some of the most common options, and points out when you may wish to use them. They are:

- ◆ Using configuration files
- ◆ Naming the server and the databases
- ◆ Performance
- ◆ Permissions
- ◆ Maximum page size
- ◆ Special modes
- ◆ Threading
- ◆ Network communications (network server only)

Using configuration files

If you use an extensive set of options, you can store them in a configuration file and invoke that file in a server command. The configuration file can contain options on several lines. For example, the following configuration file starts the personal database server and the sample database. It sets a cache of 10 Mb, and names this instance of the personal server **Elora**. Lines with # as the first comment in the line are treated as comments.

```
# Configuration file for server Elora
-n Elora
-c 10M
path\asademo.db
```

In the example, *path* is the name of your SQL Anywhere directory. On UNIX, you would use a forward slash instead of the backslash in the file path.

If you name the file *sample.cfg*, you could use these options as follows:

```
dbeng8 @sample.cfg
```

☞ For more information, see "@filename server option" on page 126.

Naming the server and the databases

You can use the `-n` as a server option (to name the server) or as a database option (to name the database).

The server and database names are among the connection parameters that client applications may use when connecting to a database. The server name appears on the desktop icon and in the title bar of the server window.

Naming databases

You may want to provide a database name to provide a more meaningful name than the file name for users of client applications. The database is identified by that name until it is stopped.

If you don't provide a database name, the default name is the root of the database file name (the file name without the *.db* extension). For example, in the following command the first database is named **asademo**, and the second **sample**.

```
dbeng8 asademo.db sample.db
```

You can name databases by supplying a **-n** option following the database file. For example, the following command starts the sample database and names it **MyDB**:

```
dbeng8 asademo.db -n MyDB
```

Naming the server

Providing a database server name helps avoid conflicts with other server names on your network. It also provides a meaningful name for users of client applications. The server keeps its name for its lifetime (until it is shut down). If you don't provide a server name, the server is given the name of the first database started.

You can name the server by supplying a **-n** option before the first database file. For example, the following command starts a server on the **asademo** database and gives the server the name **Cambridge**:

```
dbeng8 -n Cambridge asademo.db
```

If you supply a server name, you can start a database server with no database started. The following command starts a server named **Galt** with no database started:

```
dbeng8 -n Galt
```

☞ For more information about starting databases on a running server, see "Starting and stopping databases" on page 20.

Case sensitivity

Server names and database names are case insensitive as long as the character set is single-byte.

☞ For more information, see "Connection strings and character sets" on page 279.

Controlling performance and memory from the command line

Several options can have a major impact on database server performance, including:

- ◆ **Cache size** The `-c` option controls the amount of memory that Adaptive Server Anywhere uses as a cache. This can be a major factor in affecting performance.

Generally speaking, the more memory made available to the database server, the faster it performs. The cache holds information that may be required more than once. Accessing information in cache is many times faster than accessing it from disk. The default initial cache size is computed based on the amount of physical memory, the operating system, and the size of the database files. On Windows operating systems, and UNIX the database server takes additional cache when the available cache is exhausted.

🔗 For more information about performance tuning, see "Monitoring and Improving Performance" on page 143 of the book *ASA SQL User's Guide*.

🔗 For information on controlling cache size, see "Cache size" on page 124.

- ◆ **Number of processors** If you are running on a multi-processor machine, you can set the number of processors with the `-gt` option.

🔗 For more information, see "`-gt` server option" on page 143 and "Controlling threading from the command line" on page 14.

- ◆ **Other performance-related options** There are several options available for tuning network performance, including `-gb` (database process priority), and `-u` (buffered disk I/O).

🔗 For more information about startup options, see "The database server" on page 120.

Controlling permissions from the command line

Some options control the permissions required to carry out certain global operations, including permissions to start and stop databases, load and unload data, and create and delete database files.

🔗 For more information, see "Running the database server in a secure fashion" on page 397.

Setting a maximum page size

The database server cache is arranged in **pages**—fixed-size areas of memory. Since the server uses a single cache for its lifetime (until it is shut down), all pages must have the same size.

A database file is also arranged in pages, with a size that is specified on the command line. Every database page must fit into a cache page. By default, the server page size is the same as the largest page size of the databases on the command line. Once the server starts, you cannot start a database with a larger page size than the server.

To allow databases with larger page sizes to be started after startup, you can force the server to start with a specified page size using the `-gp` option. If you use larger page sizes, remember to increase your cache size. A cache of the same size will accommodate only a fraction of the number of the larger pages, leaving less flexibility in arranging the space.

The following command starts a server that reserves an 8 Mb cache and can accommodate databases of page sizes up to 4096 bytes.

```
dbsrv8 -gp 4096 -c 8M -n myserver
```

Running in special modes

You can run Adaptive Server Anywhere in special modes for particular purposes.

- ◆ **Read-only** You can run databases in read-only mode by supplying the `-r` option.

🔗 For more information, see "`-r` server option" on page 149.

- ◆ **Bulk load** This is useful when loading large quantities of data into a database using the Interactive SQL INPUT command. Do not use the `-b` option if you are using LOAD TABLE to bulk load data.

🔗 For more information, see "`-b` server option" on page 127, and "Importing and Exporting Data" on page 421 of the book *ASA SQL User's Guide*.

- ◆ **Starting without a transaction log** Use the `-f` database option for recovery—either to force the database server to start after the transaction log has been lost, or to force the database server to start using a transaction log it would otherwise not find. Note that `-f` is a database option, not a server option.

Once the recovery is complete, you should stop your server and restart without the `-f` option.

☞ For more information, see "The database server" on page 120.

Controlling threading from the command line

Adaptive Server Anywhere can use multiple operating system threads to handle requests. This allows different requests to run simultaneously on separate CPUs. Each request runs on a single thread and no request is executed concurrently on multiple CPUs.

☞ For a list of options that control threading, see "Controlling threading behavior" on page 15.

Threading in Adaptive Server Anywhere

To understand how threading support works, you need to understand **requests** and **tasks**.

Requests

Suppose Adaptive Server Anywhere is being used concurrently by two users, or by two connections from a single application. Each connection submits a query (or other SQL statement) to Adaptive Server Anywhere. Each of these SQL statements is a separate request to the server.

Tasks

A task picks up a request and handles that request until it is complete. On Windows NT/2000/XP, tasks are lightweight threads called fibers; on all other platforms, tasks are threads. The number of tasks determines the number of requests that Adaptive Server Anywhere can handle concurrently.

Tasks on Windows 95/98/Me, NetWare, and UNIX

On Windows 95/98/Me, NetWare, and UNIX machines, each task is an operating system thread. When Adaptive Server Anywhere receives a request, an operating system thread picks up the request. The operating system thread runs until the request is complete.

If the server receives a request while the first request is running, the second request is assigned to a different thread. If a thread is blocked while completing a task, it does not pick up another request; instead, it waits until it can complete the current request.

Tasks on Windows NT/XP/2000

On Windows NT/2000/XP, tasks are lightweight threads called fibers. When Adaptive Server Anywhere receives a request, a fiber picks up the request and runs until the request is complete. If the database server receives another request while the first request is running, the second request is assigned to a different fiber. If a fiber blocks while processing a request, it yields control to another fiber until it can proceed: it does not pick up another request.

Threads host fibers, and the fibers are scheduled co-operatively. A fiber must explicitly yield control to another fiber when it is waiting to process a request, for example, while waiting for an I/O operation to complete. If a fiber blocks and does not yield control, it blocks the thread that is hosting it and prevents other fibers from running on that thread. If more than one thread is hosting fibers, only the thread that is hosting the waiting fiber is blocked: other threads are still free to run fibers.

You can set the number of threads used to run fibers with the `-gx` option. Unless you are using Java or Remote Data Access, you should need only one thread per CPU.

Setting the number of threads when using Java or Remote Data Access

When you use Java or Remote Data Access and the fiber running Java or Remote Data Access blocks, that fiber may not yield control to another fiber, which in turn blocks the thread. For this reason, the number of operating system threads assigned to the database server defaults to one more thread than the number of CPUs on the machine. This ensures that there is at least one thread available to host fibers if the thread being used by Java or Remote Data Access is blocked. Specifying a higher number of threads than the default by using the `-gx` option has a minimal impact on performance.

Controlling threading behavior

There are four database server options that control threading behavior. Not all of these options are required on every platform.

- ◆ **Number of tasks** The `-gn` option controls the number of tasks used to process requests. Effectively, this is the number of requests that can be handled concurrently. Each request uses a task. When there are more requests than there are tasks, any outstanding requests must wait until a currently-running task completes. By default, there are 20 tasks for the network database server and 10 tasks for the personal database server. There is no benefit to setting a number of tasks that is greater than the maximum number of server connections.

For more information, see "`-gn` server option" on page 141.

- ◆ **Stack size per internal execution thread** You can set the stack size per internal execution thread in the server using the `-gss` option. The `-gss` option allows you to lower the memory usage of the database server in environments with limited memory. This option has no effect on Windows operating systems.

☞ For more information, see "`-gss` server option" on page 143.

- ◆ **Number of processors** If you have more than one processor, you can control how many processors the threads exploit by specifying the `-gt` option. By default, all processors available on the machine are used.

☞ For more information, see "`-gt` server option" on page 143.

- ◆ **Number of threads assigned to the database server process** On Windows NT/2000/XP, the `-gx` option controls the number of threads that are dedicated to hosting fibers to service requests. By default, this is set to one more than the number of CPUs on the machine. On UNIX, NetWare, and Windows 95/98/Me where each task is its own thread, this option is not needed.

☞ For more information, see "`-gx` server option" on page 144.

Selecting communications protocols

Any communication between a client application and a database server requires a communications protocol. Adaptive Server Anywhere supports a set of communications protocols for communications across networks and for same-machine communications.

By default, the database server starts up all available protocols. You can limit the protocols available to a database server using the `-x` option. On the client side, many of the same options can be controlled using the **CommLinks (LINKS)** connection parameter.


☞ For more information on running the server using these options, see "Supported network protocols" on page 92.

Available protocols
for the personal
server

The personal database server (*dbeng8.exe*) supports the following protocols:

- ◆ **Shared memory** This protocol is for same-machine communications, and always remains available. It is available on all platforms.
- ◆ **TCP/IP** This protocol is for same-machine communications only, from TDS clients, Open Client, or the jConnect JDBC driver. You must run TCP/IP if you wish to connect from Open Client or jConnect.

☞ For more information on TDS clients, see "Adaptive Server Anywhere as an Open Server" on page 105.


Available protocols for the network server	<ul style="list-style-type: none">◆ Named Pipes This protocol is provided on Windows NT only. Named Pipes is for same machine communications for applications that wish to run under a certified security environment. <p>The network database server (<i>dbsrv8.exe</i>) supports the following protocols:</p> <ul style="list-style-type: none">◆ Shared memory This protocol is for same-machine communications, and always remains available. It is available on all platforms.◆ SPX This protocol is supported on all platforms except for UNIX.◆ TCP/IP This protocol is supported on all platforms.◆ Named Pipes This protocol is supported on Windows NT only. Named Pipes is for same machine communications for applications that wish to run under a certified security environment.
Specifying protocols	<p>You can instruct a server to use only some of the available network protocols when starting up, using the <code>-x</code> option. The following command starts the asademo database using the TCP/IP and SPX protocols:</p> <pre>dbsrv8 -x "tcpip,spx" path\asademo.db</pre> <p>Although not strictly required in this example, the quotes are necessary if there are spaces in any of the arguments to <code>-x</code>.</p> <p>You can add additional parameters to tune the behavior of the server for each protocol. For example, the following command (typed all on one line) instructs the server to use two network cards, one with a specified port number.</p> <pre>dbsrv8 -x "tcpip{MyIP=192.75.209.12:2367,192.75.209.32}" path\asademo.db</pre> <p> For more information about available network communications parameters that can serve as part of the <code>-x</code> option, see "Network communications parameters" on page 189.</p>

Stopping the database server

You can stop the database server by:

- ◆ Clicking Shutdown on the database server window.
- ◆ Using the *dbstop* utility.

The *dbstop* utility is particularly useful in batch files, or for stopping a server on another machine. It requires a connection string in its command.

 For more information on *dbstop* options, see "Stopping a database server using the dbstop command-line utility" on page 505.

- ◆ Letting it shut down automatically by default when the application disconnects. (This only works if the server is a personal server started by an application connection string.)
- ◆ Pressing Q when the server display window has the focus on UNIX or NetWare machines.

Examples

❖ To stop a server using the dbstop utility:

- 1 Start a server. For example, the following command executed from the Adaptive Server Anywhere installation directory starts a server named Ottawa using the sample database:

```
dbsrv8 -n Ottawa asademo.db
```

- 2 Stop the server using *dbstop*:

```
dbstop -c "eng=Ottawa;uid=DBA;pwd=SQL"
```

Who can stop the server?

When you start a server, you can use the *-gk* option to set the level of permissions required for users to stop the server with *dbstop*. The default level of permissions required is **DBA**, but you can also set the value to **all** or **none**. (Interactively, of course, anybody at the machine can click Shutdown on the server window.)

Shutting down operating system sessions

If you close an operating system session where a database server is running, or if you use an operating system command to stop the database server, the server shuts down, but not cleanly. The next time the database loads, recovery is required, and happens automatically.

☞ For more information about recovery, see "Backup and Data Recovery" on page 299.

It is better to stop the database server explicitly before closing the operating system session. On NetWare, however, shutting down the NetWare server machine properly does stop the database server cleanly.

Examples of commands that will not stop a server cleanly include:

- ◆ Stopping the process in the Windows Task Manager.
- ◆ Using a UNIX **slay** or **kill** command.

Starting and stopping databases

A database server can have more than one database loaded at a time. You can start databases and start the server at the same time, as follows:

```
dbeng8 asademo sample
```

Starting a database on a running server

You can also start databases after starting a server in one of the following ways:

- ◆ Connect to a database using a **DatabaseFile (DBF)** connection parameter while connected to a server. The **DatabaseFile (DBF)** connection parameter specifies a database file for a new connection. The database file is started on the current server.

🔗 For more information, see "Connecting to an embedded database" on page 47, or "DatabaseFile connection parameter" on page 173.

- ◆ Use the **START DATABASE** statement, or choose Start Database from the File menu in Sybase Central when you have a server selected.

🔗 For more information, see "START DATABASE statement" on page 549 of the book *ASA SQL Reference Manual*.

Limitations

- ◆ The server holds database information in memory using pages of a fixed size. Once a server has been started, you cannot start a database that has a larger page size than the server.
- ◆ The `-gd` server option decides the permissions required to start databases.

Stopping a database

You can stop a database by:

- ◆ Disconnecting from a database started by a connection string. Unless you explicitly set the **AutoStop (ASTOP)** connection parameter to *NO*, this happens automatically.

🔗 For more information, see "AutoStop connection parameter" on page 167.

- ◆ Using the **STOP DATABASE** statement from Interactive SQL or Embedded SQL.

🔗 For more information, see "STOP DATABASE statement" on page 558 of the book *ASA SQL Reference Manual*.

Running the server outside the current session

When you log on to a computer using a user ID and a password, you establish a **session**. When you start a database server, or any other application, it runs within that session. When you log off the computer, all applications associated with the session terminate.

It is common to require database servers to be available all the time. To make this easier, you can run Adaptive Server Anywhere for Windows NT/2000/XP and for UNIX in such a way that, when you log off the computer, the database server remains running. The way you do this depends on your operating system.

- ◆ **UNIX daemon** You can run the UNIX database server as a daemon using the `-ud` option, enabling the database server to run in the background, and to continue running after you log off.

✍ For more information, see "Running the UNIX database server as a daemon" on page 21.

- ◆ **Windows service** You can run the Windows database server as a service. This has many convenient properties for running high availability servers.

✍ For more information, see "Understanding Windows services" on page 23.

Running the UNIX database server as a daemon

To run the UNIX database server in the background, and to enable it to run independently of the current session, you run it as a **daemon**.

Do not use '&' to run the database server in the background

If you use the UNIX `&` (ampersand) command to run the database server in the background, it will not work—the server will hang. You must instead run the database server as a daemon.

As well, attempting to start a server in the background from within a program using the typical `fork()-exec()` sequence will not work.

You can run the UNIX database server as a daemon in one of the following ways:

- 1 Use the `-ud` option when starting the database server. For example:

```
dbsrv8 -ud asademo
```

- 2 Use the *dbspawn* tool to start the database server, for example:

```
dbspawn dbsrv8 asademo
```

One advantage of using *dbspawn* is that the *dbspawn* process does not terminate until it has confirmed that the daemon has started and is ready to accept requests. If for any reason the daemon fails to start, the exit code for *dbspawn* will be non-zero.

When you start the daemon directly using the *-ud* option, the *dbeng8* and *dbsrv8* commands create the daemon process and return immediately (exiting and allowing the next command to be executed) before the daemon initializes itself or attempts to open any of the databases specified in the command.

Using *dbspawn* can be useful when writing a script or other automated process that needs to start a server as a daemon prior to running one or more applications that will use the server because you want to ensure that the daemon is running before starting the applications. The following is an example of how to test this using a csh script.

```
#!/bin/csh
# start the server as a daemon and ensure that it is
# running before we start any applications
dbspawn dbsrv8 asademo
if ( $status != 0 ) then
    echo Failed to start asademo server
    exit
endif
# ok, now we can start the applications
...
```

This example uses an sh script to test whether the daemon is running before starting the applications.

```
#!/bin/sh
# start the server as a daemon and ensure that it is
# running before we start any applications
dbspawn dbsrv8 asademo
if [ $? != 0 ]; then
    echo Failed to start asademo server
    exit
fi
# ok, now we can start the applications
...
```


- 3 Spawn a daemon from within a C program, for example:

```
...
if( fork() == 0 ) {
    /* child process = start server daemon */
    execl( "/opt/sybase/SYBSsa8/bin/dbsrv8",
        "dbsrv8", "-ud", "asademo" );
}
```



```
        exit(1);  
    }  
    /* parent process */  
    ...
```

Note that the `-ud` option is used.

 For more information, see "`-ud` server option" on page 153 and "The Spawn utility" on page 503.

Understanding Windows services

Limitations of running as a standard executable	<p>Although you can run the database server like any other Windows NT/2000/XP program rather than as a service, there are limitations to running it as a standard program, particularly in multi-user environments.</p> <p>When you start a program, it runs under your Windows NT/2000/XP login session, which means that if you log off the computer, the program terminates. Only one person logs onto Windows NT/2000/XP (on any one computer) at one time. This restricts the use of the computer if you wish to keep a program running much of the time, as is commonly the case with database servers. You must stay logged onto the computer running the database server for the database server to keep running. This can also present a security risk as the Windows NT/2000/XP computer must be left in a logged on state.</p>
Advantages of services	<p>Installing an application as a Windows service enables it to run even when you log off.</p> <p>When you start a service, it logs on using a special system account called LocalSystem (or using another account that you specify). Since the service is not tied to the user ID of the person starting it, the service remains open even when the person who started it logs off. You can also configure a service to start automatically when the Windows computer starts, before a user logs on.</p>
Managing services	<p>Sybase Central provides a more convenient and comprehensive way of managing Adaptive Server Anywhere services than the Windows services manager.</p>

Programs that can be run as Windows services

You can run the following programs as services:

- ◆ Network Database Server (*dbsrv8.exe*)
- ◆ Personal Database Server (*dbeng8.exe*)
- ◆ SQL Remote Message Agent (*dbremote.exe*)

- ◆ The MobiLink Synchronization server (*dbmlsrv8.exe*)
- ◆ A sample application

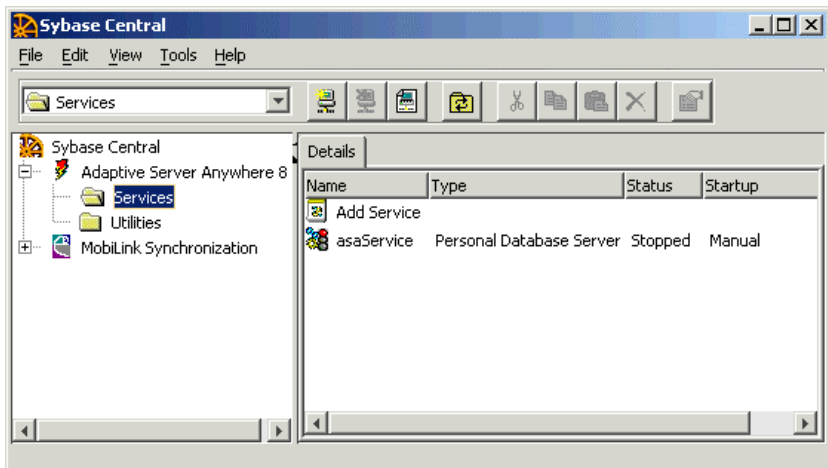
Not all these applications are supplied in all editions of Adaptive Server Anywhere.

Managing services

You can carry out the following service management tasks from the command line, or in the Services folder in Sybase Central:

- ◆ Add, edit, and remove services.
- ◆ Start, stop, and pause services.
- ◆ Modify the parameters governing a service.
- ◆ Add databases to a service so you can run several databases at one time.

The service icons in Sybase Central display the current state of each service using a traffic light icon (running, paused, or stopped).



Adding a service

This section describes how to set up services using Sybase Central and the Service Creation utility.

❖ To add a new service (Sybase Central):

- 1 In Sybase Central, open the Services folder.

- 2 Double-click Add Service.
- 3 Follow the instructions in the wizard.

❖ **To add a new service (Command line):**

- 1 Open a command prompt.
- 2 Execute the Service Creation utility using the `-w` option.

For example, to create a personal server service called `myserv`, which starts the specified engine with the specified parameters type the following. The engine runs as the `LocalSystem` user:

```
dbsvc -as -w myserv "C:\Program Files\Sybase\SQL
Anywhere 8\win32\dbeng8.exe" -n william -c 8m
"C:\Program Files\Sybase\SQL Anywhere 8\sample.db"
```

☞ For more information about the Service Creation utility and options, see "The Service Creation utility" on page 499.

Notes

- ◆ Service names must be unique within the first eight characters.
- ◆ If you choose to start a service automatically, it starts whenever the computer starts Windows. If you choose to start the service manually, you need to start the service from Sybase Central each time. You may want to select Disabled if you are setting up a service for future use.
- ◆ Type options for the executable, without the executable name itself, in the window. For example, if you want a network server to run using the sample database with a cache size of 20 Mb and the name **myserver**, you would type the following in the Parameters box of the Service Creation wizard in Sybase Central:

```
-c 20M
-n myserver c:\Program Files\Sybase\SQL
Anywhere 8\asademo.db
```


Line breaks are optional.

☞ For information on valid options, see the description of each program in "Database Administration Utilities" on page 435.

- ◆ Choose the account under which the service will run: the special `LocalSystem` account or another user ID.

☞ For more information about this choice, see "Setting the account options" on page 28.

- ◆ If you want the service to be accessible from the Windows desktop, check Allow Service to Interact with Desktop. If this option is cleared, no icon or window appears on the desktop.

 For more information on the configuration options, see "Configuring services" on page 26.

Removing a service

Removing a service removes the server name from the list of services. Removing a service does not remove any software from your hard disk.

If you wish to re-install a service you previously removed, you need to re-type the options.

❖ To remove a service (Sybase Central):


- 1 In Sybase Central, open the Services folder.
- 2 In the right pane, right-click the icon of the service you want to remove and choose Delete from the popup menu.

❖ To remove a service (Command line):

- 1 Open a command prompt.
- 2 Execute the Service Creation utility using the `-d` option.


For example, to delete the service called `myserv`, without prompting for confirmation, type the following command:

```
dbsvc -y -d myserv
```

 For more information about the Service Creation utility and options, see "The Service Creation utility" on page 499.

Configuring services

A service runs a database server or other application with a set of options.

 For a full description of the options for each of the administration utilities, see "Database Administration Utilities" on page 435.

In addition to the options, services accept other parameters that specify the account under which the service runs and the conditions under which it starts.

❖ To change the parameters for a service:

- 1 In Sybase Central, open the Services folder.

- 2 In the right pane, right-click the service you want to change and choose Properties from the popup menu.
- 3 Alter the parameters as needed on the tabs of the Service property sheet.
- 4 Click OK when finished.

Changes to a service configuration take effect the next time someone starts the service. The Startup option is applied the next time Windows is started.

Setting the startup option

The following options govern startup behavior for Adaptive Server Anywhere services. You can set them on the General tab of the Service property sheet.

- ◆ **Automatic** If you choose the Automatic setting, the service starts whenever the Windows operating system starts. This setting is appropriate for database servers and other applications running all the time.
- ◆ **Manual** If you choose the Manual setting, the service starts only when a user with Administrator permissions starts it. For information about Administrator permissions, see your Windows documentation.
- ◆ **Disabled** If you choose the Disabled setting, the service will not start.

Specifying options

The Configuration tab of the Service property sheet provides a Parameters text box for specifying options for a service. Do not type the name of the program executable in this box.

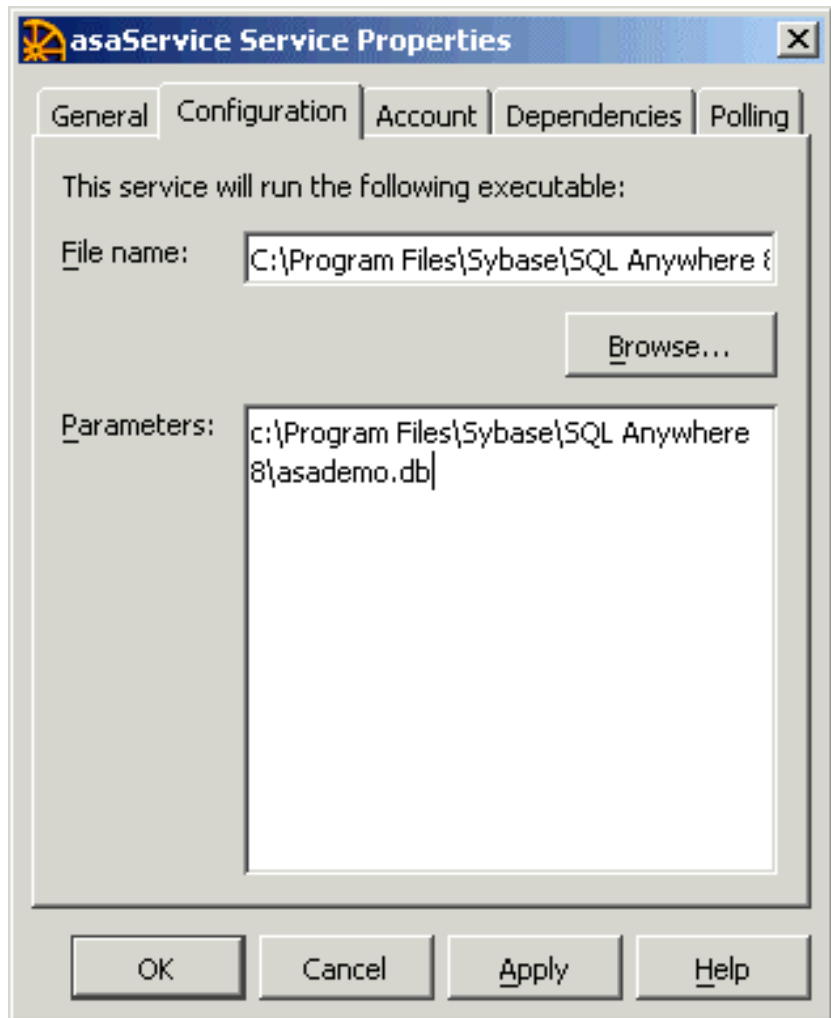
Examples

- ◆ To start a network server service running two databases, with a cache size of 20 Mb, and with a name of **my_server**, you would type the following in the Parameters box:

```
-c 20M
-n my_server
c:\Program Files\Sybase\SQL Anywhere 8\db_1.db
c:\Program Files\Sybase\SQL Anywhere 8\db_2.db
```
- ◆ To start a SQL Remote Message Agent service, connecting to the sample database as user ID DBA, you would type the following:

```
-c "uid=DBA;pwd=SQL;dbn=asademo"
```

The following figure illustrates a sample property sheet.



☞ The options for a service are the same as those for the executable. For a full description of the options for each program, see "The Database Server" on page 119.

Setting the account options

You can choose under which account the service runs. Most services run under the special LocalSystem account, which is the default option for services. You can set the service to log on under another account by opening the Account tab on the Service property sheet, and typing the account information.

If you choose to run the service under an account other than LocalSystem, that account must have the "Log On As A Service" privilege. This can be granted from the Windows User Manager application under Advanced Privileges.

When an icon appears on the taskbar

Whether or not an icon for the service appears on the taskbar or desktop depends on the account you select, and whether Allow Service to Interact with Desktop is selected, as follows:

- ◆ If a service runs under LocalSystem, and Allow Service to Interact with Desktop is selected on the Service property sheet, an icon appears on the desktop of every user logged in to Windows on the computer running the service. Consequently, any user can open the application window and stop the program running as a service.
- ◆ If a service runs under LocalSystem, and Allow Service to Interact with Desktop is cleared on the Service property sheet, no icon appears on the desktop for any user. Only users with permissions to change the state of services can stop the service.
- ◆ If a service runs under another account, no icon appears on the desktop. Only users with permissions to change the state of services can stop the service.

Changing the executable file

To change the program executable file associated with a service, click the Configuration tab on the Service property sheet and type the new path and file name in the File Name text box.

If you move an executable file to a new directory, you must modify this entry.

Adding new databases to a service

Each network server or personal server can run more than one database. If you wish to run more than one database at a time, we recommend that you do so by attaching new databases to your existing service, rather than by creating new services.

❖ To add a new database to a service:

- 1 Open the Services folder.
- 2 Right-click the service and choose Properties from the popup menu.
- 3 Click the Configuration tab.

- 4 Add the path and filename of the new database to the end of the list of options in the Parameters box.
- 5 Click OK to save the changes.

The new database starts the next time the service starts.

Databases can be started on running servers by client applications, such as Interactive SQL.

🔗 For more information about how to start a database on a server from Interactive SQL, see "START DATABASE statement" on page 549 of the book *ASA SQL Reference Manual*.

🔗 For more information about how to implement this function in an Embedded SQL application, see "db_start_database function" on page 239 of the book *ASA Programming Guide*.

Starting a database from an application does not attach it to the service. If the service is stopped and restarted, the additional database will not be started automatically.

Setting the service polling frequency

Sybase Central can poll at specified intervals to check the state (started, stopped, or paused) of each service, and update the icons to display the current state. By default, polling is off. If you leave it off, you must click Refresh to see changes to the state.

❖ To set the Sybase Central polling frequency:

- 1 Open the Services folder.
- 2 Right-click the service and choose Properties from the popup menu.
- 3 Click the Polling tab.
- 4 Select Enable Polling.
- 5 Set the polling frequency.

The frequency applies to all services, not just the one selected. The value you set in this window remains in effect for subsequent sessions, until you change it.

- 6 Click OK.

Starting, stopping, and pausing services

❖ To start, stop, or pause a service:

- 1 Open the Services folder.
- 2 Right-click the service and choose Start, Stop, or Pause from the popup menu.

To resume a paused service, right-click the service and choose Continue from the popup menu.

If you start a service, it keeps running until you stop it. Closing Sybase Central or logging off does not stop the service.

Stopping a service closes all connections to the database and stops the database server. For other applications, the program closes down.

Pausing a service prevents any further action being taken by the application. It does not shut the application down or (in the case of server services) close any client connections to the database. Most users do not need to pause their services.

The Windows Service Manager

You can use Sybase Central to carry out all the service management for Adaptive Server Anywhere. Although you can use the Windows NT Service Manager in the Control Panel for some tasks, you cannot install or configure an Adaptive Server Anywhere service from the Windows NT Service Manager.

If you open the Windows Service Manager (from the Windows Control Panel), a list of services appears. The names of the Adaptive Server Anywhere services are formed from the Service Name you provided when installing the service, prefixed by Adaptive Server Anywhere. All the installed services appear together in the list.

Running more than one service

This section describes some topics specific to running more than one service at a time.

Service dependencies

In some circumstances you may wish to run more than one executable as a service, and these executables may depend on each other. For example, you may wish to run a server and a SQL Remote Message Agent or Log Transfer Manager to assist in replication.

In cases such as these, the services must start in the proper order. If a SQL Remote Message Agent service starts up before the server has started, it fails because it cannot find the server.

You can prevent these problems by using **service groups**, which you manage from Sybase Central.

Service groups overview

You can assign each service on your system to be a member of a service group. By default, each service belongs to a group, as listed in the following table.

Service	Default group
Network server	ASANYServer
Personal server	ASANYEngine
SQL Remote Message Agent	ASANYRemote
MobiLink Synchronization Server	ASANYMobiLink
Replication Agent	ASANYLTM

Before you can configure your services to ensure that they start in the correct order, you must check that your service is a member of an appropriate group. You can check which group a service belongs to, and change this group, from Sybase Central.

❖ To check and change which group a service belongs to:

- 1 Open the Services folder.
- 2 Right-click the service and choose Properties from the popup menu.
- 3 Click the Dependencies tab. The top text box displays the name of the group the service belongs to.
- 4 Click Change to display a list of available groups on your system.
- 5 Select one of the groups, or type a name for a new group.
- 6 Click OK to assign the service to that group.

Managing service dependencies

With Sybase Central you can specify **dependencies** for a service. For example:

- ◆ You can ensure that at least one member of each of a list of service groups has started before the current service.
- ◆ You can ensure that any number of services start before the current service. For example, you may want to ensure that a particular network server has started before a SQL Remote Message Agent that is to run against that server starts.

❖ To add a service or group to a list of dependencies:

- 1 Open the Services folder.
- 2 Right-click the service and choose Properties from the popup menu.
- 3 Click the Dependencies tab.
- 4 Click Add Services or Add Service Groups to add a service or group to the list of dependencies.
- 5 Select one of the services or groups from the list.
- 6 Click OK to add the service or group to the list of dependencies.

Troubleshooting server startup

This section describes some common problems that may occur when starting the database server.

Ensure that your transaction log file is valid

The server won't start if the existing transaction log is invalid. For example, during development you may replace a database file with a new version, without deleting the transaction log at the same time. This causes the transaction log file to be different than the database, and results in an invalid transaction log file.

Ensure that you have sufficient disk space for your temporary file

Adaptive Server Anywhere uses a temporary file to store information while running. This file is stored in the directory pointed to by the TMP or TEMP environment variable, typically *c:\temp*.

If you do not have sufficient disk space available to the temporary directory, you will have problems starting the server.

Ensure that network communication software is running

Appropriate network communication software must be installed and running before you run the database server. If you are running reliable network software with just one network installed, this should be straightforward. If you experience problems, if you are running non-standard software, or if you are running multiple networks, you may want to read the full discussion of network communication issues in "Client/Server Communications" on page 91.

You should confirm that other software requiring network communications is working properly before running the database server.

If you are running under the TCP/IP protocol, you may want to confirm that ping and telnet are working properly. The ping and telnet applications are provided with many TCP/IP protocol stacks.

Debugging network communications startup problems

If you are having problems establishing a connection across a network, you can use debugging options at both the client and the server to diagnose problems. On the server, you use the `-z` option. The startup information appears on the database server window: you can use the `-o` option to log the results to an output file.

Make sure you're using the right *asasrv.ini* file

If you are having problems establishing a connection to the correct server across a network, try deleting the *asasrv.ini* file. This file contains server information, including server name, protocol and address. It is possible that the server information in this file is overriding information you specified in the connection string. Deleting this file causes Adaptive Server Anywhere to create a new *asasrv.ini* file containing the information you specify in the connection string. The *asasrv.ini* file should be located in your Adaptive Server Anywhere executable directory (the same directory as the ODBC/DBLib DLL).

If you continue to experience problems establishing a connection, you should also delete any copy of *asasrv.ini* located in any of the following places:

- ◆ the win32 subdirectory of your SQL Anywhere installation directory (you can find this in the directory listed in the HKEY_LOCAL_MACHINE\SOFTWARE\Sybase\Adaptive Server Anywhere\8.0\Location registry key)
- ◆ Windows directory
- ◆ Windows system directory
- ◆ anywhere else in your path

CHAPTER 2

Connecting to a Database

About this chapter

This chapter describes how client applications connect to databases. It contains information about connecting to databases from ODBC, OLE DB, and embedded SQL applications. It also describes connecting from Sybase Central and Interactive SQL.

For more information on connecting to a database from Sybase Open Client applications, see "Adaptive Server Anywhere as an Open Server" on page 105.

For more information on connecting via JDBC (if you are not working in Sybase Central or Interactive SQL), see "Data Access Using JDBC" on page 129 of the book *ASA Programming Guide*.

Contents

Topic	Page
Introduction to connections	38
Connecting from Sybase Central or Interactive SQL	42
Simple connection examples	45
Working with ODBC data sources	53
Connecting from desktop applications to a Windows CE database	65
Connecting to a database using OLE DB	68
Connection parameters	70
Troubleshooting connections	73
Using integrated logins	83

Introduction to connections

How connections are established

Any client application that uses a database must establish a **connection** to that database before any work can be done. The connection forms a channel through which all activity from the client application takes place. For example, your user ID determines permissions to carry out actions on the database—and the database server has your user ID because it is part of the request to establish a connection.

To establish a connection, the client application calls functions in one of the Adaptive Server Anywhere interfaces. Adaptive Server Anywhere provides the following interfaces:


- ◆ **ODBC** ODBC connections are discussed in this chapter.
- ◆ **OLE DB** OLE DB connections are discussed in this chapter.
- ◆ **Embedded SQL** Embedded SQL connections are discussed in this chapter.
- ◆ **Sybase Open Client** Open Client connections are not discussed in this chapter.

🔗 For information on connecting from Open Client applications, see "Adaptive Server Anywhere as an Open Server" on page 105.

- ◆ **JDBC** Sybase Central and Interactive SQL have the connection logic described in this chapter built into them. Other JDBC applications cannot use the connection logic discussed in this chapter.

🔗 For more information on connecting via JDBC, see "Data Access Using JDBC" on page 129 of the book *ASA Programming Guide*.

The interface uses connection information included in the call from the client application, perhaps together with information held on disk in a file data source, to locate and connect to a server running the required database. The following figure is a simplified representation of the pieces involved.



What to read

The following table identifies where you can find answers to questions.

If you want...	Consider reading...
An overview of connecting from Sybase Central or Interactive SQL (including a description of the drivers involved)	"Connecting from Sybase Central or Interactive SQL" on page 42
Some examples to get started quickly, including Sybase Central and Interactive SQL scenarios	"Simple connection examples" on page 45
To learn about data sources	"Working with ODBC data sources" on page 53
To learn what connection parameters are available	"Connection parameters" on page 70
To see an in-depth description of how connections are established	"Troubleshooting connections" on page 73
To learn about network-specific connection issues	"Client/Server Communications" on page 91
To learn about character set issues affecting connections	"Connection strings and character sets" on page 279

How connection parameters work

When an application connects to a database, it uses a set of **connection parameters** to define the connection. Connection parameters include information such as the server name, the database name, and a user ID.

A keyword-value pair (of the form *parameter=value*) specifies each connection parameter. For example, you specify the password connection parameter for the default password as follows:

```
Password=SQL
```

Connection parameters are assembled into **connection strings**. In a connection string, a semicolon separates each connection parameter, as follows:

```
ServerName=asademo;DatabaseName=asademo
```

Representing connection strings

This chapter has many examples of connection strings represented in the following form:

```
parameter1=value1  
parameter2=value2  
...
```

This is equivalent to the following connection string:

```
parameter1=value1;parameter2=value2
```

You must type a connection string on a single line with the parameter settings separated by semicolons.

Connection parameters passed as connection strings

Connection parameters are passed to the interface library as a **connection string**. This string consists of a set of parameters, separated by semicolons:

```
parameter1=value1;parameter2=value2;...
```

In general, the connection string built up by an application and passed to the interface library does not correspond directly to the way a user enters the information. Instead, a user may fill in a dialog, or the application may read connection information from an initialization file.

Many of the Adaptive Server Anywhere utilities accept a connection string as the `-c` option and pass the connection string on to the interface library without change. For example, the following is a typical Collation utility (*dbcollat*) command line (which should be typed all on one line):

```
dbcollat -c "uid=DBA;pwd=SQL;dbn=asademo"  
c:\temp\asademo.col
```

Interactive SQL connection strings

Interactive SQL processes the connection string internally. These utilities do not simply pass on the connection parameters to the interface library. Do not use Interactive SQL to test command strings from the command prompt.

Saving connection parameters in ODBC data sources

Many client applications, including application development systems, use the ODBC interface to access Adaptive Server Anywhere. When connecting to the database, ODBC applications typically use ODBC data sources. An ODBC data source is a set of connection parameters, stored in the registry or in a file.

☞ For more information on ODBC data sources, see "Working with ODBC data sources" on page 53.

For Adaptive Server Anywhere, ODBC data sources can be used not only by ODBC applications on Windows, but also by other applications:

- ◆ Adaptive Server Anywhere client applications on UNIX can use ODBC data sources, as well as those on Windows operating systems. On UNIX, the data source is stored as a file.
- ◆ Adaptive Server Anywhere client applications using the OLE DB or embedded SQL interfaces can use ODBC data sources, as well as ODBC applications.
- ◆ Interactive SQL and Sybase Central can use ODBC data sources.


Connecting from Sybase Central or Interactive SQL

To use Sybase Central or Interactive SQL for managing your database, you must first connect to it. In the Connect dialog, you tell Sybase Central or Interactive SQL what database you want to connect to, where it is located, and how you want to connect to it.

The connecting process depends on your situation. For example, if you have a server already running on your machine and this server contains only one database, all you have to do in the Connect dialog is provide a user ID and a password. Sybase Central or Interactive SQL then knows to connect immediately to the database on the running server.

If this running server has more than one database loaded on it, if it is not yet running, or if it is running on another machine, you need to provide more detailed information in the Connect dialog so that Sybase Central or Interactive SQL knows which database to connect to.

This section describes how to access the Connect dialog in Sybase Central and Interactive SQL.

 For more information about connection examples, including examples for Sybase Central and Interactive SQL, see "Simple connection examples" on page 45.

Opening the Connect dialog

A common Connect dialog is available in both Sybase Central and Interactive SQL to let you connect to a database.

When you start Sybase Central, you need to manually display this dialog. When you start Interactive SQL, the dialog automatically appears; you can also make it appear for a new connection by choosing Window►New Window.

❖ To open the Connect dialog (Sybase Central):

- ◆ In Sybase Central, choose Tools►Connect.

If you have more than one Sybase Central plug-in installed, choose Adaptive Server Anywhere from the displayed list.

or

Click the Connect button on the main toolbar.

or
Press F11.

Tip

You can make subsequent connections to a given database easier and faster using a connection profile.

❖ **To open the Connect dialog (Interactive SQL):**

- ◆ In Interactive SQL, choose SQL ► Connect.

or
Press F11.

Once the Connect dialog appears, you must specify the connection parameters you need to connect. For example, you can connect to the Adaptive Server Anywhere sample database by choosing *ASA 8.0 Sample* from the ODBC Data Source Name list and clicking OK.


Specifying a driver for your connection

When you are working with a database, all your requests and commands go through a driver to the database itself. Sybase Central and Interactive SQL support two main types of drivers: a JDBC driver (called jConnect) and an ODBC driver (called JDBC-ODBC Bridge). Both are included with Adaptive Server Anywhere.

Sybase jConnect is a fully supported, fully featured JDBC driver. This driver is platform-independent. It is enabled by default.

The JDBC-ODBC bridge driver is available as an alternative method of connecting.

When you connect to a database in the Connect dialog, you can choose which driver you want to use for the connection. This is an optional configuration; the jConnect driver is the preferred driver and is automatically used for all connections unless you specify otherwise.

 For more information on ODBC and JDBC drivers, see "Choosing a JDBC driver" on page 131 of the book *ASA Programming Guide*, "Using the jConnect JDBC driver" on page 136 of the book *ASA Programming Guide*, and "Working with ODBC data sources" on page 53.

Data sources and the jConnect driver

As a general rule, the jConnect driver cannot use ODBC data sources. However, Sybase Central and Interactive SQL are special cases. When you use the jConnect driver in either of them, you can specify an ODBC data source to establish a connection. For example, you can connect to the sample database using the ASA 8.0 Sample data source, even if you are using the jConnect driver.

This customized functionality is only available while you are working in Sybase Central or Interactive SQL. If you are constructing a JDBC application, do not try to use a data source to connect to a database.

❖ To specify a driver for the connection:

- 1 In Sybase Central, choose Tools►Connect to open the Connect dialog.
- 2 Configure the necessary settings on the Identification and Database tabs of the dialog.
- 3 On the Advanced tab of the dialog, select either jConnect5 or JDBC-ODBC Bridge.

Working with the Connect dialog

The Connect dialog lets you define parameters for connecting to a server or database. The same dialog is used in both Sybase Central and Interactive SQL.

The Connect dialog has the following tabs:

- ◆ The Identification tab lets you identify yourself to the database and specify a data source.
- ◆ The Database tab lets you identify a server and/or database to connect to.
- ◆ The Advanced tab lets you add additional connection parameters and specify a driver for the connection.

In Sybase Central, after you connect successfully, the database name appears in the left pane of the main window, under the server that it is running on. The user ID for the connection appears after the database name.


In Interactive SQL, the connection information (including the database name, your user ID, and the database server) appears in the title bar above the SQL Statements pane.

Simple connection examples

Although the connection model for Adaptive Server Anywhere is configurable, and can become complex, in many cases connecting to a database is very simple.

Who should read this section?

This section describes some simple cases of applications connecting to an Adaptive Server Anywhere database. This section may be all you need to get started.

 For more information about available connection parameters and their use, see "Connection parameters" on page 70.

Connecting to the sample database from Sybase Central or Interactive SQL

Many examples and exercises throughout the documentation start by connecting to the sample database from Sybase Central or Interactive SQL.

❖ To connect to the sample database (Sybase Central):

- 1 To start Sybase Central: from the Start menu, choose Programs►Sybase SQL Anywhere 8►Sybase Central.
- 2 To open the Connect dialog: from the Tools menu, choose Connect.
- 3 Select the ODBC Data Source Name option and click Browse.
- 4 Select *ASA 8.0 Sample* and click OK.

❖ To connect to the sample database (Interactive SQL):

- 1 To start Interactive SQL: from the Start menu, choose Programs►Sybase SQL Anywhere 8►Adaptive Server Anywhere►Interactive SQL.
- 2 To open the Connect dialog: from the SQL menu, choose Connect.
- 3 Select the ODBC Data Source Name option and click Browse.
- 4 Select *ASA 8.0 Sample* and click OK.

The database name, user ID, and server name appear in the title bar above the SQL Statements pane.

Note

You do not need to type a user ID and a password for this connection because the data source already contains this information.

Connecting to a database on your own machine from Sybase Central or Interactive SQL

The simplest connection scenario is when the database you want to connect to resides on your own machine. If this is the case for you, ask yourself the following questions:

- ◆ Is the database already running on a server? If so, you can specify fewer parameters in the Connect dialog. If not, you need to identify the database file so that Sybase Central or Interactive SQL can start it for you.
- ◆ Are there multiple databases running on your machine? If so, you need to tell Sybase Central or Interactive SQL which database in particular to connect to. If there is only one, Sybase Central or Interactive SQL assumes that it is the one you want to connect to, and you don't need to specify it in the Connect dialog.

The procedures below depend on your answers to these questions.

❖ To connect to a database on an already-running local server:

- 1 Start Sybase Central or Interactive SQL and open the Connect dialog (if it doesn't appear automatically).
- 2 On the Identification tab of the dialog, type a user ID and a password.
- 3 Do one of the following:
 - ◆ If the server only contains the one database, click OK to connect to it.
 - ◆ If the server contains multiple databases, click the Database tab of the dialog and specify a database name. This is usually the database file name, without the path or extension.


❖ To start and connect to a database:

- 1 Start Sybase Central or Interactive SQL and open the Connect dialog (if it doesn't appear automatically).
- 2 On the Identification tab of the dialog, type a user ID and a password.
- 3 Click the Database tab of the dialog.
- 4 Specify a file in the Database File field (including the full path, name, and extension). You can search for a file by clicking Browse.
- 5 If you want the database name for subsequent connections to be different from the file name, type a name in the Database Name field (without including a path or extension).

Tips

If the database is already loaded (started) on the server, you only need to provide a database name for a successful connection. The database file is not necessary.

You can connect using a data source (a stored set of connection parameters) for either of the above scenarios by selecting the appropriate data source option at the bottom of the Identification tab of the Connect dialog.

 For information about using data sources in conjunction with the JDBC driver (jConnect), see "Specifying a driver for your connection" on page 43.

 See also

- ◆ "Opening the Connect dialog" on page 42
- ◆ "Simple connection examples" on page 45

Connecting to an embedded database

An **embedded database**, designed for use by a single application, runs on the same machine as the application and is largely hidden from the application user.

When an application uses an embedded database, the personal server is generally not running when the application connects. In this case, you can start the database using the connection string, and by specifying the database file in the **DatabaseFile (DBF)** parameter of the connection string.

Using the DBF parameter

The **DatabaseFile (DBF)** parameter specifies which database file to use. The database file automatically loads onto the default server, or starts a server if none are running.

The database unloads when there are no more connections to the database (generally when the application that started the connection disconnects). If the connection started the server, it stops once the database unloads.

The following connection parameters show how to load the sample database as an embedded database:

```
dbf=path\asademo.db
uid=DBA
pwd=SQL
```

where *path* is the name of your Adaptive Server Anywhere installation directory.

Using the StartLine [Start] parameter

The following connection parameters show how you can customize the startup of the sample database as an embedded database. This is useful if you wish to use options, such as the cache size:

```
Start=dbeng8 -c 8M
dbf=path\asademo.db
uid=DBA
pwd=SQL
```


 See also

- ◆ "Opening the Connect dialog" on page 42
- ◆ "Simple connection examples" on page 45

Connecting using a data source

You can save sets of connection parameters in a **data source**. ODBC and Embedded SQL applications use data sources. You can create data sources from the ODBC Administrator.

If you are constructing an application, you should only use data sources for ODBC applications. It is possible to specify data sources when you are using the JDBC driver (jConnect), but only within Sybase Central or Interactive SQL.

 For more information, see "Specifying a driver for your connection" on page 43.

❖ To connect from using a data source (Sybase Central or Interactive SQL):

- 1 Start Sybase Central or Interactive SQL and open the Connect dialog (if it doesn't appear automatically).
- 2 On the Identification tab, type a user ID and password.
- 3 On the lower half of the Identification tab, do one of the following:
 - ◆ Select the ODBC Data Source Name option and specify a data source name (equivalent to the **DataSourceName (DSN)** connection parameter, which references a data source in the registry). You can view a list of data sources by clicking Browse.
 - ◆ Select the ODBC Data Source File option and specify a data source file (equivalent to the **FileDataSourceName (FILEDSN)** connection parameter, which references a data source held in a file). You can search for a file by clicking Browse.

The ASA 8.0 Sample data source holds a set of connection parameters, including the database file and a **StartLine (START)** parameter to start the database.

See also


- ◆ "Opening the Connect dialog" on page 42
- ◆ "Simple connection examples" on page 45

Connecting to a server on a network

To connect to a database running on a network server somewhere on a local or wide area network, the client software must locate the database server. Adaptive Server Anywhere provides a network library to handle this task.

Network connections occur over a **network protocol**. Several protocols are supported, including TCP/IP and SPX.

For more information about client/server communications over a network, see "Client/Server Communications" on page 91.



Specifying the server

Adaptive Server Anywhere server names must be unique on a local domain for a given network protocol. The following connection parameters provide a simple example for connecting to a server running elsewhere on a network:

```
eng=svr_name
dbn=db_name
uid=user_id
pwd=password
CommLinks=all
```

The client library first looks for a personal server of the given name, and then looks on the network for a server of the given name.

Specifying the protocol

☞ The above example finds any server started using the default port number. However, you can start servers using other port numbers by providing more information in the **CommLinks (LINKS)** parameter. For information, see "CommLinks connection parameter" on page 169.

If several protocols are available, you can instruct the network library which ones to use to improve performance. The following parameters use only the TCP/IP protocol:

```
eng=svr_name
dbn=db_name
uid=user_id
pwd=password
CommLinks=tcPIP
```

The network library searches for a server by broadcasting over the network, which can be a time-consuming process. Once the network library locates a server, the client library stores its name and network address in a file (*asasrv.ini*), and reuses this entry for subsequent connection attempts to that server using the specified protocol. Subsequent connections can be many times faster than a connection achieved by broadcast.


☞ Many other connection parameters are available to assist Adaptive Server Anywhere in locating a server efficiently over a network. For more information see "Network communications parameters" on page 189.

❖ To connect to a database on a network server (Sybase Central or Interactive SQL):


- 1 Start Sybase Central or Interactive SQL and open the Connect dialog (if it does not appear automatically).
- 2 On the Identification tab of the dialog, type a user ID and a password.
- 3 On the Database tab of the dialog, type the Server Name. You can search for a server by clicking Find.
- 4 Identify the database by specifying a Database Name.

Tips

You can connect using a data source (a stored set of connection parameters) by selecting the appropriate data source option at the bottom of the Identification tab of the Connect dialog.

 For information about using data sources in conjunction with the JDBC driver (jConnect), see "Specifying a driver for your connection" on page 43.

By default, all network connections in Sybase Central and Interactive SQL use the TCP/IP network protocol.

 See also

- ◆ "Opening the Connect dialog" on page 42
- ◆ "Simple connection examples" on page 45

Using default connection parameters

You can leave many connection parameters unspecified, and instead use the default behavior to make a connection. Be cautious about relying on default behavior in production environments, especially if you distribute your application to customers who may install other Adaptive Server Anywhere applications on their machine.

Default database
server and
database

If a single personal server is running, with a single loaded database, you can connect using entirely default parameters:

```
uid=user_id
pwd=password
```

Default database
server

If more than one database is loaded on a single personal server, you can leave the server as a default, but you need to specify the database you wish to connect to:

```
dbn=db_name
uid=user_id
pwd=password
```

Default database

If more than one server is running, you need to specify which server you wish to connect to. If only one database is loaded on that server, you do not need to specify the database name. The following connection string connects to a named server, using the default database:

```
eng=server_name
uid=user_id
pwd=password
```

No defaults

The following connection string connects to a named server, using a named database:

```
eng=server_name  
dbn=db_name  
uid=user_id  
pwd=password
```

☞ For more information about default behavior, see "Troubleshooting connections" on page 73.

Connecting from Adaptive Server Anywhere utilities

All Adaptive Server Anywhere database utilities that communicate with the server (rather than acting directly on database files) do so using Embedded SQL. They follow the procedure outlined in "Troubleshooting connections" on page 73 when connecting to a database.

How database tools obtain connection parameter values

Many of the administration utilities obtain the connection parameter values by:

- 1 Using values specified on the command line (if there are any). For example, the following command starts a backup of the default database on the default server using the user ID DBA and the password SQL:

```
dbbackup -c "uid=DBA;pwd=SQL" c:\backup
```

☞ For more information about options for each database tool, see the chapter "Database Administration Utilities" on page 435.

- 2 Using the SQLCONNECT environment variable settings if any values are missing. Adaptive Server Anywhere does not set this variable automatically.

☞ For more information about the SQLCONNECT environment variable, see "Environment variables" on page 208.

- 3 Prompting you for a user ID and password to connect to the default database on the default server, if parameters are not set in the command line, or the SQLCONNECT environment variable.

Working with ODBC data sources

Microsoft Corporation defines the **Open Database Connectivity (ODBC)** interface, which is a standard interface for connecting client applications to database-management systems in the Windows 95/98/Me and Windows NT/2000/XP environments. Many client applications, including application development systems, use the ODBC interface to access a wide range of database systems.

Where data
sources are held

You connect to an ODBC database using an ODBC data source. You need an ODBC data source on the client computer for each database you want to connect to.

The ODBC data source contains a set of connection parameters. You can store sets of Adaptive Server Anywhere connection parameters as an ODBC data source, in either the system registry or as files.

If you have a data source, your connection string can simply name the data source to use:

- ◆ **Data source** Use the **DataSourceName (DSN)** connection parameter to reference a data source in the registry:

`DSN=my data source`

- ◆ **File data source** Use the **FileDataSourceName (FILEDSN)** connection parameter to reference a data source held in a file:

`FileDSN=mysource.dsn`

For Adaptive Server Anywhere, the use of ODBC data sources goes beyond Windows applications using the ODBC interface:

- ◆ Adaptive Server Anywhere client applications on UNIX can use ODBC data sources, as well as those on Windows operating systems.
- ◆ Adaptive Server Anywhere client applications using the OLE DB or embedded SQL interfaces can use ODBC data sources, as well as ODBC applications.
- ◆ Interactive SQL and Sybase Central can use ODBC data sources.

Creating an ODBC data source

You can create ODBC data sources on Windows 95/98/Me and Windows NT/2000/XP operating systems using the ODBC Administrator, which provides a central place for creating and managing ODBC data sources.

Adaptive Server Anywhere also includes a cross-platform utility named *dbdsn* to create data sources.

Before you begin

This section describes how to create an ODBC data source. Before you create a data source, you need to know which connection parameters you want to include in it.

☞ For more information, see "Simple connection examples" on page 45, and "Connection parameters" on page 70.

ODBC Administrator

On Windows 95/98/Me and Windows NT/2000/XP, you can use the Microsoft ODBC Administrator to create and edit data sources. You can work with User Data Sources, File Data Sources, and System Data Sources in this utility.

❖ **To create an ODBC data source (ODBC Administrator):**

- 1 Start the ODBC Administrator:

In Sybase Central, choose Tools►Adaptive Server Anywhere 8►Open ODBC Administrator.

or

From the Windows Start menu, choose Programs►Sybase SQL Anywhere 8►Adaptive Server Anywhere►ODBC Administrator.

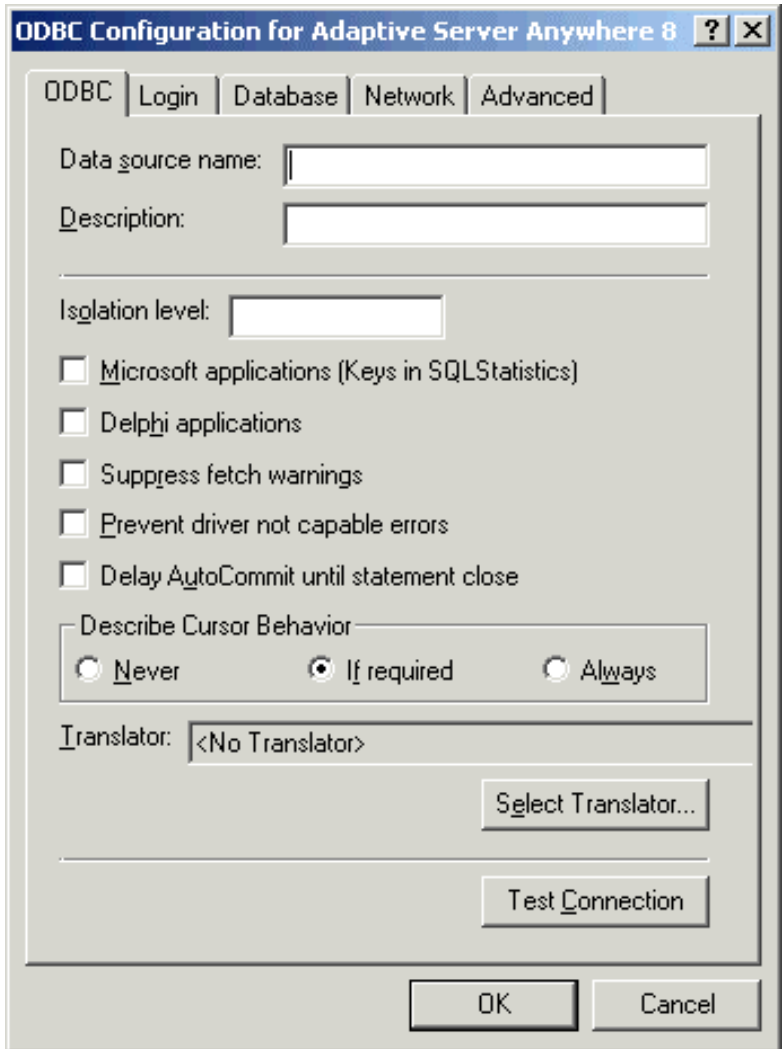
The ODBC Data Source Administrator dialog appears.

- 2 Click Add.


The Create New Data Source wizard appears.

- 3 From the list of drivers, choose Adaptive Server Anywhere 8.0, and click Finish.

The ODBC Configuration for Adaptive Server Anywhere dialog appears.



Most of the fields in this window are optional. Click the question mark at the top right of the window and click a dialog field to find more information about that field.

 For more information about the fields in the dialog, see "Configuring ODBC data sources using the ODBC Administrator" on page 56.

- 4 When you have specified the parameters you need, click OK to close the window and create the data source.

Creating an ODBC data source from the command line

To edit a data source, find and select one in the ODBC Administrator main window and click Configure.

You can create User Data Sources using the *dbdsn* utility. You cannot create File Data Sources or System Data Sources. File and System Data Sources are limited to Windows operating systems only, and you can use the ODBC Administrator to create them.

❖ To create an ODBC data source (Command line):

- 1 Open a command prompt.
- 2 Type a *dbdsn* command, specifying the connection parameters you wish to use.

For example, the following command creates a data source for the Adaptive Server Anywhere sample database. The command must be typed on one line:

```
dbdsn -w "My DSN"  
"uid=DBA;pwd=SQL;dbf=c:\Program Files\Sybase\SQL Anywhere 8\asademo.db"
```

🔗 For more information on the *dbdsn* utility, see "The Data Source utility" on page 451.

Configuring ODBC data sources using the ODBC Administrator

This section describes the meaning of each of the options on the ODBC configuration dialog.

ODBC tab

This tab provides fields related to ODBC data sources and appears only when you are creating or modifying ODBC data sources.

Data source name The Data Source Name is used to identify the ODBC data source. You can use any descriptive name for the data source (spaces are allowed) but it is recommended that you keep the name short, as you may need to type it in connection strings.

🔗 For more information, see "DataSourceName connection parameter" on page 175.

Description You may type an optional longer description of the Data Source to help you or end users to identify this data source from among their list of available data sources. This field is optional.

Isolation level Type the desired isolation level for this data source:

- ◆ **0** Dirty reads, non-repeatable reads and phantom rows may occur. This is the default isolation level.
- ◆ **1** Non-repeatable rows and phantom rows may occur. Dirty reads are prevented.
- ◆ **2** Phantom rows may occur. Dirty reads and non-repeatable rows are prevented.
- ◆ **3** Dirty reads, non-repeatable reads and phantom rows are prevented.

☞ For more information, see "Choosing isolation levels" on page 102 of the book *ASA SQL User's Guide*.

Microsoft applications (Keys in SQLStatistics) Check this box if you want foreign keys to be returned by the **SQLStatistics** function. The ODBC specification states that SQLStatistics should not return primary and foreign keys. However, some Microsoft applications (such as Visual Basic and Access) assume that primary and foreign keys are returned by **SQLStatistics**.

Delphi applications Check this box to improve performance for Borland Delphi applications. When this option is checked, one bookmark value is assigned to each row, instead of the two that are otherwise assigned (one for fetching forwards and a different one for fetching backwards).

Delphi cannot handle multiple bookmark values for a row. If the option is unchecked, scrollable cursor performance can suffer since scrolling must always take place from the beginning of the cursor to the row requested in order to get the correct bookmark value.

Suppress fetch warnings Check this box to suppress warning messages that are returned from the database server on a fetch.

Versions 8.0 and later of the database server return a wider range of fetch warnings than earlier versions of the software. For applications that are deployed with an earlier version of the software, you can select this option to ensure that fetch warnings are handled properly.

Prevent driver not capable errors The Adaptive Server Anywhere ODBC driver returns a Driver not capable error code because it does not support qualifiers. Some ODBC applications do not handle this error properly. Check this box to disable this error code, allowing such applications to work.

Delay AutoCommit until statement close Check this box if you want the Adaptive Server Anywhere ODBC driver to delay the commit operation until a statement has been closed.

Describe cursor behavior Choose how often you wish a cursor to be redescrbed when a procedure is executed or resumed. The default setting is If Required.

- ◆ **Never** Select this option if you know that your cursors do not have to be redescrbed because redescrbing cursors is expensive and can decrease performance.
- ◆ **If required** When you select this option, the ODBC driver determines whether a cursor must be redescrbed. The presence of a **RESULT** clause in your procedure prevents ODBC applications from redescrbing the result set after a cursor is opened.
- ◆ **Always** The cursor is redescrbed each time it is opened. If you use Transact-SQL procedures or procedures that return multiple result sets, you must redescrbe the cursor each time it is opened.

Select Translator Opens the Select Translator dialog where you can select the ODBC translator that you want to use from a list of installed translators. Use this option sparingly, since the ODBC driver and database server do automatic character set translation.

Test Connection Tests whether the information provided results in a proper connection. In order for the test to work, a user ID and password must have been specified.

Login tab


This tab provides fields related to user authentication and logins.

Use integrated login Connects using an integrated login. The user ID and password do not need to be specified: instead, your operating system user id and password are supplied to the Adaptive Server Anywhere integrated login mechanism.

To use this type of login users must have been granted integrated login permission. The database being connected to must also be set up to accept integrated logins. Only users with DBA access may administer integrated login permissions.

 For more information, see "Using integrated logins" on page 83.

User ID Provides a place for you to type the user ID for the connection.

 For more information, see "Userid connection parameter" on page 188.

Password Provides a place for you to type the password for the connection.

For more information, see "Password connection parameter" on page 184.

Encrypt password Check this box if you want the password to be stored in encrypted form in the profile.

For more information, see "EncryptedPassword connection parameter" on page 177.

Database tab

This tab provides fields for identifying the database server and database to which you are connecting.

Server name Provides a place for you to type the name of the Adaptive Server Anywhere personal or network server to which you want to connect.

For more information, see "EngineName connection parameter" on page 176.

Start line Type the name of the server executable that should be started. Only provide a **StartLine (START)** connection parameter if the database server you are connecting to is not currently running. For example:

```
C:\Program Files\Sybase\SQL Anywhere 8\win32\dbeng8.exe -c 8m
```

For more information, see "StartLine connection parameter" on page 187.

Database name Provides a place for you to type the name of the Adaptive Server Anywhere database that you wish to connect to.

For more information, see "DatabaseName connection parameter" on page 174.

Database file Provides a place for you to type the full path and name of the Adaptive Server Anywhere database file on the server PC. You may also click Browse to locate the file. For example:

```
C:\Program Files\Sybase\SQL Anywhere 8\asademo.db
```

For more information, see "DatabaseFile connection parameter" on page 173.

Encryption key If your database is strongly encrypted, you must supply the correct encryption key to access any part of the database.

Automatically start the database if it isn't running Causes the database to start automatically (if it is not already running) when you start a new session. This option is enabled when you supply a database file or start line.

Automatically shut down database after last disconnect Causes the automatic shutdown of the server after the last user has disconnected.

🔗 For more information, see "AutoStop connection parameter" on page 167.

Network tab

This tab provides fields for controlling client/server communications over a network.

Select the network protocols and options These checkboxes specify what protocol or protocols the ODBC DSN uses to access a network database server. In the adjacent boxes, you may enter communication parameters that establish and tune connections from your client application to a database.

- ◆ **TCP/IP** If you want to use ECC_TLS (Certicom) or RSA_TLS encryption for network packets, you must select the TCP/IP protocol to access the network database server.

For example, you could enter the following communication parameters for a TCP/IP connection **HOST=my_server;PORT=4563**.

- ◆ **SPX** You can select the SPX protocol to connect to databases on Novell NetWare networks. NetWare also supports the TCP/IP protocol.
- ◆ **Named Pipes** The Named Pipes protocol is used for client/server communication on the same machine. If you want to run under a certified security environment, you can use the Named Pipes protocol. It is only provided on Windows NT/2000/XP.
- ◆ **Shared Memory** The shared memory protocol is used for communication between a client and server running under the same operating system on the same machine.


🔗 For more information, see "CommLinks connection parameter" on page 169, and "Network communications parameters" on page 189.

Liveness timeout A liveness packet is sent across a client/server to confirm that a connection is intact. If the client runs for the liveness timeout period without detecting a liveness packet, the communication will be severed. This parameter works only with the network server and TCP/IP and SPX communications protocols. The default is 120 seconds.


🔗 For more information, see "LivenessTimeout connection parameter" on page 182.

Idle timeout Set the amount of client idle time before the connection is terminated. If a client runs for the idle timeout period without submitting a request, the connection is severed.

The default client idle time is 240 minutes.

 For more information, see "Idle connection parameter" on page 181.

Buffer size Sets the maximum size of communication packets, in bytes. The default buffer size is 1460.

 For more information, see "CommBufferSize connection parameter" on page 168.

Compress network packets Sets compression for the connection. Using compression for a connection can significantly improve Adaptive Server Anywhere performance under some circumstances.

 For more information, see "Compress connection parameter" on page 170.


Select the method for encryption of network packets Allows the encryption of packets transmitted from the client machine over the network.

- ◆ **None** Communication packets transmitted from the client are not encrypted. This is the default setting.
- ◆ **Simple** Communication packets transmitted from the client are encrypted with simple encryption. Simple encryption is supported on all platforms, as well as on previous versions of Adaptive Server Anywhere.
- ◆ **ECC_TLS** Communication packets transmitted from the client are encrypted with ECC_TLS (formerly Certicom) encryption. This is a type of strong encryption. It is only available over the TCP/IP protocol.

In the adjacent field you must supply a trusted certificate value.

- ◆ **RSA_TLS** Communication packets transmitted from the client are encrypted with RSA_TLS encryption. This is a type of strong encryption. It is only available over the TCP/IP protocol.

In the adjacent field you must supply a trusted certificate value.

 For more information, see "Encryption connection parameter" on page 177.

Advanced tab

This tab provides advanced fields that are for special purposes and are not required by the majority of users.

Connection name The name of the connection that is being created.

Character set Lets you specify a character set (a set of 256 letters, numbers, and symbols specific to a country or language). The ANSI character set is used by Microsoft Windows. An OEM character set is any character set except the ANSI character set.

Allow multiple record fetching Enables multiple records to be retrieved at one time instead of individually. By default, multiple record fetching is allowed.

Display debugging information in a log file Select this option to specify the name of the file in which the debugging information is to be saved.

Additional connection parameters Type any additional options here. Parameters set throughout the remainder of this dialog take precedence over parameters typed here.

 For more information, see "Connection parameters" on page 164.

Using file data sources on Windows

On Windows operating systems, ODBC data sources are typically stored in the system registry. File data sources are an alternative, which are stored as files. In Windows, file data sources typically have the extension *.dsn*. They consist of sections, each section starting with a name enclosed in square brackets. DSN files are very similar in layout to initialization files.

To connect using a File Data Source, use the **FileDataSourceName (FILEDSN)** connection parameter. You cannot use both **DataSourceName (DSN)** and **FileDataSourceName (FILEDSN)** in the same connection.

File data sources
can be distributed

One benefit of file data sources is that you can distribute the file to users. If the file is placed in the default location for file data sources, it is picked up automatically by ODBC. In this way, managing connections for many users can be made simpler.

Embedded SQL applications can also use ODBC file data sources.

❖ To create an ODBC file data source (ODBC Administrator):

- 1 Start the ODBC Administrator, click the File DSN tab and click Add.
- 2 Select Adaptive Server Anywhere 8.0 from the list of drivers, and click Next.
- 3 Follow the instructions to create the data source.

Using ODBC data sources on UNIX

On UNIX operating systems, ODBC data sources are held in a file named *.odbc.ini*. A sample file looks like this:

```
[My Data Source]
ENG=myserver
CommLinks=tcipip(Host=hostname)
uid=DBA
pwd=SQL
```

You can type any connection parameter in the *.odbc.ini* file.

☞ For a complete list of connection parameters, see "Connection parameters" on page 164.

Network communications parameters are added as part of the **CommLinks (LINKS)** parameter.

☞ For a complete list of network communications parameters, see "Network communications parameters" on page 189.

You can create and manage ODBC data sources on UNIX using the *dbdsn* utility.

☞ For more information, see "Creating an ODBC data source" on page 53, and "The Data Source utility" on page 451.

File location

The database server looks for the *.odbc.ini* file in the following locations:

- 1 ODBCINI environment variable
- 2 ODBCHOME and HOME environment variables
- 3 The user's home directory
- 4 The path

Using ODBC data sources on Windows CE

Windows CE does not provide an ODBC driver manager or an ODBC Administrator. On this platform, Adaptive Server Anywhere uses ODBC data sources stored in files. You can specify either the **DSN** or the **FILEDSN** keyword to use these data source definitions—on Windows CE (only), **DSN** and **FILEDSN** are synonyms.

Data source location


Windows CE searches for the data source files in the following locations:

- 1 The directory from which the ODBC driver (*dbodbc8.dll*) was loaded. This is usually the Windows directory.

- 2 The directory specified in Location key of the Adaptive Server Anywhere section of the registry. This is usually the same as the Adaptive Server Anywhere installation directory. The default installation directory is:

```
\Program Files  
  \Sybase  
    \ASA
```

Each data source itself is held in a file. The file has the same name as the data source, with an extension of *.dsn*.

 For more information about file data sources, see "Using file data sources on Windows" on page 62.

Connecting from desktop applications to a Windows CE database

You can connect from applications running on a desktop PC, such as Sybase Central or Interactive SQL, to a database server running on a Windows CE device. The connection uses TCP/IP over the ActiveSync link between the desktop machine and the Windows CE device.

If you are using an Ethernet connection between your desktop machine and the Windows CE device, or if you are using Windows CE Services 2.2, the following procedure works. If you are using a serial cable connection and ActiveSync 3.0, see the following section.

❖ To connect from a desktop application to a Windows CE database server:

- 1 Determine the IP address of the server. Start the server on the Windows CE device with the `-z` option (output extra debug information).

For example:

```
dbsrv8 -z -x tcpip -n TestServer asademo.db
```

With the `-z` option, the server writes out its IP address during startup. The address may change if you disconnect your HPC from the network and then re-connect it.

To change between static and dynamic IP assignment for the HPC, configure the settings in the Windows Control Panel. Open Network, and choose the Services tab. Select Remote Access Service and click Properties►Network►TCP/IP Configuration.

- 2 Create an ODBC profile on your desktop machine.

Open the ODBC Administrator, and click Add. Choose Adaptive Server Anywhere 8.0 from the list of drivers and click Finish. The ODBC Configuration for Adaptive Server Anywhere dialog appears.

- ◆ On the Login tab, type a user ID and password.
- ◆ On the Database tab, type the server name.
- ◆ On the Network tab, check TCP/IP, and type the following in the adjacent field:

```
dobroadcast=DIRECT;host=XXX.XXX.XXX.XXX
```

where XXX.XXX.XXX.XXX is the server IP address.

- ◆ On the ODBC tab, click Test Connection to confirm that your ODBC data source is properly configured.
- 3 Exit the ODBC Administrator.
- 4 Ensure the database server is running on your Windows CE machine.
- 5 On your desktop machine, start an application such as Interactive SQL and select the ODBC data source you have created. The application connects to the Windows CE database.

Using ActiveSync 3.0 and a serial cable

To connect to a Windows CE device from your desktop over a serial cable, Adaptive Server Anywhere uses the TCP/IP protocol. For TCP/IP to work in this context, your ActiveSync installation must be set up to provide a Remote Access Service (RAS) link between desktop machine and Windows CE device.

ActiveSync 2.2 automatically installs and configures RAS, and you can connect in a straightforward manner.

For information, see "Connecting from desktop applications to a Windows CE database" on page 65.

ActiveSync 3.0 does not install and configure RAS. You must install RAS yourself to obtain TCP/IP connectivity to your device over a serial connection.

Instructions for installing RAS are provided by Microsoft. They are available at <http://support.microsoft.com/support/kb/articles/Q241/2/16.ASP> (Microsoft Knowledge Base article Q241216). You must follow these instructions *exactly*, including re-installing any Windows NT service packs, and granting your user account dial-in access using User Manager.

As you follow the instructions, where it says to install your modem, choose Dial-up Networking Serial Cable between 2 PCs instead.

Using RAS with ActiveSync

The following list includes suggestions for enabling ActiveSync 3.0 connections over a serial connection using RAS:

- 1 In the ActiveSync **Connection Settings**, select the checkbox Allow network (Ethernet) and Remote Access Service (RAS) server connection with desktop computer. You may need to turn off the checkbox Allow serial cable or infrared connection to this COM port.
- 2 On the desktop, using Remote Access Administrator (under Administrative Tools on Windows NT), start RAS on COM1.

- 3 On the Windows CE device, run the ActiveSync client (*repllog.exe* on a Windows CE PC). Choose serial connection.
- 4 Wait for up to one minute for a connection to be established.
- 5 As a test, run the *ipconfig* utility on Windows NT, and see the 192.168.55.100 static IP of the device. This is the IP you would use when connecting to an Adaptive Server Anywhere database server (for example) running on the CE device.
- 6 If you switch devices, Stop and Restart the RAS service (or reboot).
- 7 If everything is set up as above, but you still fail to get a connection from the device to the desktop, you should make sure your Port settings match the baud rates in the Modems Control Panel applet.

Connecting to a database using OLE DB

OLE DB uses the Component Object Model (COM) to make data from a variety of sources available to applications. Relational databases are among the classes of data sources that you can access through OLE DB.

This section describes how to connect to an Adaptive Server Anywhere database using OLE DB from the following environments:

- ◆ Sybase PowerBuilder can access OLE DB data sources, and you can use Adaptive Server Anywhere as a PowerBuilder OLE DB database profile.
- ◆ Microsoft ActiveX Data Objects (ADO) provides a programming interface for OLE DB data sources. You can access Adaptive Server Anywhere from programming tools such as Microsoft Visual Basic.

This section is an introduction to how to use OLE DB from Sybase PowerBuilder and Microsoft ADO environments such as Visual Basic. It is not complete documentation on how to program using ADO or OLE DB. The primary source of information on development topics is your development tool documentation.

☞ For more information about OLE DB, see "Introduction to OLE DB" on page 338 of the book *ASA Programming Guide*.

OLE DB providers

You need an OLE DB provider for each type of data source you wish to access. Each **OLE DB provider** is a dynamic-link library. There are two OLE DB providers you can use to access Adaptive Server Anywhere:

- ◆ **Sybase ASA OLE DB provider** The Adaptive Server Anywhere OLE DB provider provides access to Adaptive Server Anywhere as an OLE DB data source without the need for ODBC components. The short name for this provider is **ASAProv**.

When the **ASAProv** provider is installed, it registers itself. This registration process includes making registry entries in the COM section of the registry so that ADO can locate the DLL when the **ASAProv** provider is called. If you change the location of your DLL, you must reregister it.

☞ For more information about OLE DB providers, see "Introduction to OLE DB" on page 338 of the book *ASA Programming Guide*.

- ◆ **Microsoft OLE DB provider for ODBC** Microsoft provides an OLE DB provider with a short name of **MSDASQL**.

The **MSDASQL** provider makes ODBC data sources appear as OLE DB data sources. It requires the Adaptive Server Anywhere ODBC driver.

Connecting from ADO

ADO is an object-oriented programming interface. In ADO, the **Connection** object represents a unique session with a data source.

You can use the following **Connection** object features to initiate a connection:

- ◆ The **Provider** property that holds the name of the provider. If you do not supply a Provider name, ADO uses the MSDASQL provider.
- ◆ The **ConnectionString** property that holds a connection string. This property holds an Adaptive Server Anywhere connection string, which is used in just the same way as the ODBC driver. You can supply ODBC data source names, or explicit UserID, Password, DatabaseName, and other parameters, just as in other connection strings.

☞ For more information about connection parameters, see "Connection parameters" on page 70.

- ◆ The **Open** method initiates a connection.

☞ For more information about ADO, see "ADO programming with Adaptive Server Anywhere" on page 340 of the book *ASA Programming Guide*.


Example

The following Visual Basic code initiates an OLE DB connection to Adaptive Server Anywhere:

```
' Declare the connection object
Dim myConn as New ADODB.Connection
myConn.Provider = "ASAProv"
myConn.ConnectionString = "Data Source=ASA 8.0 Sample"
myConn.Open
```

Connection parameters

The following table lists the Adaptive Server Anywhere connection parameters.

 For more information about each of these connection parameters, see "Connection and Communication Parameters" on page 163. For character set issues in connection strings, see "Connection strings and character sets" on page 279.

Parameter	Short form	Argument
AppInfo	APP	String
AutoStart	ASTART	Boolean
AutoStop	ASTOP	Boolean
CharSet	CS	String
CommBufferSize	CBSIZE	Integer
CommLinks	LINKS	String
Compress	COMP	String
CompressionThreshold	COMPTH	Integer
ConnectionName	CON	String
DatabaseFile	DBF	String
DatabaseName	DBN	String
DatabaseSwitches	DBS	String
DataSourceName	DSN	String
DisableMultiRowFetch	DMRF	Boolean
EncryptedPassword	ENP	Encrypted string
Encryption	ENC	String
EngineName / ServerName	ENG	String
FileDataSourceName	FILEDSN	String
ForceStart	FORCE	Boolean
Idle connection	IDLE	Integer
Integrated	INT	Boolean
LazyClose	LCLOSE	Boolean
LivenessTimeout	LTO	Integer
Logfile	LOG	String

Parameter	Short form	Argument
Password *	PWD	String
PrefetchBuffer	PBUF	Integer
PrefetchRows	PROWS	Integer
StartLine	START	String
Unconditional	UNC	Boolean
Userid *	UID	String

* Verbose form of keyword not used by ODBC connection parameters

Notes

- ◆ **Boolean values** Boolean (true or false) arguments are either YES, ON, 1, or TRUE if true, or NO, OFF, 0, or FALSE if false.
- ◆ **Case sensitivity** Connection parameters are case insensitive.
- ◆ The connection parameters used by the interface library can be obtained from the following places (in order of precedence):
 - ◆ **Connection string** You can pass parameters explicitly in the connection string.
 - ◆ **SQLCONNECT environment variable** The SQLCONNECT environment variable can store connection parameters.
 - ◆ **Data sources** ODBC data sources can store parameters.
- ◆ **Character set restrictions** The server name must be composed of the ASCII character set in the range 1 to 127. There is no such limitation on other parameters.

☞ For more information on the character set issues, see "Connection strings and character sets" on page 279.

- ◆ **Priority** The following rules govern the priority of parameters:
 - ◆ The entries in a connect string are read left to right. If the same parameter is specified more than once, the last one in the string applies.
 - ◆ If a string contains a data source or file data source entry, the profile is read from the configuration file, and the entries from the file are used if they are not already set. For example, if a connection string contains a data source name and sets some of the parameters contained in the data source explicitly, then in case of conflict the explicit parameters are used.

Connection parameter tips

Connection parameters often provide more than one way of accomplishing a given task. This is particularly the case with embedded databases, where the connection string starts a database server. For example, if your connection starts a database, you can specify the database name using the **DatabaseName (DBN)** connection parameter or using the (**DatabaseSwitch DBS**) parameter.

Here are some recommendations and notes for situations where connection parameters conflict:

- ◆ **Specify database files using DBF** You can specify a database file on the **StartLine (START)** parameter or using the **DatabaseFile (DBF)** connection parameter (recommended).
- ◆ **Specify database names using DBN** You can specify a database name on the **StartLine (START)** parameter, the **DatabaseSwitch (DBS)** connection parameter, or using the **DatabaseName (DBN)** connection parameter (recommended).
- ◆ **Use the Start parameter to specify cache size** Even though you use the **DatabaseFile (DBF)** connection parameter to specify a database file, you may still want to tune the way in which it starts. You can use the **StartLine (START)** connection parameter to do this.

For example, if you are using the Java features of Adaptive Server Anywhere, you should provide additional cache memory on the **StartLine (START)** connection parameter. The following sample set of embedded database connection parameters describes a connection that may use Java features:


```
DBF=path\asademo.db
dbn=Sample
ENG=Sample Server
uid=DBA
pwd=SQL
Start=dbeng8 -c 8M
```

Troubleshooting connections

Who needs to read this section?

In many cases, establishing a connection to a database is straightforward using the information presented in the first part of this chapter.

However, if you are having problems establishing connections to a server, you may need to understand the process by which Adaptive Server Anywhere establishes connections in order to resolve your problems. This section describes how Adaptive Server Anywhere connections work.

 For more information about network-specific issues, including connections across firewalls, see "Client/Server Communications" on page 91.

The software follows exactly the same procedure for each of the following types of client application:

- ◆ **ODBC** Any ODBC application using the *SQLDriverConnect* function, which is the common method of connection for ODBC applications. Many application development systems, such as Sybase PowerBuilder and Power++, belong to this class of application.
- ◆ **Embedded SQL** Any client application using Embedded SQL and using the recommended function for connecting to a database (*db_string_connect*).

The SQL CONNECT statement is available for Embedded SQL applications and in Interactive SQL. It has two forms: CONNECT AS... and CONNECT USING. All the database administration tools, including Interactive SQL, use *db_string_connect*.

The steps in establishing a connection

To establish a connection, Adaptive Server Anywhere carries out the following steps in order:

- 1 **Locate the interface library** The client application must locate the ODBC driver or Embedded SQL interface library.
- 2 **Assemble a list of connection parameters** Since connection parameters may appear in several places (such as data sources, a connection string assembled by the application, or an environment variable) Adaptive Server Anywhere assembles the parameters into a single list.
- 3 **Locate a server** Using the connection parameters, Adaptive Server Anywhere locates a database server on your machine or over a network.

- 4 **Locate the database** Adaptive Server Anywhere locates the database you want to connect to.
- 5 **Start a personal server** If Adaptive Server Anywhere fails to locate a server, it attempts to start a personal database server and load the database.

The following sections describe each of these steps in detail.

Locating the interface library

The client application makes a call to one of the Adaptive Server Anywhere interface libraries. In general, the location of this DLL or shared library is transparent to the user. Here we describe how to locate the library in case of problems.

ODBC driver location

For ODBC, the interface library is also called an ODBC driver. An ODBC client application calls the ODBC driver manager, and the driver manager locates the Adaptive Server Anywhere driver.

The ODBC driver manager looks in the supplied data source in the *odbc.ini* file or registry to locate the driver. When you create a data source using the ODBC Administrator, Adaptive Server Anywhere fills in the current location for your ODBC driver.

Embedded SQL interface library location

Embedded SQL applications call the interface library by name. The name of the Adaptive Server Anywhere Embedded SQL interface library is as follows:

- ◆ **Windows NT/2000/XP and Windows 95/98/Me** *dblib8.dll*
- ◆ **UNIX** *dblib8* with an operating-system-specific extension
- ◆ **NetWare** *dblib8.nlm*

The locations that are searched depend on the operating system:

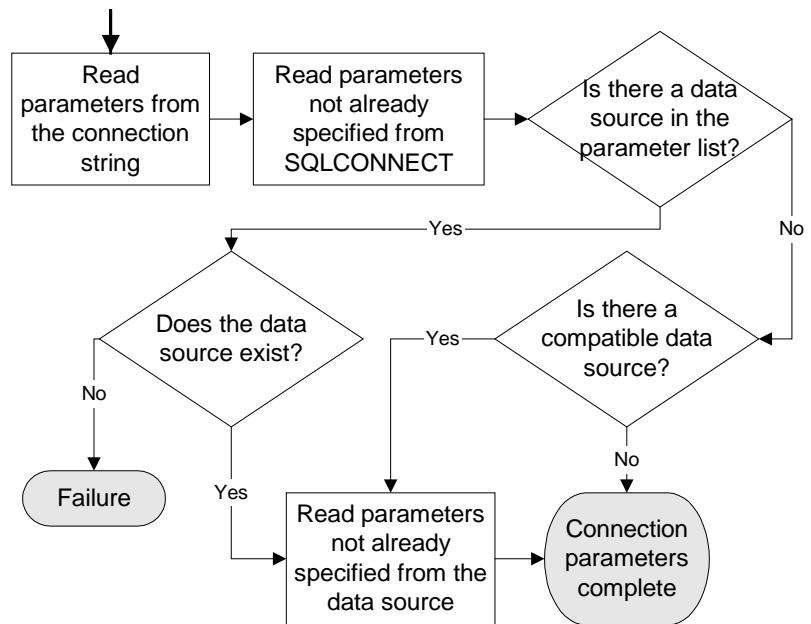
- ◆ **PC operating systems** On PC operating systems such as Windows, files are looked for in the current directory, in the system path, and in the *Windows* and *Windows\system* directories.
- ◆ **UNIX operating systems** On UNIX, files are looked for in the system path and the user library path.
- ◆ **NetWare** On NetWare, files are looked for in the search path, and in the *sys:system* directory.

When the library is located

Once the client application locates the interface library, it passes a connection string to it. The interface library uses the connection string to assemble a list of connection parameters, which it uses to establish a connection to a server.

Assembling a list of connection parameters

The following figure illustrates how the interface libraries assemble the list of connection parameters they use to establish a connection.



Notes

Key points from the figure include:

- ◆ **Precedence** Parameters held in more than one place are subject to the following order of precedence:

- 1 Connection string
- 2 SQLCONNECT
- 3 Data source

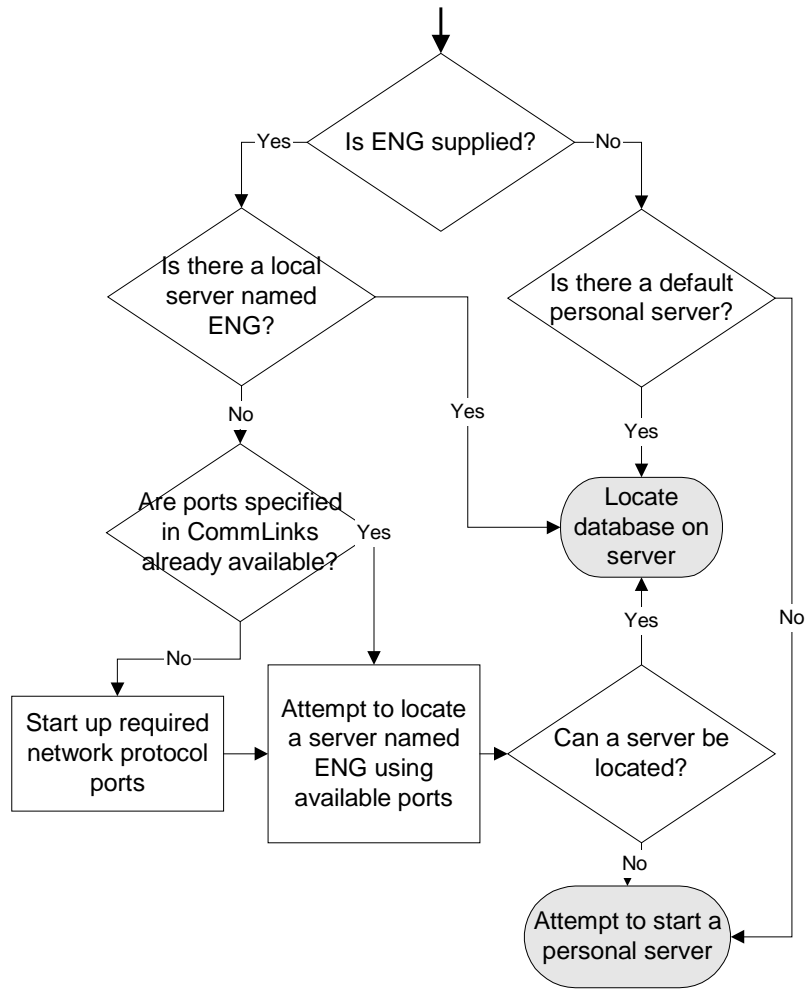
That is, if a parameter is supplied both in a data source and in a connection string, the connection string value overrides the data source value.

- ◆ **Failure** Failure at this stage occurs only if you specify in the connection string or in `SQLCONNECT` a data source that does not exist in the client connection file.
- ◆ **Common parameters** Depending on other connections already in use, some connection parameters may be ignored, including:
 - ◆ **AutoStop** Ignored if the database is already loaded.
 - ◆ **CommLinks** The specifications for a network protocol are ignored if another connection has already set parameters for that protocol.
 - ◆ **CommBufferSize** If `CommBufferSize` is specified on the client, the connection uses the minimum value of either the engine's buffer size or the `CommBufferSize` value. If it is not specified on the client, the connection uses the engine's buffer size.
 - ◆ **Unconditional** Ignored if the database is already loaded or if the server is already running.

The interface library uses the completed list of connection parameters to attempt to connect.

Locating a server

In the next step towards establishing a connection, Adaptive Server Anywhere attempts to locate a server. If the connection parameter list includes a server name (**EngineName (ENG)** connection parameter), it carries out a search first for a local server (personal server or network server running on the same machine) of that name, followed by a search over a network. If no **EngineName (ENG)** connection parameter is supplied, Adaptive Server Anywhere looks for a default server.



In the next step towards establishing a connection, Adaptive Server Anywhere attempts to locate a server. If the connection parameter list includes a server name (**EngineName (ENG)** connection parameter), it carries out a search first for a local server (personal server or network server running on the same machine) of that name, followed by a search over a network. If no **EngineName (ENG)** connection parameter is supplied, Adaptive Server Anywhere looks for a default server.

☞ If Adaptive Server Anywhere locates a server, it tries to locate or load the required database on that server. For information, see "Locating the database" on page 78.

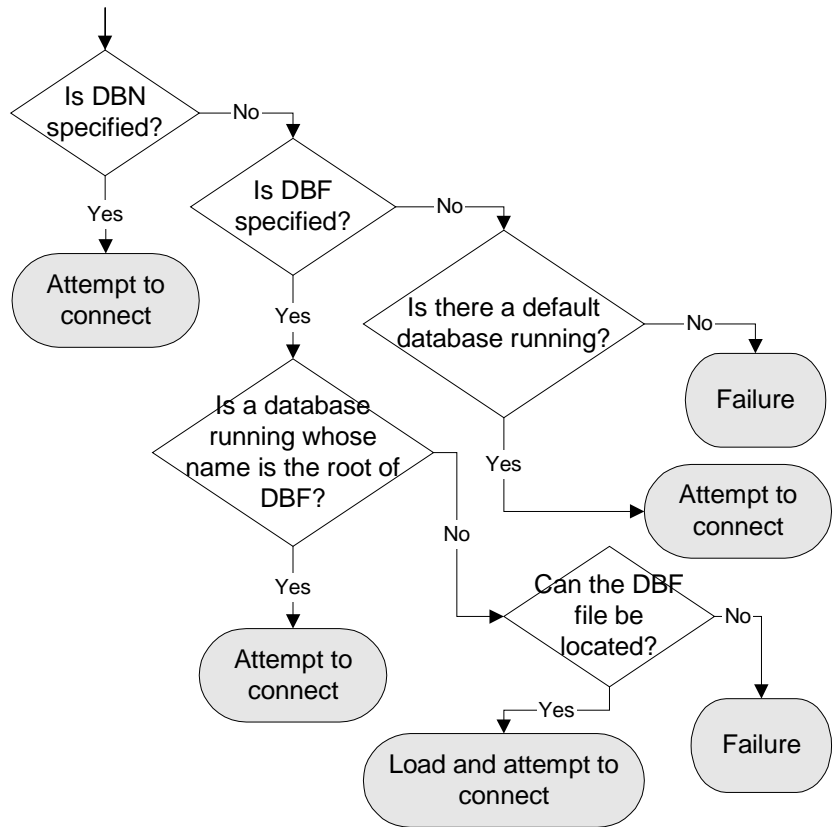
☞ If Adaptive Server Anywhere cannot locate a server, it attempts to start a personal server. For information, see "Starting a personal server" on page 79.

Notes

- ◆ For local connections, locating a server is simple. For connections over a network, you can use the **CommLinks (LINKS)** connection parameter to tune the search in many ways by supplying network communication parameters.
- ◆ The network search involves a search over one or more of the protocols supported by Adaptive Server Anywhere. For each protocol, the network library starts a single port. All connections over that protocol at any one time use a single port.
- ◆ You can specify a set of network communication parameters for each network port in the argument to the **CommLinks (LINKS)** connection parameter. Since these parameters are necessary only when the port first starts, the interface library ignores any connection parameters specified in the **CommLinks (LINKS)** connection parameter for a port already started.
- ◆ Each attempt to locate a server (the local attempt and the attempt for each network port) involves two steps. First, Adaptive Server Anywhere looks in the server name cache to see if a server of that name is available. Second, it uses the available connection parameters to attempt a connection.

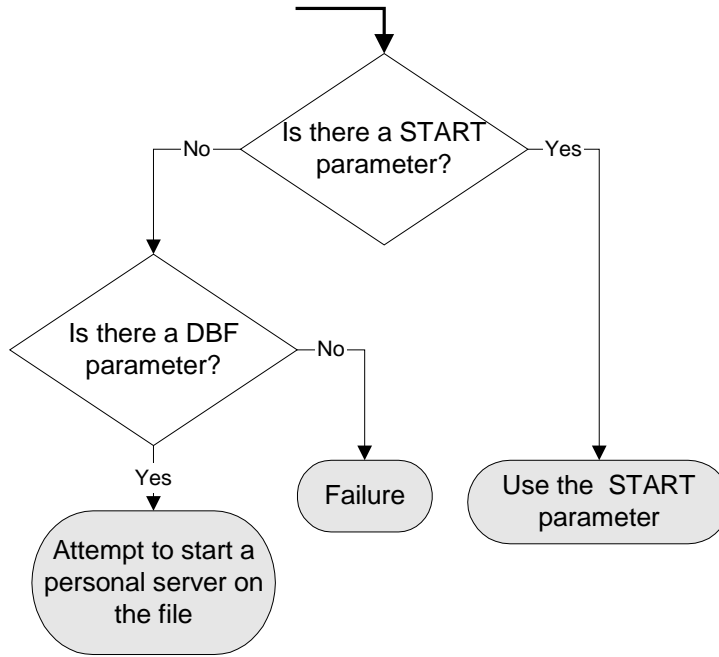
Locating the database

If Adaptive Server Anywhere successfully locates a server, it then tries to locate the database. For example:



Starting a personal server

If no server can be located, the interface libraries attempt to start a personal server using other parameters. The START and DBF parameters can be used to start a personal server.




The **StartLine (START)** connection parameter takes a personal database server command line. If a Start parameter is unavailable, Adaptive Server Anywhere attempts to start a personal server on the file indicated by the **DatabaseFile (DBF)**. If both an **EngineName (ENG)** parameter and a **DatabaseFile (DBF)** parameter appear, the **EngineName (ENG)** parameter becomes the name of the server.

Server name caching for faster connections

When the **DoBroadcast (DOBROAD)** communication parameter is set to **DIRECT** or **ALL**, the network library looks for a database server on a network by broadcasting over the network using the **CommLinks (LINKS)** connection parameter.

Tuning the broadcast

The **CommLinks (LINKS)** parameter takes as argument a string listing the protocols to use and, optionally for each protocol, a variety of network communication parameters that tune the broadcast.

 For more information about network communications parameters, see "Network communications parameters" on page 189.

Caching server information

Broadcasting over large networks searching for a server of a specific name can be time-consuming. Leaving the **DoBroadcast (DOBROAD)** communication parameter at the default value (**ALL**) speeds up network connections by saving the protocol the first connection to a server was found on, and its address, to a file and using that information for subsequent connections.

The server information is saved in a cached file named *asasrv.ini*. The file contains a set of sections, each of the following form:

```
[Server name]
Link=protocol_name
Address=address_string
```

The *asasrv.ini* file is located in your Adaptive Server Anywhere executable directory (the same directory as the ODBC/ DBLib DLL).

Note:


It is very important that each server has a unique name. Giving different servers the same name can lead to identification problems.

How the cache is used

If the server name and protocol in the cache match the connection string, Adaptive Server Anywhere tries to connect using the cached address first. If that fails, or if the server name and protocol in the cache do not match the connection string, the connection string information is used to search for the server using a broadcast. If the broadcast is successful, the server name entry in the cache is overwritten. If no server is found, the server name entry in the cache is removed.

Interactive SQL connections

The Interactive SQL utility has a different behavior from the default Embedded SQL behavior when a **CONNECT** statement is issued while already connected to a database. If no database or server is specified in the **CONNECT** statement, Interactive SQL connects to the current database, rather than to the default database. This behavior is required for database reloading operations.

 For more information, see "CONNECT statement [ESQL] [Interactive SQL]" on page 268 of the book *ASA SQL Reference Manual*.

Testing that a server can be found

The *dbping* utility is provided to help in troubleshooting connections. In particular, you can use it to test if a server with a particular name is available on your network.

The *dbping* utility takes a connection string as an option, but by default only those pieces required to locate a server are used. It does not attempt to start a server.

Examples

The following command line tests to see if a server named Waterloo is available over a TCP/IP connection:

```
dbping -c "eng=Waterloo;CommLinks=tcpip"
```

The following command tests to see if a default server is available on the current machine.

```
dbping
```

🔗 For more information on *dbping* options, see "The Ping utility" on page 494.

Using integrated logins

The **integrated login** feature allows you to maintain a single user ID and password for both database connections and operating system and/or network logins. This section describes the integrated login feature.

Operating systems supported

Integrated login capabilities are available for the Windows NT server only. It is possible for Windows 95/98/Me clients, as well as Windows NT/2000/XP clients, to use integrated logins to connect to a network server running on Windows NT/2000/XP.

Benefits of an integrated login

An integrated login is a mapping from one or more Windows NT/2000/XP user profiles to an existing user in a database. A user who has successfully navigated the security for that user profile and logged in to their machine can connect to a database without providing an additional user ID or password.

To accomplish this, the database must be enabled to use integrated logins and a mapping must have been granted between the user profile used to log in to the machine and/or network, and a database user.

Using an integrated login is more convenient for the user and permits a single security system for database and network security. Its advantages include:

- ◆ When connecting to a database using an integrated login, the user does not need to type a user ID or password.
- ◆ If you use an integrated login, the user authentication is done by the operating system, not the database: a single system is used for database security and machine or network security.
- ◆ Multiple user profiles can be mapped to a single database user ID.
- ◆ The name and password used to login to the Windows NT/2000/XP machine do not have to match the database user ID and password.

Caution

Integrated logins offer the convenience of a single security system but there are important security implications which database administrators should be familiar with.

🔗 For more information about security and integrated logins, see "Security concerns: unrestricted database access" on page 87.

Using integrated logins

Several steps must be implemented in order to connect successfully via an integrated login.

❖ To use an integrated login:

- 1 Enable the integrated login feature in a database by setting the value of the LOGIN_MODE database option to either Mixed or Integrated (the option is case insensitive), in place of the default value of Standard. This step requires DBA authority).
- 2 Create an integrated login mapping between a user profile and an existing database user. This can be done using a SQL statement or a wizard in Sybase Central. In Sybase Central, all users with integrated login permission appear in the Integrated Logins folder.
- 3 Connect from a client application in such a way that the integrated login facility is triggered.

Each of these steps is described in the sections below.

Enabling the integrated login feature

The LOGIN_MODE database option determines whether the integrated login feature is enabled. As database options apply only to the database in which they are found, different databases can have a different integrated login setting even if they are loaded and running within the same server.

The LOGIN_MODE database option accepts one of following three values (which are case insensitive).

- ◆ **Standard** With this setting, integrated logins are not permitted. This is the default setting. An error occurs if an integrated login connection is attempted.
- ◆ **Mixed** With this setting, both integrated logins and standard logins are allowed.
- ◆ **Integrated** With this setting, all logins to the database must be made using integrated logins.

Caution

Setting the LOGIN_MODE database option to Integrated restricts connections to only those users who have been granted an integrated login mapping. Attempting to connect using a user ID and password generates an error. The only exception to this is users with DBA authority (full administrative rights).

Example

The following SQL statement sets the value of the LOGIN_MODE database option to Mixed, allowing both standard and integrated login connections:

```
SET OPTION Public.LOGIN_MODE = Mixed
```

Creating an integrated login

User profiles can only be mapped to an existing database user ID. When that database user ID is removed from the database, all integrated login mappings based on that database user ID are automatically removed.

A user profile does not have to exist for it to be mapped to a database user ID. More than one user profile can be mapped to the same user ID.

Only users with DBA authority are able to create or remove an integrated login mapping.

An integrated login mapping is made either using a wizard in Sybase Central or a SQL statement.

❖ To map an integrated login (Sybase Central):

- 1 Connect to a database as a user with DBA authority.
- 2 Open the Integrated Logins folder for the database and double-click Add Integrated Login.
- 3 On the first page of the wizard, specify the name of the system (computer) user for whom the integrated login is to be created.

Also, select the database user ID this user maps to. The wizard displays the available database users. You must select one of these. You cannot add a new database user ID.

- 4 Follow the remaining instructions in the wizard.

❖ To map an integrated login (SQL):

- 1 Connect to a database with DBA authority.
- 2 Execute a GRANT INTEGRATED LOGIN TO statement.

Example

The following SQL statement allows Windows NT users **fran_whitney** and **matthew_cobb** to log in to the database as the user **DBA**, without having to know or provide the DBA user ID or password.

```
GRANT INTEGRATED LOGIN  
TO fran_whitney, matthew_cobb  
AS USER DBA
```

For more information, see "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*.

Revoking integrated login permission

You can remove an integrated login mapping using either Sybase Central or Interactive SQL.

❖ To revoke an integrated login permission (Sybase Central):

- 1 Connect to a database with DBA authority.
- 2 Open the Integrated Logins folder.
- 3 In the right pane, right-click the user/group you would like to remove the integrated login permission from and choose Delete from the popup menu.

❖ To revoke an integrated login permission (SQL):

- 1 Connect to a database with DBA authority.
- 2 Execute a REVOKE INTEGRATED LOGIN FROM statement.

Example

The following SQL statement removes integrated login permission from the Windows NT user pchin.

```
REVOKE INTEGRATED LOGIN  
FROM pchin
```

For more information, see the "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*.

Connecting from a client application

A client application can connect to a database using an integrated login in one of the following ways:

- ◆ Set the **Integrated (INT)** parameter in the list of connection parameters to YES.
- ◆ Specify neither a user ID nor a password in the connection string or Connect dialog. This method is available only for Embedded SQL applications, including the Adaptive Server Anywhere administration utilities.

If Integrated=YES is specified in the connection string, an integrated login is attempted. If the connection attempt fails and the LOGIN_MODE database option is set to Mixed, the server attempts a standard login.

If an attempt to connect to a database is made without providing a user ID or password, an integrated login is attempted. The attempt succeeds or fails depending on whether the current user profile name matches an integrated login mapping in the database.

Interactive SQL Examples

For example, a connection attempt using the following Interactive SQL statement will succeed, providing the user has logged on with a user profile name that matches a integrated login mapping in a default database of a server:

```
CONNECT USING 'INTEGRATED=yes'
```

The Interactive SQL statement

```
CONNECT
```

can connect to a database if all the following are true:

- ◆ A server is currently running.
- ◆ The default database on the current server is enabled to accept integrated login connections.
- ◆ An integrated login mapping has been created that matches the current user's user profile name.
- ◆ If the user is prompted with a dialog by the server for more connection information (such as occurs when using the Interactive SQL utility), the user clicks OK *without* providing more information.

Integrated logins via ODBC

A client application connecting to a database via ODBC can use an integrated login by including the **Integrated (INT)** parameter among other attributes in its Data Source configuration.

Setting the attribute `Integrated=YES` in an ODBC data source causes database connection attempts using that DSN to attempt an integrated login. If the `LOGIN_MODE` database option is set to Standard, the ODBC driver prompts the user for a database user ID and password.

Security concerns: unrestricted database access

The integrated login feature works using the login control system of Windows NT/2000/XP in place of the Adaptive Server Anywhere security system. Essentially, the user passes through the database security if they can log in to the machine hosting the database, and if other conditions, outlined in "Using integrated logins" on page 83, are met.

If the user successfully logs in to the Windows NT/2000/XP server as "dsmith", they can connect to the database without further proof of identification provided there is either an integrated login mapping or a default integrated login user ID.

When using integrated logins, database administrators should give special consideration to the way Windows NT/2000/XP enforces login security in order to prevent unwanted access to the database.

In particular, be aware that by default a "Guest" user profile is created and enabled when Windows NT Workstation or Server is installed.

Caution

Leaving the user profile Guest enabled can permit unrestricted access to a database that is hosted by that server.

If the Guest user profile is enabled and has a blank password, any attempt to log in to the server will be successful. It is not required that a user profile exist on the server, or that the login ID provided have domain login permissions. Literally any user can log in to the server using any login ID and any password: they are logged in by default to the Guest user profile.

This has important implications for connecting to a database with the integrated login feature enabled.

Consider the following scenario, which assumes the Windows NT server hosting a database has a Guest user profile that is enabled with a blank password.

- ◆ An integrated login mapping exists between the user **fran_whitney** and the database user ID **DBA**. When the user **fran_whitney** connects to the server with her correct login ID and password, she connects to the database as **DBA**, a user with full administrative rights.

But anyone else attempting to connect to the server as **fran_whitney** will successfully log in to the server regardless of the password they provide because Windows NT will default that connection attempt to the Guest user profile. Having successfully logged in to the server using the **fran_whitney** login ID, the unauthorized user successfully connects to the database as **DBA** using the integrated login mapping.

Disable the Guest user profile for security

The safest integrated login policy is to disable the Guest user profile on any Windows NT machine hosting an Adaptive Server Anywhere database. This can be done using the Windows NT User Manager utility.

Setting temporary public options for added security

Setting the value of the LOGIN_MODE option for a given database to Mixed or Integrated using the following SQL statement permanently enables integrated logins for that database.

```
SET OPTION Public.LOGIN_MODE = Mixed
```


If the database is shut down and restarted, the option value remains the same and integrated logins are still enabled.

Changing the LOGIN_MODE option temporarily will still allow user access via integrated logins. The following statement will change the option value temporarily:

```
SET TEMPORARY OPTION Public.LOGIN_MODE = Mixed
```

If the permanent option value is Standard, the database will revert to that value when it is shut down.

Setting temporary public options can be considered an additional security measure for database access since enabling integrated logins means that the database is relying on the security of the operating system on which it is running. If the database is shut down and copied to another machine (such as a user's machine) access to the database reverts to the Adaptive Server Anywhere security model and not the security model of the operating system of the machine where the database has been copied.

 For more information on using the SET OPTION statement see "SET OPTION statement" on page 539 of the book *ASA SQL Reference Manual*.

Network aspects of integrated logins

If the database is located on a network server, then one of two conditions must be met for integrated logins to be used:

- ◆ The user profile used for the integrated login connection attempt must exist on both the local machine and the server. As well as having identical user profile names on both machines, the passwords for both user profiles must also be identical.

For example, when the user jsmith attempts to connect using an integrated login to a database loaded on a network server, identical user profile names and passwords must exist on both the local machine and application server hosting the database. jsmith must be permitted to login to both the local machine and the server hosting the network server.

- ◆ If network access is controlled by a Microsoft Domain, the user attempting an integrated login must have domain permissions with the Domain Controller server and be logged in to the network. A user profile on the network server matching the user profile on the local machine is not required.

Creating a default integrated login user

A default integrated login user ID can be created so that connecting via an integrated login will be successful even if no integrated login mapping exists for the user profile currently in use.

For example, if no integrated login mapping exists for the user profile name JSMITH, an integrated login connection attempt will normally fail when JSMITH is the user profile in use.

However, if you create a user ID named Guest in a database, an integrated login will successfully map to the Guest user ID if no integrated login mapping explicitly identifies the user profile JSMITH.

The default integrated login user permits anyone attempting an integrated login to successfully connect to a database if the database contains a user ID named Guest. The authorities granted to the Guest user ID determine the permissions and authorities granted to the newly-connected user.

CHAPTER 3

Client/Server Communications

About this chapter

Each network environment has its own peculiarities. This chapter describes those aspects of network communication that are relevant to the proper functioning of your database server, and provides some tips for diagnosing network communication problems. It describes how networks operate, and provides hints on running the network database server under each protocol.

Network database server only
The material in this chapter applies only to the network server. You do not need to read this chapter if you are using the personal database server.

Contents

Topic	Page
Supported network protocols	92
Using the TCP/IP protocol	93
Using the SPX protocol	96
Using Named Pipes	97
Adjusting communication compression settings to improve performance	98
Troubleshooting network communications	101

Supported network protocols

Properly configured Adaptive Server Anywhere servers run under the following networks and protocols:

- ◆ **Windows NT/2000/XP and Windows 95/98/Me** TCP/IP or SPX protocols.
- ◆ **Windows CE** TCP/IP protocol.
- ◆ **NetWare** All Novell networks using the SPX or TCP/IP protocols.
- ◆ **UNIX** TCP/IP protocol.

In addition, the Named Pipes protocol is provided for communication from a client application running on the same machine but built for a different operating system.

- ◆ **Windows NT/2000/XP** Named Pipes is used for same-machine communication between any application and the server, but is generally not recommended for this purpose. Named Pipes is not used for network communications.

The client library for each platform supports the same protocols as the corresponding server. In order for Adaptive Server Anywhere to run properly, the network protocols (TCP/IP and/or SPX) must be installed and configured properly on both the client and server computers.

Using the TCP/IP protocol

TCP/IP is a suite of protocols originally implemented by the University of California at Berkeley for BSD UNIX. TCP/IP has gained widespread use with the expansion of the Internet and the World wide web.

UDP is a transport layer protocol that sits on top of IP. Adaptive Server Anywhere uses UDP on top of IP to do initial server name resolution and TCP for connection and communication after that.

When you use the TCP/IP protocol, you can force strong encryption of client/server communications using ECC_TLS (formerly Certicom) or RSA_TLS encryption technology.

☞ For more information about strong encryption, see "Encrypting client/server communications" on page 398.

Using TCP/IP with Windows

The TCP/IP implementation for Windows NT/2000/XP is written to the Winsock 2.0 standard, and the implementation for Windows 95/98/Me and Windows CE uses the Winsock 1.1 standard.

If you do not have TCP/IP installed, chose TCP/IP Protocol from the Control Panel, Network Settings.

Tuning TCP/IP performance

Increasing the packet size may improve query response time, especially for queries transferring a large amount of data between a client and a server process. You can set the packet size using the `-p` option in the database server command, or by setting the **CommBufferSize (CBSIZE)** connection parameter in your connection profile.

☞ For more information, see "`-p` server option" on page 147, or "CommBufferSize connection parameter" on page 168.

Connecting across a firewall

There are restrictions on connections when the client application is on one side of a firewall, and the server is on the other. Firewall software filters network packets according to network port. Also, it is common to disallow UDP packets from crossing the firewall.


When connecting across a firewall, you must use a set of communication parameters in the **CommLinks (LINKS)** connection parameter of your application's connection string.

- ◆ **ClientPort** Set this parameter to a range of allowed values for the client application to use. You can then configure your firewall to allow these packets across. You can use the short form CPort.
- ◆ **Host** Set this parameter to the host name on which the database server is running. You can use the short form IP.
- ◆ **ServerPort** If your database server is not using the default port of 2638, you must specify the port it is using. You can use the short form Port.
- ◆ **DoBroadcast=NONE** Set this parameter to prevent UDP from being used when connecting to the server.

Example

The following connection string fragment restricts the client application to ports 5050 through 5060, and connects to a server named **myeng** running on the machine at address **myhost** using the server port 2020. No UDP broadcast is carried out because of the DoBroadcast option.

```
Eng=myeng;Links=tcip(ClientPort=5050-5060;Host=myhost;Port=2020;DoBroadcast=NONE)
```

 For more information, see the following:


- ◆ "CommLinks connection parameter" on page 169
- ◆ "ClientPort communication parameter" on page 191
- ◆ "ServerPort communication parameter" on page 198
- ◆ "Host communication parameter" on page 194
- ◆ "DoBroadcast communication parameter" on page 192

Connecting on a dial-up network connection

You can use connection and communications parameters to assist with connecting to a database across a dial-up link.

On the client side, you should specify the following communications parameters:

- ◆ **Host parameter** You should specify the host name or IP address of the database server using the **Host (IP)** communication parameter.

 For more information, see "Host communication parameter" on page 194.

- ◆ **DoBroadcast parameter** If you specify the **Host (IP)** communication parameter, there is no need to do a broadcast search for the database server. For this reason, use direct broadcasting.

🔗 For more information, see "DoBroadcast communication parameter" on page 192.

- ◆ **MyIP parameter** You should set **MyIP=NONE** on the client side.

🔗 For more information, see "MyIP communication parameter" on page 195.

- ◆ **TIMEOUT parameter** Set the **TIMEOUT (TO)** communication parameter to increase the time the client will wait while searching for a server.

🔗 For more information, see the "Timeout communication parameter" on page 199.

A typical **CommLinks (LINKS)** connection parameter may look as follows:

```
Links=tcpip(MyIP=NONE;DoBroadcast=DIRECT;Host=server_ip)
```

Encrypting client/server communications over TCP/IP

By default, communication packets are not encrypted; this poses a potential security risk. For greater security, you can force client/server network communications over the TCP/IP port to be encrypted using ECC_TLS (formerly Certicom) or RSA_TLS encryption technology. TCP/IP is the only network protocol that supports this strong encryption.

🔗 For information about forcing encryption of client/server communications from the server, see "-ec server option" on page 135.

🔗 For information about forcing encryption of client/server communications from a particular client, see "Encryption connection parameter" on page 177.

Using the SPX protocol

Connecting via SPX

SPX is a protocol from Novell. Adaptive Server Anywhere for NetWare, Windows NT/2000/XP, and Windows 95/98/Me can all employ the SPX protocol. This section provides some tips for using SPX under different operating systems.


Some machines can use the NetWare bindery. These machines are NetWare servers or Windows NT/2000/XP or Windows 95/98/Me machines where the Client Service for NetWare is installed. A client application on one of these machines does not need to use broadcasts to connect to the server if the server to which it is connecting is also using the bindery.

Applications running on machines not using the bindery must connect using one of the following:

- ◆ **An explicit address** You can specify an explicit address using the **HOST (IP)** communication parameter.

 For more information, see the "Host communication parameter" on page 194.

- ◆ **Broadcast** If a server is not found in the bindery, the client will attempt to find it using a broadcast. This can be disabled with by setting the **DoBroadcast (DOBROAD)** communication parameter to **NONE** (client-side) or **NO** (server-side).

 For more information, see the "DoBroadcast communication parameter" on page 192.

Using Named Pipes

Named Pipes is a facility for interprocess communication. Named Pipes can be used for communication between processes on the same computer or on different computers.

Adaptive Server Anywhere for Windows NT/2000/XP uses local Named Pipes for same-machine communication.

 For more information, see "Supported network protocols" on page 92.

Adaptive Server Anywhere does not use Named Pipes for client/server communications between different machines.


Adjusting communication compression settings to improve performance

Enabling compression for one or all connections, as well as setting the minimum size at which packets are compressed can significantly improve Adaptive Server Anywhere performance in some circumstances.

To determine if enabling compression will help in your particular situation, we recommend that you conduct a performance analysis on your particular network and using your particular application before using communication compression in a production environment. Performance results will vary according to the type of network you are using, your applications, and the data you transfer.

The most basic way of tuning compression is as simple as enabling or disabling the **Compression (COMP)** connection parameter on either a connection or server level. More advanced methods of fine tuning compression performance include adjusting the **CommBufferSize (CBSIZE)** and/or the **CompressionThreshold (COMPT)** connection parameters.

Enabling compression increases the quantity of information stored in data packets, thereby reducing the number of packets required to transmit a particular set of data. By reducing the number of packets, the data can be transmitted more quickly.

 For more information about performance analysis, see "Performance Monitor statistics" on page 610, or "sa_conn_compression_info system procedure" on page 687 of the book *ASA SQL Reference Manual*.

Enabling Compression

Enabling compression for a connection (or all connections) can significantly improve Adaptive Server Anywhere performance under some circumstances, including:

- ◆ when used over slow networks such as some wireless networks, some modems, serial links and some WANs.
- ◆ when used in conjunction with Adaptive Server Anywhere encryption over a slow network with built-in compression, since packets are compressed before they are encrypted.

Enabling compression, however, can sometimes also cause slower performance. For instance,

- ◆ communication compression uses more memory and more CPU. It may cause slower performance, especially for LANs and other fast networks.

- ◆ most modems and some slow networks already have built-in compression. In these cases, Adaptive Server Anywhere communication compression will not likely provide additional performance benefits unless you are also encrypting the data.

☞ For more information about compression, see "Compress connection parameter" on page 170, "-pc server option" on page 147.

Modifying the CommBufferSize setting

While **CommBufferSize (CBSIZE)** is not a compression parameter in itself, adjusting this parameter can benefit compression, especially if you are transferring a large amount of data between a client and server. A larger packet size can improve performance for multi-row fetches and fetches of larger rows, as well as inserting or retrieving BLOBs.

As always, however, there are tradeoffs with performance improvements. Since each connection has its own pool of buffers, a large buffer size also increases the memory usage. Be sure to analyze your particular situation to make sure that the benefits of increasing the CommBufferSize outweigh the costs.

☞ For more information see "Tuning TCP/IP performance" on page 93, the "CommBufferSize connection parameter" on page 168, or the "-p server option" on page 147.

Modifying the compression threshold

You can also adjust the compression threshold to improve Adaptive Server Anywhere performance. For most networks, the compression threshold does not need to be changed.

When compression is enabled, individual packets may or may not be compressed, depending on their size. For example, Adaptive Server Anywhere does not compress packets smaller than the compression threshold, even if communication compression is enabled. As well, small packets (less than about 100 bytes) usually do not compress at all. Since CPU time is required to compress packets, attempting to compress small packets could actually decrease performance.

Generally speaking, lowering the compression threshold value may improve performance on very slow networks, while raising the compression threshold may improve performance by reducing CPU usage. However, since lowering the compression threshold value will increase CPU usage on both the client and server, a performance analysis should be done to determine whether or not changing the compression threshold is beneficial.

☞ For more information see "CompressionThreshold connection parameter" on page 172 and "-pt server option" on page 148.

❖ To adjust Adaptive Server Anywhere compression settings:

- 1 Enable communication compression.

Large data transfers with highly compressible data and larger packet sizes tend to get the best compression rates.

🔗 For more information about enabling compression, see "Compress connection parameter" on page 170 and "-pc server option" on page 147.

2 Adjust the CommBufferSize setting.

Increasing Adaptive Server Anywhere's packet size can improve compression performance.

🔗 For more information about adjusting the CommBufferSize setting, see "CommBufferSize connection parameter" on page 168 or "-p server option" on page 147.

3 Adjust the CompressionThreshold setting.

Lowering the compression threshold value may improve performance on very slow networks, while raising the compression threshold may improve performance by reducing CPU usage.

🔗 For more information about adjusting the **CompressionThreshold (COMPT)** connection parameter, see "CompressionThreshold connection parameter" on page 172 and "-pt server option" on page 148.

Troubleshooting network communications


Network software involves several different components, increasing the likelihood of problems. Although we provide some tips concerning network troubleshooting here, the primary source of assistance in network troubleshooting should be the documentation and technical support for your network communications software, as provided by your network communications software vendor.


Ensure that you are using compatible protocols

Ensure that the client and the database server are using the same protocol. The `-x` option for the server selects a list of protocols for the server to use, and the **CommLinks (LINKS)** connection parameter does the same for the client application.

You can use these options to ensure that each application is using the same protocol.

By default, both the database server and client library use all available protocols. The server supports client requests on any active protocol, and the client searches for a server on all active protocols.

 For more information about the `-x` option, see the "`-x` server option" on page 153.

 For information about the **CommLinks (LINKS)** connection parameter, see "CommLinks connection parameter" on page 169.

Ensure that you have current drivers

Old network adapter drivers can be a source of communication problems. You should ensure that you have the latest version of your network adapter. You should be able to obtain current network adapter drivers from the manufacturer or supplier of the card.

Testing the TCP/IP protocol

The **ping** utility can be useful for testing that TCP/IP is installed and configured properly.

Using ping to test the IP layer

Each IP layer has an associated address—a four-integer, period-separated number (such as 191.72.109.12). Ping takes as an argument an IP-address and attempts to send a single packet to the address.

First, determine if your own machine is configured correctly by "pinging" yourself. If your IP-address is 191.72.109.12, you would type:

```
ping 191.72.109.12
```

at the command prompt and wait to see if the packets are routed at all. If they are, the output will appear similar to the following:

```
c:> ping 191.72.109.12

Pinging 191.72.109.12 with 32 bytes of data:
Reply from 191.72.109.12: bytes=32 time<.10ms TTL=32
Reply from 191.72.109.12: bytes=32 time<.10ms TTL=32
Reply from 191.72.109.12: bytes=32 time<.10ms TTL=32
...
```

If this works, it means that the computer is able to route packets to itself. This is reasonable assurance that the IP layer is set up correctly. You could also ask someone else running TCP/IP for their IP address and try pinging them.

You should ensure that you can ping the computer running the database server from the client computer before proceeding.

Diagnosing wiring problems

Faulty network wiring or connectors can cause problems that are difficult to track down. Try recreating problems on a similar machine with the same configuration. If a problem occurs on only one machine, it may be a wiring problem or a hardware problem.

A checklist of common problems


The following list presents some common problems and their solutions.

If you receive the message Unable to start—server not found when trying to start the client, the client cannot find the database server on the network.

Check for the following problems:


- ◆ Under the TCP/IP protocol, clients search for database servers by broadcasting a request. Such broadcasts will typically not pass through gateways, so any database server on a machine in another (sub)network, will not be found. If this is the case, you must supply the host name of the machine on which the server is running using the `-x` option.


- ◆ A firewall between the client and server may be preventing the connection. For more information, see "Connecting across a firewall" on page 93.
- ◆ Your network drivers are not installed properly or the network wiring is not installed properly.
- ◆ If you receive the message Unable to initialize any communication links, no link can be established. The probable cause is that your network drivers have not been installed. The server tries to start communication links using all available protocols unless you have specified otherwise using the `-x` option. Check your network documentation to find out how to install the driver you wish to use.

 For more information about network communication parameters, see "Network communications parameters" on page 189.

Adjusting timeout values

If you are experiencing problems with connections unexpectedly terminating, consider adjusting either the liveness or the idle timeout values.

 For more information about liveness timeout values, see "LivenessTimeout connection parameter" on page 182, and the "`-tl` server option" on page 151.

 For more information about connection timeouts, see the "Idle connection parameter" on page 181, and the "`-ti` server option" on page 150.


CHAPTER 4

Adaptive Server Anywhere as an Open Server

About this chapter

Adaptive Server Anywhere can appear to client applications as an Open Server. This feature enables Sybase Open Client applications to connect natively to Adaptive Server Anywhere databases.

This chapter describes how to use Adaptive Server Anywhere as an Open Server, and how to configure Open Client and Adaptive Server Anywhere to work together.

 For information on developing Open Client applications for use with Adaptive Server Anywhere, see "The Open Client Interface" on page 353 of the book *ASA Programming Guide*.

Contents

Topic	Page
Open Clients, Open Servers, and TDS	106
Setting up Adaptive Server Anywhere as an Open Server	108
Configuring Open Servers	110
Characteristics of Open Client and jConnect connections	116

Open Clients, Open Servers, and TDS

This chapter describes how Adaptive Server Anywhere fits into the Sybase Open Client/Open Server client/server architecture. This section describes the key concepts of this architecture, and provides the conceptual background for the rest of the chapter.

If you simply wish to use a Sybase application with Adaptive Server Anywhere, you do not need to know any details of Open Client, Open Server, or TDS. However, an understanding of how these pieces fit together may be helpful for configuring your database and setting up applications. This section explains how the pieces fit together, but avoids any discussion of the internal features of the pieces.

Open Clients and Open Servers

Adaptive Server Anywhere and other members of the Adaptive Server family act as **Open Servers**. This means that you can develop client applications using the **Open Client** libraries available from Sybase. Open Client includes both the Client Library (CT-Library) and the older DB-Library interfaces.

Tabular Data Stream

Open Clients and Open Servers exchange information using an application protocol called the **tabular data stream** (TDS). All applications built using the Sybase Open Client libraries are also TDS applications because the Open Client libraries handle the TDS interface. However, some applications (such as Sybase jConnect) are TDS applications even though they do not use the Sybase Open Client libraries—they communicate directly using TDS protocol.

While many Open Servers use the Sybase Open Server libraries to handle the interface to TDS, some applications have a direct interface to TDS of their own. Sybase Adaptive Server Enterprise and Adaptive Server Anywhere both have internal TDS interfaces. They appear to client applications as an Open Server, but do not use the Sybase Open Server libraries.

Programming interfaces and application protocols

Adaptive Server Anywhere supports two application protocols. Open Client applications and other Sybase applications such as Replication Server and OmniConnect use TDS. ODBC and Embedded SQL applications use a separate application protocol specific to Adaptive Server Anywhere.

TDS uses TCP/IP

Application protocols such as TDS sit on top of lower-level communications protocols that handle network traffic. Adaptive Server Anywhere supports TDS only over the TCP/IP network protocol. In contrast, the Adaptive Server Anywhere-specific application protocol supports several network protocols, as well as a shared memory protocol designed for same-machine communication.


Sybase applications and Adaptive Server Anywhere

The ability of Adaptive Server Anywhere to act as an Open Server enables Sybase applications such as Replication Server and OmniConnect to work with Adaptive Server Anywhere.

Replication Server support

The Open Server interface enables support for Sybase Replication Server: Replication Server connects through the Open Server interface, enabling databases to act as replicate sites in Replication Server installations.

For your database to act as a primary site in a Replication Server installation, you must also use the Replication Agent for Sybase Adaptive Server Anywhere, also called a **Log Transfer Manager**.

 For information on the Replication Agent, see "Replicating Data with Replication Server" on page 405.

OmniConnect support

Sybase OmniConnect provides a unified view of disparate data within an organization, allowing users to access multiple data sources without having to know what the data looks like or where to find it. In addition, OmniConnect performs heterogeneous joins of data across the enterprise, enabling cross-platform table joins of targets such as DB2, Sybase Adaptive Server Enterprise, Oracle, and VSAM.

Using the Open Server interface, Adaptive Server Anywhere can act as a data source for OmniConnect.

Setting up Adaptive Server Anywhere as an Open Server

This section describes how to set up an Adaptive Server Anywhere server to receive connections from Open Client applications.

System requirements

There are separate requirements at the client and server for using Adaptive Server Anywhere as an Open Server.

Server-side requirements

You must have the following elements at the server side to use Adaptive Server Anywhere as an Open Server:

- ◆ **Adaptive Server Anywhere server components** You must use the network server (*dsrv8.exe*) if you want to access an Open Server over a network. You can use the personal server (*dbeng8.exe*) as an Open Server only for connections from the same machine.
- ◆ **TCP/IP** You must have a TCP/IP protocol stack to use Adaptive Server Anywhere as an Open Server, even if you are not connecting over a network.

Client-side requirements

You need the following elements to use Sybase client applications to connect to an Open Server (including Adaptive Server Anywhere):

- ◆ **Open Client components** The Open Client libraries provide the network libraries your application needs to communicate via TDS, if your application uses Open Client.
- ◆ **jConnect** If your application uses JDBC, you need jConnect and a Java runtime environment.
- ◆ **DSEdit** You need *DSEdit*, the directory services editor, to make server names available to your Open Client application. On UNIX platforms, this utility is called *sybinit*.

DSEdit is not included with SQL Anywhere Studio, but is included with Open Server software.

Starting the database server as an Open Server

If you wish to use Adaptive Server Anywhere as an Open Server, you must ensure that you start it using the TCP/IP protocol. By default, the server starts all available communications protocols, but you can limit the protocols started by listing them explicitly in the command. For example, the following commands are both valid:

```
dbsrv8 -x tcpip,spx asademo.db
```

```
dbsrv8 -x tcpip -n myserver asademo.db
```

The first command uses both TCP/IP and SPX protocols, of which TCP/IP is available for use by Open Client applications. The second line uses only TCP/IP.

You can use the personal database server as an Open Server for communications on the same machine because it supports the TCP/IP protocol.

The server can serve other applications through the TCP/IP protocol or other protocols using the Adaptive Server Anywhere-specific application protocol at the same time as serving Open Client applications over TDS.

Port numbers


Every application using TCP/IP on a machine uses a distinct TCP/IP **port** so that network packets end up at the right application. The default port for Adaptive Server Anywhere is port 2638. It is recommended that you use the default port number as Adaptive Server Anywhere has been granted that port number by the Internet Assigned Numbers Authority (IANA). If you wish to use a different port number, you can specify which one using the **ServerPort (PORT)** communication parameter:

```
dbsrv8 -x tcpip(ServerPort=2629) -n myserver asademo.db
```

You may also need to supply an EngineName if more than one local database server is running, or if you wish to connect to a network server.

Open Client settings

To connect to this server, the interfaces file at the client machine must contain an entry specifying the machine name on which the database server is running, and the TCP/IP port it uses.

 For details on setting up the client machine, see "Configuring Open Servers" on page 110.

Configuring Open Servers

Adaptive Server Anywhere can communicate with other Adaptive Servers, Open Server applications, and client software on the network. Clients can talk to one or more servers, and servers can communicate with other servers via remote procedure calls. For products to interact with one another, each needs to know where the others reside on the network. This network service information is stored in the interfaces file.

The interfaces file

The interfaces file is usually named *SQL.ini* on PC operating systems and *interfaces*, or *interfac* on UNIX operating systems.


Like an address book, the interfaces file lists the name and address of every database server known to Open Client applications on your machine. When you use an Open Client program to connect to a database server, the program looks up the server name in the interfaces file and then connects to the server using the address.

The name, location, and contents of the interfaces file differ between operating systems. Also, the format of the addresses in the interfaces file differs between network protocols.

When you install Adaptive Server Anywhere, the setup program creates a simple interfaces file that you can use for local connections to Adaptive Server Anywhere over TCP/IP. It is the System Administrator's responsibility to modify the interfaces file and distribute it to users so that they can connect to Adaptive Server Anywhere over the network.

Using the DSEdit utility

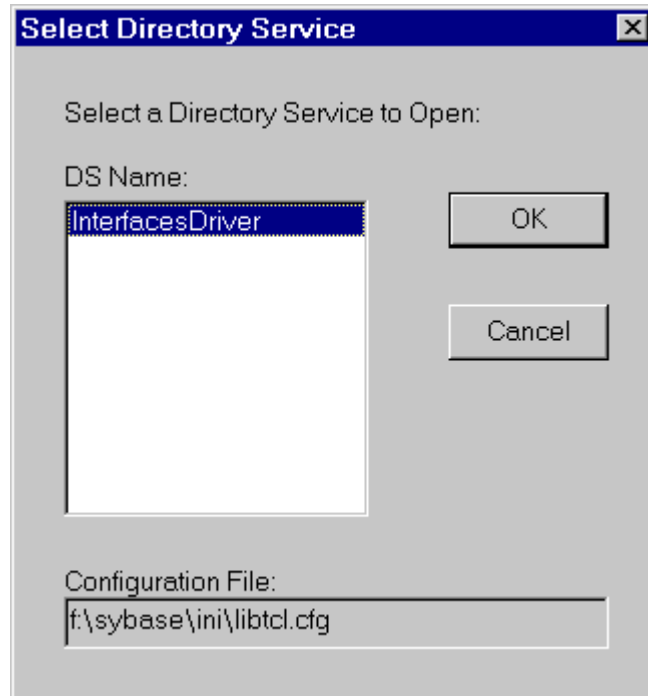
The *DSEdit* utility is a Windows utility that allows you to configure the interfaces file (*SQL.ini*). The following sections explain how to use the *DSEdit* utility to configure the interfaces file.

 These sections describe how to use *DSEdit* for those tasks required for Adaptive Server Anywhere. It is not complete documentation for the *DSEdit* utility. For more information on *DSEdit*, see the *Utility Programs* book for your platform, included with other Sybase products.

Starting DSEdit

The *DSEdit* executable is in the *SYBASE\bin* directory, which is added to your path on installation. You can start *DSEdit* either from the command prompt or from the Explorer in the standard fashion.

When you start *DSEdit*, the Select Directory Service dialog appears.



Opening a Directory Services session

The Select Directory Service window allows you to open a session with a directory service. You can open a session to edit the interfaces file (*SQL.ini*), or any directory service that has a driver listed in the *libtcl.cfg* file.

❖ To open a session:

- ◆ Click the local name of the directory service you want to connect to, as listed in the DS Name box, and click OK.

For Adaptive Server Anywhere, select the Interfaces Driver.

SYBASE environment variable must be set

The *DSEdit* utility uses the SYBASE environment variable to locate the *libtcl.cfg* file. If the SYBASE environment variable is incorrect, *DSEdit* cannot locate the *libtcl.cfg* file.



You can add, modify, or delete entries for servers, including Adaptive Server Anywhere servers, in this window.

Adding a server entry

❖ **To add a server entry:**

- 1 Choose Add from the Server Object menu.

The Input Server Name window appears.

- 2 Type a server name in the Server Name box, and click OK to enter the server name.

The server name entry must match the name you used to start Adaptive Server Anywhere with. The server *address* is used to identify and locate the server. The server name field is an identifier for Open Client. For Adaptive Server Anywhere, if the server has more than one database loaded, the *DSEdit* server name entry identifies which database to use.

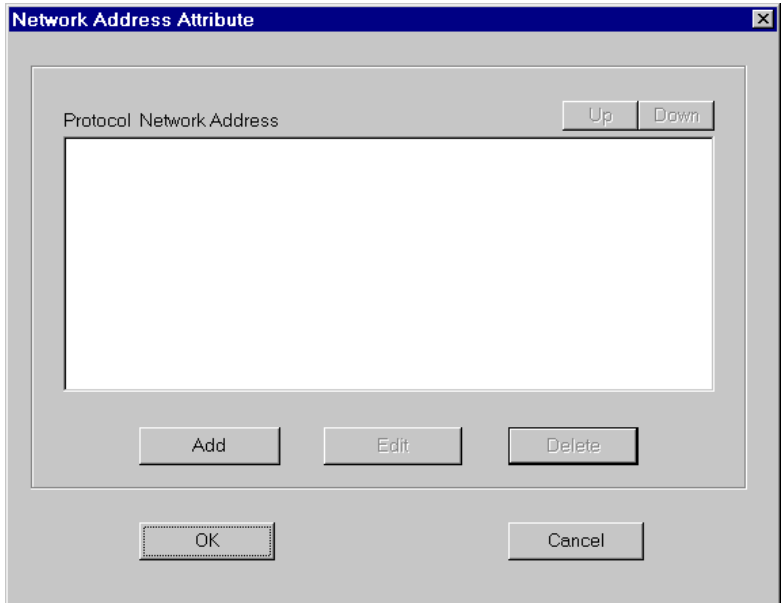
The server entry appears in the Server box. To specify the attributes of the server, you must modify the entry.

Adding or changing the server address

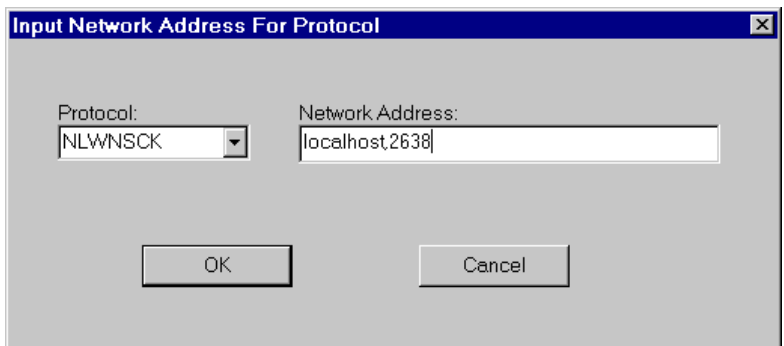
Once you have entered a Server Name, you need to modify the Server Address to complete the interfaces file entry.

❖ **To enter a Server Address:**

- 1 Select a server entry in the Server box.
- 2 Right-click the Server Address in the Attributes box.
- 3 Choose Modify Attribute from the popup menu. A window appears, showing the current value of the address. If you have no address entered, the box is empty.



- 4 Click Add. The Network Address for Protocol window appears.
- 5 Select NLWNSCK from the Protocol list box (this is the TCP/IP protocol) and enter a value in the Network Address text box.



For TCP/IP addresses, use one of the following two forms:

- ◆ computer name,port number
- ◆ IP-address,portnumber

The address or computer name is separated from the port number by a comma.

Machine name

A name (or an IP address) identifies the machine on which the server is running. On Windows operating systems, you can find the machine name in Network Settings, in the Control Panel.

If your client and server are on the same machine, you must still enter the machine name. In this case, you can use **localhost** to identify the current machine.

Port Number

The port number you enter must match the port number you used to start the Adaptive Server Anywhere database server with, as described in "Starting the database server as an Open Server" on page 109. The default port number for Adaptive Server Anywhere servers is 2638. This number has been assigned to Adaptive Server Anywhere by the Internet Adapter Number Authority (IANA), and use of this port is recommended unless you have good reasons for explicitly using another port.

The following are valid server address entries:

```
elora,2638
123.85.234.029,2638
```

Verifying the server address

You can verify your network connection using the Ping command from the Server Object menu.

Database connections not verified

Verifying a network connection confirms that a server is receiving requests on the machine name and port number specified. It does not verify anything about database connections.

❖ To ping a server:

- 1 Ensure that the database server is running.
- 2 Click the server entry in the Server box of the *DSEdit* session window.
- 3 Choose Ping from the Server Object menu. The Ping window appears.
- 4 Select the address you want to ping. Click Ping.

A message box appears, notifying you whether or not the connection is successful. A message box for a successful connection states that both Open Connection and Close Connection succeeded.

Renaming a server entry

You can rename server entries from the *DSEdit* session window.

❖ To rename a server entry:

- 1 Select a server entry in the Server box.
- 2 Choose Rename from the Server Object menu. The Input Server Name window appears.
- 3 Type a new name for the server entry in the Server Name box.
- 4 Click OK to make the change.

Deleting server entries

You can delete server entries from the *DSEdit* session window.

❖ To delete a server entry:

- 1 Click a server entry in the Server box.
- 2 Choose Delete from the Server Object menu.

Configuring servers for JDBC

The JDBC connection address (URL) contains all the information required to locate the server.

 For information on the JDBC URL, see "Supplying a URL for the server" on page 138 of the book *ASA Programming Guide*.

Characteristics of Open Client and jConnect connections

When Adaptive Server Anywhere is serving applications over TDS, it automatically sets relevant database options to values compatible with Adaptive Server Enterprise default behavior. These options are set temporarily, for the duration of the connection only. The client application can override them at any time.

Default settings

The database options set on connection using TDS include:

Option	Set to
ALLOW_NULLS_BY_DEFAULT	OFF
ANSINULL	OFF
ANSI_BLANKS	ON
ANSI_INTEGER_OVERFLOW	ON
AUTOMATIC_TIMESTAMP	ON
CHAINED	OFF
CONTINUE_AFTER_RAISERROR	ON
DATE_FORMAT	YYYY-MM-DD
DATE_ORDER	MDY
ESCAPE_CHARACTER	OFF
ISOLATION_LEVEL	1
FLOAT_AS_DOUBLE	ON
QUOTED_IDENTIFIER	OFF
TIME_FORMAT	HH:NN:SS.SSS
TIMESTAMP_FORMAT	YYYY-MM-DD HH:NN:SS.SSS
TSQL_HEX_CONSTANT	ON
TSQL_VARIABLES	ON

How the startup options are set

The default database options are set for TDS connections using a system procedure named *sp_tsql_environment*. This procedure sets the following options:

```
SET TEMPORARY OPTION TSQL_VARIABLES='ON' ;
SET TEMPORARY OPTION ANSI_BLANKS='ON' ;
SET TEMPORARY OPTION TSQL_HEX_CONSTANT='ON' ;
```

```
SET TEMPORARY OPTION CHAINED='OFF' ;
SET TEMPORARY OPTION QUOTED_IDENTIFIER='OFF' ;
SET TEMPORARY OPTION ALLOW_NULLS_BY_DEFAULT='OFF' ;
SET TEMPORARY OPTION AUTOMATIC_TIMESTAMP='ON' ;
SET TEMPORARY OPTION ANSINULL='OFF' ;
SET TEMPORARY OPTION CONTINUE_AFTER_RAISERROR='ON' ;
SET TEMPORARY OPTION FLOAT_AS_DOUBLE='ON' ;
SET TEMPORARY OPTION ISOLATION_LEVEL='1' ;
SET TEMPORARY OPTION DATE_FORMAT='YYYY-MM-DD' ;
SET TEMPORARY OPTION TIMESTAMP_FORMAT='YYYY-MM-DD
HH:NN:SS.SSS' ;
SET TEMPORARY OPTION TIME_FORMAT='HH:NN:SS.SSS' ;
SET TEMPORARY OPTION DATE_ORDER='MDY' ;
SET TEMPORARY OPTION ESCAPE_CHARACTER='OFF'
```

Do not edit the `sp_tsql_environment` procedure

Do not alter the `sp_tsql_environment` procedure yourself. It is for system use only.

The procedure sets options only for connections that use the TDS communications protocol. This includes Open Client and JDBC connections using jConnect. Other connections (ODBC and Embedded SQL) have the default settings for the database.

You can change the options for TDS connections.

❖ **To change the option settings for TDS connections:**

- 1 Create a procedure that sets the database options you want. For example, you could use a procedure such as the following:

```
CREATE PROCEDURE my_startup_procedure()
BEGIN
    IF connection_property('CommProtocol')='TDS' THEN
        SET TEMPORARY OPTION QUOTED_IDENTIFIER='OFF' ;
    END IF
END
```

This particular procedure example changes only the `QUOTED_IDENTIFIER` option from the default setting.

- 2 Set the `LOGIN_PROCEDURE` option to the name of a new procedure:

```
SET OPTION LOGIN_PROCEDURE=
'DBA.my_startup_procedure'
```

Future connections will use the procedure. You can configure the procedure differently for different user IDs.

🔗 For more information about database options, see "Database Options" on page 535.

CHAPTER 5

The Database Server

About this chapter

This chapter describes the line options for the Adaptive Server Anywhere database server.

It also contains information for the options of the client executable (provided for compatibility with Version 5 software).

Contents

Topic	Page
The database server	120

The database server

- Function

Start a personal database server or network database server.
- Syntax

`{ dbeng8 | dbsrv8 }
[server-options] [database-file [database-options] ...]`
- NetWare syntax

`load dbsrv8 [server-options] [database-file [database-options] ...]`

Server options

Server Option	Description
<i>@environment-variable</i>	Read in options from an environment variable. See " <i>@environment-variable</i> server option" on page 125.
<i>@filename</i>	Read in options from a configuration file. See " <i>@filename</i> server option" on page 126.
-?	Display usage information. See " -? server option" on page 127.
-b	Run in bulk operations mode. See " -b server option" on page 127.
-c size	Set initial cache size. See " -c server option" on page 127.
-ca 0	Disable dynamic cache sizing [Windows NT/2000/XP, Windows 95/98/Me, UNIX]. See " -ca server option" on page 129.
-ch size	Set the cache size upper limit [Windows NT/2000/XP, Windows 95/98/Me]. See " -ch server option" on page 130.
-cl size	Set the cache size lower limit [Windows NT/2000/XP]. See " -cl server option" on page 130.
-cs	Display cache usage in database server window. See " -cs server option" on page 131.
-ct { + - }	Turn character-set translation on and off [not NetWare or Windows CE]. See " -ct server option" on page 131.
-cw	Enable use of Address Windowing Extensions on Windows 2000 and Windows XP for setting the size of the database server cache. See " -cw server option" on page 131.
-d	Use POSIX I/O [NetWare]. See " -d server option" on page 134.
-ec encryption-options	Enable packet encryption [network server]. See " -ec server option" on page 135.
-ep	Prompt for encryption key. See " -ep server option" on page 137.
-ga	Automatically unload the database after the last connection closed. In addition, shut down after the last database is closed

Server Option	Description
	[Not NetWare]. See "-ga server option" on page 138.
-gb level	Set database process priority class to <i>level</i> [Windows NT/2000/XP]. See "-gb server option" on page 138.
-gc num	Set maximum checkpoint timeout period to <i>num</i> minutes. See "-gc server option" on page 138.
-gd level	Set database starting permission. See "-gd server option" on page 139.
-ge size	Set the stack size for threads that run external functions [not UNIX]. See "-ge server option" on page 140.
-gf	Disable firing of triggers. See "-gf server option" on page 140.
-gk level	Set the permission required to stop the server. See "-gk server option" on page 140.
-gl level	Set the permission required to load or unload data. See "-gl server option" on page 141.
-gm num	Set the maximum number of connections. See "-gm server option" on page 141.
-gn num	Set the maximum number of concurrent requests the database server can handle at one time. See "-gn server option" on page 141.
-gp size	Set the maximum page size to <i>size</i> bytes. See "-gp server option" on page 142.
-gr minutes	Set the maximum recovery time to <i>num</i> minutes. See "-gr server option" on page 143.
-gss size	Set the thread stack size to <i>size</i> bytes [not applicable to Windows]. See "-gss server option" on page 143.
-gt num	Set the number of operating processors used by the database server. See "-gt server option" on page 143.
-gu level	Set the permission level for utility commands: utility_db , all , none , or DBA . See "-gu server option" on page 143.
-gx	Set the number of operating system threads assigned to the database server process. [Windows NT/2000/XP, Windows 95/98/Me]. See "-gx server option" on page 144.
-m	Truncate the transaction log after each checkpoint, for all databases. See "-m server option" on page 144.
-n name	Use <i>name</i> as the name of the database server. Note that the -n switch is positional. See "-n server option" on page 145.
-o filename	Output messages to the specified file. See "-o server option" on page 146.

Server Option	Description
-os <i>size</i>	Limit the size of the log file for messages. See "-os server option" on page 146.
-p <i>packet-size</i>	Set the maximum network packet size [network server]. See "-p server option" on page 147.
-pc	Compress all connections except same-machine connections. See "-pc server option" on page 147.
-pt <i>size_in_bytes</i>	Set the minimum network packet size to compress. See "-pt server option" on page 148.
-q	Quiet mode—suppress output. See "-q server option" on page 148.
-qp	Suppress messages about performance in the database server window. See "-qp server option" on page 148.
-r	Opens database in read-only mode. See "-r server option" on page 149.
-s	Set the syslog facility ID [UNIX]. See "-s server option" on page 149.
-sb { 0 1 }	Specify how the server reacts to broadcasts. See "-sb server option" on page 150.
-sc	Disable the shared memory port, and enable Named Pipes. [Windows NT/2000/XP]. See "-sc server option" on page 150.
-ti <i>minutes</i>	Client idle time before shutdown—default 240 minutes [network server]. See "-ti server option" on page 150.
-tl <i>seconds</i>	Default liveness timeout for clients in seconds—default 120 seconds. See "-tl server option" on page 151.
-tmf	Force transaction manager recovery for distributed transactions [Windows NT/2000/XP]. See "-tmf server option" on page 151.
-tmt <i>milliseconds</i>	Set the reenlistment timeout for distributed transactions [Windows NT/2000/XP]. See "-tmt server option" on page 151.
-tq <i>time</i>	Set quitting time [network server]. See "-tq time server option" on page 152.
-u	Use buffered disk I/O. See "-u server option" on page 152.
-ud	Run as a daemon [UNIX]. See "-ud server option" on page 153.
-ut <i>minutes</i>	Touch temporary files every <i>min</i> minutes [UNIX]. See "-ut server option" on page 153.
-v	Display database server version and stop. See "-v server option" on page 153.

Server Option	Description
-x <i>list</i>	Comma-separated list of communication links to try. See "-x server option" on page 153.
-y	Run as a Windows 95/98/Me service [Windows 95/98/Me]. See "-y server option" on page 155.
-z	Provide diagnostic information on communication links [network server]. See "-z server option" on page 156.
-zl	Turn on capturing of the most recently-prepared SQL statement for each connection. See "-zl server option" on page 156.
-zo <i>filename</i>	Redirect request-level logging information to a separate file. See "-zo server option" on page 156.
-zr { all SQL none }	Turn on logging of SQL operations. The default is NONE. See "-zr server option" on page 157.
-zs <i>size</i>	Limit the size of the log file used for request-level logging. See "-zs server option" on page 157.

Recovery options

Recovery Option	Description
-a <i>filename</i>	Apply the named transaction log file. See "-a recovery option" on page 157.
-f	Force the database to start without a transaction log. See "-f recovery option" on page 158.

Database options

Database Option	Description
-ek <i>key</i>	Specify encryption key. See "-ek database option" on page 159.
-m	Truncate (delete) the transaction log after each checkpoint for the specified database. See "-m database option" on page 159.
-n <i>name</i>	Name the database. Note that the -n switch is positional. See "-n database option" on page 160.
-r	Opens the specified database(s) in read-only mode. Database modifications not allowed. See "-r database option" on page 161.

Description

The **dbeng8** command starts a personal database server. The **dbsrv8** command starts a network database server.

Cache size

The amount of cache memory available to the database server can be a key factor in affecting performance. The database server takes an initial amount of cache memory that is either specified by the `-c` option or is a default value.

🔗 For information on the default cache size, see "`-c` server option" on page 127.

On Windows NT/2000/XP, Windows 95/98/Me, and UNIX, the database server automatically takes more memory for use in the cache as needed, determined by a heuristic algorithm.

🔗 For more information, see "Using the cache to improve performance" on page 152 of the book *ASA SQL User's Guide*.

You can use database options to configure the upper cache limit: see "`-ch` server option" on page 130. As well, you can force the cache to remain at its initial amount: see "`-ca` server option" on page 129.

Server differences

The personal database server has a maximum of ten concurrent connections, uses at most two CPUs for request processing, and does not support network client/server connections.

In addition, there are other minor differences, such as the default permission level that is required to start new databases, or the permissions required to execute the CHECKPOINT statement.

Platform availability

Both personal and network database servers are supplied for each supported operating system, with the following exceptions:

- ◆ **Novell NetWare** Only the network server is supplied.
- ◆ **Windows CE** Only the network server is supplied. The support for TCP/IP in the network server enables you to carry out tasks from your desktop machine, including database management, with Sybase Central.

NetWare notes

In NetWare, the database file and the transaction log file must be on a NetWare volume, and the paths must be fully specified. NetWare allows you to have volumes that span two or more hard disks.

Database file The *database-file* specifies the database filename. If *database-file* is specified without a file extension, Adaptive Server Anywhere looks first for *database-file* with extension *.wrt* (a write file) followed by *database-file* with extension *.db*.

If you use a relative path, it is read relative to the current working directory. You can supply a full path. Also, you can supply a path that conforms to the Universal Naming Convention (UNC) format:

```
\\server\volume\path\file.ext
```

In addition, users of Novell NetWare version 4 and later can use NetWare Directory Services (NDS) volumes, which have the following format:

```
\\treename\volume.org_unit.org\path\file.ext
```

where *volume.org_unit.org* is the name of an NDS volume object.

Caution

The database file must be on the same machine as the database server. Managing a database file that is located on a network drive can lead to file corruption.

Suppressing Windows event log messages

If you run the database server as a Windows service, you can suppress Windows event log entries by setting a registry entry. The registry entry is

```
Software\Sybase\Adaptive Server Anywhere\8.0
```

To control event log entries, set the EventLogMask key, which is of type REG_DWORD. The value is a bit mask containing the internal bit values for the different types of event messages:

```
errors      EVENTLOG_ERROR_TYPE      0x0001
warnings    EVENTLOG_WARNING_TYPE     0x0002
information  EVENTLOG_INFORMATION_TYPE 0x0004
```

For example, if the EventLogMask key is set to zero, no messages appear at all. A better setting would be 1, so that informational and warning messages do not appear, but errors do. The default setting (no entry present) is for all message types to appear.

See also

"Running the Database Server" on page 3

"Network communications parameters" on page 189

Database server options

These options apply to the server as a whole, not just to an individual database.

@environment-variable server option

Function	Read in server options from the supplied environment variable.
Syntax	{ dbsrv8 dbeng8 } @ <i>env-var</i> ...
Applies to	All operating systems and servers.

Description

The environment variable may contain any set of options. For example, this statement (typed all on one line) sets an environment variable that holds options for a database server that starts with a cache size of 4 Mb and loads the sample database.

```
set envvar=-c 4096 "c:\Program Files\Sybase\SQL Anywhere
8\asademo.db"
```

This statement starts the database server using the environment variable.

```
dbsrv8 @envvar
```

Environment variable given priority

If you have both a file and an environment variable with the value of your @ option, the environment variable is used.

@filename server option

Function

Read in server options from the supplied file. The number sign (#) must precede comments.

Syntax

```
{ dbsrv8 | dbeng8 } @filename ...
```


Applies to

All operating systems and servers.

Description

The command file may contain line breaks, and may contain any set of options.

If you want to protect the information in a configuration file (for example, because it contains passwords) you can use the File Hiding [dbfhide] utility to obfuscate the contents of configuration files.

 For more information about the File Hiding utility, see "The File Hiding utility" on page 446.

Examples

The following command file holds a set of options for a server named **myserver** that starts with a cache size of 4 Mb and loads the sample database:

```
-c 4096
-n myserver
"c:\Program Files\Sybase\SQL Anywhere 8\asademo.db"
```

If this configuration file is saved as *c:\config.txt*, it can be used in a command as follows:

```
dbsrv8 @c:\config.txt
```

The following command file contains comments:


```
#This is the server name:
-n MyServer
#These are the protocols:
-x tcpip
#This is the database file
my.db
```

-? server option

Function	Display usage information.
Syntax	{ dbsrv8 dbeng8 } -?
Applies to	All operating systems and servers.
Description	Display a short description of each server option. The database does not carry out any other task.

-b server option

Function	Use bulk operation mode.
Syntax	{ dbsrv8 dbeng8 } -b ...
Applies to	All operating systems and servers.
Description	<p>This is useful for using the Interactive SQL INPUT command to load large quantities of data into a database.</p> <p>The -b option should not be used if you are using LOAD TABLE to bulk load data.</p> <p>When you use this option, the database server allows only one connection by one application. It does not keep a rollback log or a transaction log, and the multi-user locking mechanism is turned off.</p> <p>When you first start the database server after loading data with the -b option, you should use a new log file.</p> <p>Bulk operation mode does not disable the firing of triggers.</p>

-c server option

Function	Set the initial memory reserved for caching database pages and other server information.
Syntax	{ dbsrv8 dbeng8 } -c { <i>integer</i> <i>integer</i> G <i>integer</i> K <i>integer</i> M <i>integer</i> P } ...
Applies to	All operating systems and servers.

Description

The amount of memory available for use as a database server cache is one of the key factors controlling performance. You can set the initial amount of cache memory using the `-c` server option

The more cache memory that can be given the server, the better its performance.

The units **G**, **K**, and **M** can be either lower case or upper case. If **G**, **K**, or **M** is not supplied, any integer less than 10000 is assumed to be in kilobytes, and any integer 10000 or greater is assumed to be in bytes. For example, `-c 4096` means 4096 kb or 4 194 304 bytes, whereas `-c 200 000` means (an unreasonably small) cache of 200 000 bytes.

The unit **P** is a percentage of the physical system memory, and if you use this, the argument is a percentage. You can use `%` as an alternative to **P**, but as most non-UNIX operating systems use `%` as an environment variable escape character, you must escape the `%` character. To use 50 percent of the physical system memory, you would use the following:

```
dbeng8 -c 50%% ...
```

If no `-c` option is provided, the database server computes the initial cache allocation as follows:

- 1 It uses the following operating-system-specific default cache sizes:
 - ◆ **Windows CE** 600K
 - ◆ **Windows NT/2000/XP, Windows 95/98/Me, NetWare** 2 Mb
 - ◆ **UNIX** 8 Mb
- 2 It computes a runtime-specific minimum default cache size, which is the lesser of the following items:
 - ◆ 25% of the machine's physical memory
 - ◆ The sum of the sizes of the main database files specified on the command line. Additional dbspaces apart from the main database files are not included in the calculation. If no files are specified, this value is zero.
- 3 It allocates the greater of the two values computed.

NetWare database server

There is a tradeoff between memory for the database server and memory for the NetWare file system buffers. A larger database server cache will improve database server performance at the expense of NetWare file system performance. If the database server cache is too big, NetWare will report an error that there is insufficient memory for cache buffers.

NetWare memory requirements increase with every new directory and file on the file server. To track memory usage on the NetWare server, load *monitor.nlm* (if it is not already loaded) and select "Resource Utilization". Extra memory for your NetWare server computer could improve database performance and/or file server performance dramatically.

Example

The following example, entered all on one line, starts a server named **myserver** that starts with a cache size of 3 Mb and loads the sample database:

```
dbeng8 -c 3M -n myserver "C:\Program Files\Sybase\SQL
Anywhere 8\asdemo.db"
```

See also

"-ch server option" on page 130

-ca server option**Function**

Enforces a static cache size. The zero argument is required.

Syntax

```
{ dbsrv8 | dbeng8 } -ca 0 ...
```

Applies to

Windows NT/2000/XP, Windows 95/98/Me, UNIX

Description

You can disable automatic cache increase due to high server load by specifying **-ca 0** on the command line. If you do not set **-ca 0**, the database server automatically takes additional cache as needed. The cache size still increases if the database server would otherwise run into the error Fatal Error: dynamic memory exhausted, or if the Java VM requires memory that would lead to a fatal error.

This server option should be used only in the form **-ca 0**. If **-ca** is omitted, automatic cache growth is enabled.

Example

The following example starts a server named **myserver** that has a static cache that is 40% of the available physical memory and loads the sample database:

```
dbsrv8 -c 40P -ca 0 -n myserver "C:\Program
Files\Sybase\SQL Anywhere 8\asdemo.db"
```

See also

"-c server option" on page 127

"-ch server option" on page 130

–ch server option

Function	Sets a maximum cache size, as a limit to automatic cache growth.
Syntax	<code>{ dbsrv8 dbeng8 } -ch { <i>integer</i> <i>integer</i>G <i>integer</i>K <i>integer</i>M <i>integer</i>P }</code> ...
Applies to	Windows NT/2000/XP, Windows 95/98/Me, and UNIX
Description	<p>This option limits the cache that the database server can take during automatic cache growth. By default the upper limit is approximately the lower of 256 Mb and 90% of the physical memory of the machine.</p> <p>☞ For the meaning and usage of the cache size arguments and the G, K, M, and P characters, see "–c server option" on page 127.</p>
Example	<p>The following example starts a server named silver that has a maximum cache size of 2 Mb and loads the sample database:</p> <pre>dbeng8 -ch 2M -n silver "C:\Program Files\Sybase\SQL Anywhere 8\asademo.db"</pre>
See also	<p>"–c server option" on page 127 "–ca server option" on page 129 "–cl server option" on page 130</p>

–cl server option

Function	Sets a minimum cache size as a lower limit to automatic cache resizing.
Syntax	<code>{ dbsrv8 dbeng8 } -cl { <i>integer</i> <i>integer</i>G <i>integer</i>K <i>integer</i>M <i>integer</i>P }</code> ...
Applies to	Windows NT/2000/XP, Windows 95/98/Me, UNIX
Description	<p>This option sets a lower limit to the cache. The default minimum cache size is the initial cache size.</p> <p>☞ For the meaning and usage of the cache size arguments and the G, K, M, and P characters, see "–c server option" on page 127.</p>
Example	<p>The following example starts a server named silver that has a minimum cache size of 5 Mb and loads the sample database:</p> <pre>dbeng8 -cl 5M -n silver "C:\Program Files\Sybase\SQL Anywhere 8\asademo.db"</pre>
See also	<p>"–c server option" on page 127 "–ca server option" on page 129 "–ch server option" on page 130</p>

-cs server option

Function	Display cache size changes in the database server window.
Syntax	{ dbsrv8 dbeng8 } -cs ...
Applies to	All platforms with dynamic cache sizing.
Description	For troubleshooting purposes, display cache information in the database server window whenever the cache size changes.

-ct server option

Function	To turn character set translation on and off.
Syntax	{ dbsrv8 dbeng8 } -ct { + - } ...
Applies to	All operating systems except Windows CE. NetWare supports character set translation only when both the database server and the client use single-byte character sets.
Description	<p>By default, character set translation is turned on. Character set translation converts strings between character sets that represent the same characters, but at different values. This is useful when the client machine and the database use different character sets. Character set translation is disabled with the -ct- server option. If this argument is supplied as -ct+, character set translation is turned on.</p> <p>In version 7.x and earlier of Adaptive Server Anywhere, the + or - values are not accepted: specifying the -ct option enabled character set translation.</p>
Example	<p>The following example starts a server with character set translation disabled and loads the sample database:</p> <pre>dbeng8 -ct- "C:\Program Files\Sybase\SQL Anywhere 8\asademo.db"</pre>
See also	"Turning off character set translation on a database server" on page 292 "Starting a database server using character set translation" on page 291

-cw server option

Function	Enables use of Address Windowing Extensions (AWE) on Windows 2000, Windows XP, and Windows .NET Server for setting the size of the database server cache.
Syntax	{ dbsrv8 dbeng8 } -cw ...
Applies to	Windows 2000, Windows XP, Windows .NET Server and higher.

Description

The amount of memory available for use as a database server cache is one of the key factors controlling performance. Because Windows 2000, Windows XP, and Windows .NET Server support Address Windowing Extensions, you can use the `-cw` option to take advantage of large cache sizes based on the maximum amount of physical memory in the system.

Operating system	Maximum non-AWE cache size	Maximum amount of physical memory supported by Windows
Windows 2000 Professional	1.6 Gb	4 Gb
Windows 2000 Server	2.3 Gb*	4 Gb
Windows 2000 Advanced Server	2.3 Gb*	8 Gb
Windows 2000 Datacenter Server	2.3 Gb*	64 Gb
Windows XP Home Edition	1.6 Gb	2 Gb
Windows XP Professional	1.6 Gb	4 Gb
Windows .NET Web Server	1.6 Gb	2 Gb
Windows .NET Standard Server	1.6 Gb	4 Gb
Windows .NET Enterprise Server	2.3 Gb*	32 Gb
Windows .NET Datacenter Server	2.3 Gb*	64 Gb

*You must boot the operating system using the `/3GB` option to use a cache of this size.

Note

At the time of printing, Windows .NET Web Server, Windows .NET Standard Server, Windows .NET Enterprise Server, and Windows .NET Datacenter Server had not been released.

When using an AWE cache, almost all of the available physical memory in the system can be allocated for the cache.

If you can set a cache of the desired size using a non-AWE cache, this is recommended because AWE caches allocate memory that can only be used by the database server. This means that while the database server is running, the operating system and other applications cannot use the memory allocated for the database server cache. AWE caches do not support dynamic cache sizing. Therefore, if an AWE cache is used and you specify the `-ch` or `-cl` options to set the upper and lower cache size, they are ignored.

☞ For information about specifying the cache size, see "`-c` server option" on page 127.

To start a database server with an AWE cache, you must do the following

- ◆ Have at least 3 Gb of memory available on your system.
- ◆ If your system has between 2 Gb and 16 Gb of memory, add the `/3GB` option to the Windows boot line in the "[operating systems]" section of the *boot.ini* file.

If your system has more than 16 Gb of memory, do not add the `/3GB` option to the Windows boot line in the "[operating systems]" section of the *boot.ini* file because Windows will not be able to address memory beyond 16 Gb.
- ◆ If your system has more than 4 Gb of memory, add the `/PAE` option to the Windows boot line in the "[operating systems]" section of the *boot.ini* file.
- ◆ Grant the "Lock pages in memory" privilege to the user ID under which the server is run. The following steps explain how to do this on Windows 2000.
 - 1 Log on to Windows as Administrator.
 - 2 From the Start menu, choose Settings►Control Panel.
 - 3 Open the Administrative Tools folder.
 - 4 Double-click Local Security Policy.
 - 5 Open Local Policies in the left pane.
 - 6 Double-click User Rights Assignment in the left pane.
 - 7 Double-click the Lock Pages In Memory policy in the right pane.

The Local Security Policy Setting dialog appears.
 - 8 In the Local Security Policy Setting dialog, click Add.

The Select Users or Groups dialog appears.
 - 9 Select the user ID from the list and click Add.
 - 10 In the Local Security Policy Setting dialog, click OK.

- 11 Restart the computer for the setting to take effect.

If you specify the `-cw` option and the `-c` option on the command line, the database server attempts the initial cache allocation as follows:

- 1 The AWE cache is no larger than the cache size specified by the `-c` option. If the value specified by the `-c` option is less than 3 Gb-128 Mb, AWE is not used.
- 2 The AWE cache is no larger than all available physical memory less 64 Mb.
- 3 The AWE cache is no smaller than 3 Gb-128 Mb. If this minimum amount of physical memory is not available, AWE is not used.

When you specify the `-cw` option and do not specify the `-c` option, the database server attempts the initial cache allocation as follows:

- 1 The AWE cache uses 100% of all available memory except for 128 Mb that is left free for the operating system.
- 2 The AWE cache is the sum of the sizes of the main database files specified on the command line. Additional dbspaces apart from the main database files are not included in the calculation. If no files are specified, this value is zero.
- 3 The AWE cache is no smaller than 3 Gb-128 Mb. If this minimum amount of physical memory is not available, an AWE cache is not used.

When the server uses an AWE cache, the cache page size is at least 4 Kb and dynamic cache sizing is disabled.

☞ For more information about dynamic cache sizing, see "Using the cache to improve performance" on page 152 of the book *ASA SQL User's Guide*.

Example

The following example starts a server named **myserver** that starts with a cache size of 12 Gb and loads the sample database:

```
dbeng8 -c 12G -cw asademo.db
```

See also

"`-c` server option" on page 127

-d server option

Function

Use POSIX I/O.

Syntax

```
{ dbsrv8 | dbeng8 } -d ...
```

Applies to

NetWare. The use of this option on Windows platforms has been deprecated.

Description

The `-d` option forces the use of POSIX I/O rather than DFS (Direct File System) I/O. This switch is provided as a workaround for bugs in old versions of DFS. Note that asynchronous I/O is not available when using POSIX I/O.

This option applies to NetWare systems only.

-ec server option**Function**

Encrypt all native Adaptive Server Anywhere packets (DBLib, ODBC, and OLEDB) transmitted to and from all clients. TDS packets are not encrypted.

Syntax

`{ dbsrv8 | dbeng8 } -ec encryption-options ...`

encryption-options:

`{ NONE`

`| SIMPLE`

`| ECC_TLS (CERTIFICATE=filename;
CERTIFICATE_PASSWORD=password)`

`| RSA_TLS (CERTIFICATE=filename;
CERTIFICATE_PASSWORD=password)`

`| ALL }, ...`

Description

The `-ec` option instructs the database server to accept *only* connections from ODBC, OLE DB, or embedded SQL interfaces that are encrypted using one of the specified types. Connections over the TDS protocol, which include Java applications using jConnect, are always accepted, regardless of encryption.

By default, communication packets are not encrypted, which poses a potential security risk. If you are concerned about the security of network packets, use the `-ec` option. Encryption affects performance only marginally. The `-ec` option controls the server's encryption settings and requires one or more of the following parameters in a comma-separated list:

none accepts only connections that are not encrypted.

simple accepts connections that are encrypted with simple encryption. This type of encryption is supported on all platforms, as well as on previous versions of Adaptive Server Anywhere. Simple encryption is not as strong as Certicom encryption.

ECC_TLS Formerly called Certicom encryption, this parameter accepts connections that are encrypted using the elliptic curve-based Certicom encryption technology. To use this type of encryption, both the server and the client must be operating on Solaris, Linux, NetWare, or any supported Windows platform except Windows CE, and the connection must be over the TCP/IP port. UNIX platforms, except for Solaris and Linux, do not recognize the client or server **ECC_TLS** parameter. Specifying **CERTICOM** is accepted to mean **ECC_TLS**. This parameter accepts the following arguments:

- ◆ **certificate** the file name of the certificate. The default value is *sample.crt*.
- ◆ **certificate_password** the password for the certificate named above. The password for *sample.crt* is **tJ1#m6+W**.

RSA_TLS this parameter accepts connections that are encrypted using RSA-based encryption technology. To use this type of encryption, both the server and the client must be operating on Solaris, Linux, NetWare, or any supported Windows platform except Windows CE, and the connection must be over the TCP/IP port. UNIX platforms, except for Solaris and Linux, do not recognize the client or server **RSA_TLS** parameter. This parameter accepts the following arguments:

- ◆ **certificate** the file name of the certificate. The default value is *rsaserver.crt*.
- ◆ **certificate_password** the password for the certificate named above. The password for *rsaserver.crt* is **test**.

Caution

The sample certificate should be used for testing purposes only. The sample certificate provides no security in deployed situations because it and the corresponding password are widely distributed with Sybase software. To protect your system, you must create your own certificate.

You can use the **gencert** utility provided with SQL Anywhere Studio to generate new certificates in Adaptive Server Anywhere.

For information about the **gencert** utility and creating certificates, see "Self-signed certificates" on page 294 of the book *MobiLink Synchronization User's Guide*.

all accepts connections that are not encrypted (none), or encrypted with simple encryption (simple). This is the default.

The *dbtls8.dll* and *dbrsa8.dll* files contain the Certicom code used for encryption and decryption. When you connect to the server, if the appropriate file cannot be found, or if an error occurs, a message appears on the server console in debug mode. The server does not start if the types of encryption specified cannot be initiated.

The client's and the server's encryption settings must match or the connection will fail. The server automatically encrypts client transmissions that request encryption if the **none** parameter is not used with the **-ec** option.

Example 1

```
dbsrv8 -ec simple,certicom(certificate=sample.crt;  
certificate_password=tJl#m6+W) -x tcpip asademo.db
```

Example 2

```
dbsrv8 -ec ecc_tls(certificate=sample.crt;  
certificate_password=tJl#m6+W) -x tcpip asademo.db
```

Example 3

```
dbsrv8 -ec rsa_tls(certificate=rsaserver.crt;  
certificate_password=test) -x tcpip asademo.db
```

See also

"Encryption connection parameter" on page 177

"Encrypting client/server communications" on page 398

-ep server option**Function**

To prompt the user for the encryption key upon starting a strongly encrypted database.

Syntax

```
{ dbsrv8 | dbeng8 } -ep ...
```

Description

The **-ep** option instructs the database server to display a dialog box for the user to enter the encryption key. This server option provides an extra measure of security by never allowing the encryption key to be seen in clear text.

When used with supported tools, this option always prompts the user for the encryption key, even if a key is not necessary. If the dialog box prompt appears and you know that a key is not necessary, you can click Cancel to continue.

When used with the engine, this option prompts the user for the encryption key only if *all* of the following are true:

- ◆ the **-ep** option is specified
- ◆ the engine is either a Win32 personal server, or the engine is just starting up
- ◆ a key is required to start a database
- ◆ the engine is either not a Windows service, or it is a Windows service with the interact with desktop option turned ON (Win32)

- ◆ the engine is not a daemon (UNIX)

Example

```
dbsrv8 -ep -x tcpip myencrypted.db
```

See also

"Encryption connection parameter" on page 177
"Encrypting client/server communications" on page 398
"-ek database option" on page 159

–ga server option

Function

Unload database after last connection is dropped.

Syntax

```
{ dbsrv8 | dbeng8 } -ga ...
```

Applies to

All operating systems except NetWare.

Description

Specifying this option on the network server causes each database to be unloaded after the last connection to it is dropped. In addition to unloading each database after the last connection is dropped, the personal server shuts down when the last database is stopped.

–gb server option

Function

Set the database process priority class.

Syntax

```
{ dbsrv8 | dbeng8 } -gb { idle | normal | high | maximum } ...
```

Applies to

Windows NT/2000/XP

Description

Set the database process priority class. The value **idle** is provided for completeness, and **maximum** may interfere with the running of your computer. **Normal** and **high** are the commonly used settings.

–gc server option

Function

Set maximum desired interval between checkpoints.

Syntax

```
{ dbsrv8 | dbeng8 } -gc integer ...
```

Applies to

All operating systems and servers.

Description

Set the *maximum* desired length of time in minutes that the database server runs without doing a checkpoint on each database.

The default value is 60 minutes.

When a database server is running with multiple databases, the checkpoint time specified by the first database started is used unless overridden by this option. If a value of 0 is entered, the default value of 60 minutes is used.

Checkpoints generally occur more frequently than the specified time.

☞ For more information, see "How the database server decides when to checkpoint" on page 325.

See also

"CHECKPOINT_TIME option" on page 557

"Checkpoints and the checkpoint log" on page 322

–gd server option

Function

Set permissions required to start or stop a database.

Syntax

```
{ dbsrv8 | dbeng8 } -gd { DBA | all | none } ...
```

Applies to

All operating systems and servers.

Description

This is the permission required for a user to cause a new database file to be loaded by the server, or to stop a database on a running database server. The level can be one of the following:

- ◆ **DBA** Only users with DBA authority can start or stop databases.
- ◆ **all** All users can start or stop databases.
- ◆ **none** Starting and stopping databases is not allowed apart from when the database server itself is started and stopped.

The default setting is **ALL** for the personal database server and **DBA** for the network database server. Both uppercase and lowercase syntax is acceptable.

Example

The following set of steps illustrates how to use the –gd option for the network database server.

- 1 Enter a password in the *util_db.ini* file in your SQL Anywhere *win32* directory:

```
[UTILITY_DB]
pwd=mypwd
```

- 2 Start the network database server:

```
dbsrv8 -x tcpip -n myserver -gd DBA
```

- 3 Connect to the utility database from Interactive SQL.

The following command assumes that **myserver** is the default database server:

```
dbisql -c "uid=DBA;pwd=mypwd;dbn=utility_db "
```

- 4 Start a database:

```
start database asademo
on myserver;
```

5 Connect to the database you have started:

```
connect
to myserver
database asademo
user DBA identified by SQL
```

–ge server option

Function	Set stack size for external functions.
Syntax	{ dbsrv8 dbeng8 } -ge <i>integer</i> ...
Applies to	Windows 95/98/Me, Windows NT/2000/XP, NetWare
Description	Sets the stack size for threads running external functions, in bytes. The default is 32K.

–gf server option

Function	Disable firing of triggers by the server.
Syntax	{ dbsrv8 dbeng8 } -gf ...
Applies to	All operating systems and servers.
Description	The –gf server option instructs the server to disable the firing of triggers.
See also	"FIRE_TRIGGERS option" on page 568

–gk server option

Function	Set the permission required to stop the network server and personal server using <i>dbstop</i> .
Syntax	{ dbsrv8 dbeng8 } -gk { DBA all none } ...
Applies to	All operating systems and servers.
Description	<p>The allowed values include:</p> <ul style="list-style-type: none">◆ DBA Only users with DBA authority can use <i>dbstop</i> to stop the server. This is the default for the network server.◆ all All users can use <i>dbstop</i> to stop the server. This is the default for the personal server.◆ none The server cannot be stopped using <i>dbstop</i>. <p>Both uppercase and lower case syntax is acceptable.</p>

-gl server option

Function	Set the permission required to load data using LOAD TABLE, and to unload data using UNLOAD or UNLOAD TABLE.
Syntax	{ dbsrv8 dbeng8 } -gl { DBA all none } ...
Applies to	All operating systems and servers.
Description	<p>Using the UNLOAD TABLE or UNLOAD statements places data in files on the database server machine, and the LOAD TABLE statement reads files from the database server machine.</p> <p>To control access to the file system using these statements, the -gl server option allows you to control the level of database permission that is required to use these statements.</p> <p>The allowed values are as follows:</p> <ul style="list-style-type: none">◆ DBA Only users with DBA authority can load or unload data from the database.◆ all All users can load or unload data from the database.◆ none Data cannot be unloaded or loaded. <p>Both uppercase and lower case syntax is acceptable.</p> <p>The default settings are all for personal database servers on non-UNIX operating systems, and DBA for the network database server and the UNIX personal server. These settings reflect the fact that, on non-UNIX platforms, the personal database server is running on the current machine, and so the user already has access to the file system.</p>

-gm server option

Function	Limit the number of concurrent connections to the server.
Syntax	{ dbsrv8 dbeng8 } -gm <i>integer</i> ...
Applies to	All operating systems and servers.
Description	If this number is greater than the number that is allowed under licensing and memory constraints, it has no effect.

-gn server option

Function	Set the number of requests the database server can handle at one time.
Syntax	{ dbsrv8 dbeng8 } -gn <i>integer</i> ...
Applies to	All operating systems and servers.

Description

Set the number of execution threads to be used in the database server. If all execution threads are busy and the database server receives a new request, the new request must wait until another request is completed.

Each connection uses a thread for each request, and when the request is completed the thread is returned to the pool for use by other connections. As no connection can have more than one request in progress at one time, no connection uses more than one thread at a time. An exception to this rule is if a Java application uses threads. Each thread in the Java application is a database server execution thread.

The default is 20 threads for the network database server, and 10 threads for the personal database server. The number of threads cannot be greater than the number of server connections, so the maximum number of threads for the personal database server is also 10.

You may want adjust `-gn` based on Performance Monitor readings for **Requests: Active** and **Requests: Unscheduled**. If the number of active requests is always less than `-gn`, you can lower `-gn`. If the number of total requests (active + unscheduled) is often larger than `-gn`, then you might want to increase `-gn`.

See also

"Controlling threading from the command line" on page 14

`-gp` server option

Function

Set the maximum allowed database page size.

Syntax

`{ dbsrv8 | dbeng8 } -gp { 1024 | 2048 | 4096 | 8192 | 16384 | 32768 } ...`

Applies to

All operating systems and servers.

Description


Database files with a page size larger than the page size of the server cannot be loaded. This option explicitly sets the page size of the server, in bytes.

If you do not use this option, then the page size of the first database on the command line is used.

The minimum page size on all UNIX platforms is 2048 bytes. You can still use databases with smaller page sizes, but cache memory is used very inefficiently. If you do not use this option and start a server with no databases loaded, the default value is 2048.

On all other platforms, if you do not use this option and start a server with no databases loaded, the default value is 1024.

-gr server option

Function	Set the maximum length of time (in minutes) for recovery from system failure.
Syntax	{ dbsrv8 dbeng8 } -gr <i>integer</i> ...
Applies to	All operating systems and servers.
Description	<p>When a database server is running with multiple databases, the recovery time that is specified by the first database started is used unless overridden by this option.</p> <p> For more information, see "RECOVERY_TIME option" on page 595.</p>

-gss server option

Function	Set the stack size per internal execution thread in the server.
Syntax	{ dbsrv8 dbeng8 } -gss { <i>integer</i> <i>integer</i> K <i>integer</i> M } ...
Applies to	This option has no effect on Windows operating systems.
Description	<p>The number of internal execution threads is controlled by the -gn option, and has a default value of 20. The default stack size per internal execution thread is 64 kb for NetWare and UNIX, and the maximum stack size is 1 M. The -gss option allows you to lower the memory usage of the database server in environments with limited memory.</p>
See also	"Controlling threading from the command line" on page 14

-gt server option

Function	Sets the number of processors used by the database server.
Syntax	{ dbsrv8 dbeng8 } -gt <i>integer</i> ...
Applies to	All operating systems and servers except NetWare.
Description	<p>By default, on Windows operating systems, the network database server uses all CPUs available on the machine, the personal database server is limited to two processors, and the runtime database server uses a single processor. On UNIX, all database servers use all available processors by default.</p>
See also	"Controlling threading from the command line" on page 14

-gu server option

Function	Set permission levels for utility commands.
-----------------	---

Syntax { **dbsrv8** | **dbeng8** } **-gu** { **all** | **none** | **DBA** | **utility_db** } ...

Applies to All operating systems and servers.

Description Sets permission levels for utility commands such as CREATE DATABASE and DROP DATABASE. The level can be set to one of the following: **utility_db**, **all**, **none**, **DBA**. The default is **DBA**.

The **utility_db** level restricts the use of these commands to only those users who can connect to the utility database. The **all**, **none**, and **DBA** levels permit all users, no users, or users with DBA authority to execute utility commands.

-gx server option

Function Set the number of operating system threads assigned to the database server process.

Syntax { **dbsrv8** | **dbeng8** } **-gx** *integer* ...

Applies to Windows 95/98/Me, Windows NT/2000/XP

Description By default, this is set to one more than the number of CPUs on the machine. The additional process permits the use of remote data access between databases on a single database server.

You may want to increase the option setting beyond the default to reserve threads for external tasks, separate from standard database tasks. For example, you may want to add an extra thread for each concurrent connection using remote data access, or for connections using Java in the database to listen on an external port. In other circumstances, this option should be left at the default value.

On UNIX, each task is executed in its own thread, so that the number of tasks (**-gn**) also determines the number of threads.

See also "Controlling threading from the command line" on page 14

-m server option

Function Delete the transaction log when a checkpoint is done.

Syntax { **dbsrv8** | **dbeng8** } **-m** ...

Applies to All operating systems and servers.

Description This option deletes the transaction log when a checkpoint is done, either at shutdown or as a result of a checkpoint scheduled by the server.

Caution

When this option is selected, there is no protection against media failure on the device that contains the database files.

This provides a way to automatically limit the growth of the transaction log. Checkpoint frequency is still controlled by the CHECKPOINT_TIME and RECOVERY_TIME options (which you can also set on the command line).

This option is useful where high volume transactions that require fast response times are being processed, and the contents of the transaction log are not being relied upon for recovery or replication.

To avoid database file fragmentation, it is recommended that where this option is used, the transaction log be placed on a separate device or partition from the database itself.

Replicated databases

Do not use the -m option with databases that are being replicated. Replication inherently relies on transaction log information.

-n server option**Function**

Set the name of the database server.

Syntax

{ **dbsrv8** | **dbeng8** } -n *database file-name...*

Applies to

All operating systems and servers.

Description

By default, the database server receives the name of the database file with the path and extension removed. For example, if the server is started on the file *c:\Program Files\Sybase\SQL Anywhere 8\asademo.db* and no -n option is specified, the name of the server is **asademo**.

The server name is interpreted according to the character set of the machine, as no database collation exists at startup time. Multi-byte characters are not recommended in server names.

Names must be a valid identifier. Long engine names are truncated to different lengths depending on the protocol.

Protocol	Truncation Length
Unix shared memory	31 bytes
non-Unix shared memory	40 bytes
TCP/IP	40 bytes
SPX	32 bytes
Named Pipes	8 bytes


The server name specifies the name to be used in the **EngineName (ENG)** connection parameter of client application connection strings or profiles. With shared memory, there is a default database server that will be used if no server name is specified provided that at least one database server is running on the computer.

There are two -n options

The `-n` option is positional. If it appears before a database file name, it is a server option and names the server. If it appears after a database file name, it is a database option and names the database.

For example, the following command names the server `SERV` and the database `DATA`:

```
dbsrv8 -n SERV asademo.db -n DATA
```

 For more information, see "`-n` database option" on page 160.

See also "Identifiers" on page 7 of the book *ASA SQL Reference Manual*
"EngineName connection parameter" on page 176

-o server option

Function Print all server window output to a file.

Syntax { `dbsrv8` | `dbeng8` } `-o filename ...`

Applies to All operating systems and servers.

Description Print all server message window output to a file.

-os server option

Function Limit the file size for the server window output.

Syntax	{ dbsrv8 dbeng8 } -os <i>size</i> ...
Applies to	All operating systems and servers.
Description	Limit the size of the log file used by the -o option. The default is no limit, and the value is in bytes.
See also	" -o server option" on page 146

-p server option

Function	Set the maximum size of communication packets.
Syntax	{ dbsrv8 dbeng8 } -p <i>integer</i> ...
Applies to	All operating systems and servers.
Description	The default is 1460 bytes. The minimum value is 300 bytes and the maximum is 16000.
See also	"CommBufferSize connection parameter" on page 168

-pc server option

Function	Specifying this option compresses all connections except for same-machine connections. This option can be overridden for a particular client by specifying COMPRESS=NO in the client's connection parameters.
Syntax	{ dbsrv8 } -pc ...
Applies to	All operating systems and network servers.
Description	<p>The packets sent between an Adaptive Server Anywhere client and server can be compressed using the -pc option. Compressing a connection may improve performance under some circumstances. Large data transfers with highly compressible data tend to get the best compression rates.</p> <p>The default is not to compress connections. Specifying the -pc option compresses all connections except same-machine connections and TDS connections. TDS connections (including jConnect) do not support Adaptive Server Anywhere communication compression.</p> <p>Same-machine connections over any communication link will not enable compression, even if the -pc option or COMPRESS=YES connection parameter is used.</p>
See also	"Adjusting communication compression settings to improve performance" on page 98 "Compress connection parameter" on page 170

"Try using Adaptive Server Anywhere's compression features" on page 151
of the book *ASA SQL User's Guide*

–pt server option

Function	To increase or decrease the size limit at which packets are compressed.
Syntax	<code>{ dbsrv8 } -pt size_in_bytes ...</code>
Applies to	All operating systems and network servers.
Description	<p>This parameter takes an integer value representing the minimum byte-size of packets to be compressed. Values less than 80 are not recommended. The default is 120 bytes.</p> <p>Under some circumstances, changing the compression threshold can help performance of a compressed connection by allowing you to only compress packets when compression will increase the speed at which the packets are transferred. The default setting should be appropriate for most cases.</p> <p>If both client and server specify different compression threshold settings, the client setting applies.</p>
See also	<p>"Adjusting communication compression settings to improve performance" on page 98</p> <p>"CompressionThreshold connection parameter" on page 172</p> <p>"Try using Adaptive Server Anywhere's compression features" on page 151 of the book <i>ASA SQL User's Guide</i></p>

–q server option

Function	Do not display the server screen or its output.
Syntax	<code>{ dbsrv8 dbeng8 } -q ...</code>
Applies to	All operating systems and servers, except NetWare.

–qp server option

Function	<p>Do not display messages about performance in the database server window. Messages that are suppressed include the following:</p> <ul style="list-style-type: none">◆ No unique index or primary key for table 'table_name'◆ Database file "mydatabase.db" consists of nnn fragments
Syntax	<code>{ dbsrv8 dbeng8 } -qp ...</code>
Applies to	All operating systems and servers.

-r server option

Function	All databases started on the server are read-only. No changes to the database(s) are allowed: the server does not modify the database file(s) and transaction log files. This option is position dependent.
Syntax	<code>{ dbsrv8 dbeng8 } -r ...</code>
Applies to	All operating systems and servers.
Description	<p>Opens all database files as read-only with the exception of the temporary file when the option is specified before any database names on the command line. If the <code>-r</code> server option is specified after a database name, only that specific database is read-only. You can make changes on temporary tables, but ROLLBACK has no effect, since the transaction and rollback logs are disabled.</p> <p>Databases distributed on CD-ROM devices, and compressed databases are examples of database files that cannot be modified. You can either create a write file to allow changes to the database outside the database file, or run in read-only mode.</p> <p>If you attempt to modify the database, for example with an INSERT or DELETE statement, a <code>SQLSTATE_READ_ONLY_DATABASE</code> error is returned.</p> <p>Databases that require recovery cannot be started in read-only mode. For example, database files created using an online backup cannot be started in read-only mode if there were any open transactions when the backup was started, since these transactions would require recovery when the backup copy is started.</p>

Example

To open two databases in read-only mode

```
dbeng8 -r database1.db database2.db
```

To open only the first of two databases in read-only mode.

```
dbeng8 database1.db -r database2.db
```

-s server option

Function	Set the user ID for syslog messages.
Syntax	<code>{ dbsrv8 dbeng8 } -s { none user <i>login-id</i> } ...</code>
Applies to	UNIX
Description	Sets the system user ID used in messages to the syslog facility. The default is user , which uses the user ID for the database server process. A value of none prevents any syslog messages from being logged.

–sb server option

Function	Specify how the server reacts to broadcasts.
Syntax	{ dbsrv8 dbeng8 } -sb { 0 1 } ...
Applies to	SPX, TCP/IP
Description	<p>Using -sb 0 causes the server not to start up any TCP/UDP broadcast listeners. In addition to forcing clients to use the DoBroadcast=NONE and HOST= options to connect to the server, this option causes the server to be unlisted when using dblocate.</p> <p>Using -sb 1 causes the server to not respond to broadcasts from dblocate, while leaving connection logic unaffected. You can connect to the server by specifying LINKS=tcip and ENG=<name>.</p>

–sc server option

Function	Disable the shared memory communications protocol and use Named Pipes.
Syntax	{ dbsrv8 dbeng8 } -sc ...
Applies to	Windows NT/2000/XP
Description	<p>This option disables the shared memory communications protocol that is used for same-machine communications, and starts the Named Pipes protocol.</p> <p>This option is implemented as part of an initiative to obtain C2 security certification. It is likely to be of general use only for customers wanting to run in a C2-certified environment.</p>

–ti server option

Function	Disconnect inactive connections.
Syntax	{ dbsrv8 dbeng8 } -ti <i>minutes</i> ...
Applies to	All operating systems and servers.
Description	<p>Disconnect connections that have not submitted a request for <i>minutes</i>. The default is 240 (4 hours). A client machine in the middle of a database transaction holds locks until the transaction is ended or the connection is terminated. The -ti option is provided to disconnect inactive connections, freeing their locks.</p> <p>The -ti option does not disconnect clients that use the shared memory communications link.</p>

Setting the value to zero disables checking of inactive connections, so that no connections are disconnected.

See also "sa_server_option system procedure" on page 716 of the book *ASA SQL Reference Manual*

-tl server option

Function Set the period at which to send liveness packets.

Syntax { **dbsrv8** | **dbeng8** } **-tl** *seconds* ...

Applies to All database servers using TCP/IP or SPX.

See also "sa_server_option system procedure" on page 716 of the book *ASA SQL Reference Manual*

Description A liveness packet is sent periodically across a client/server TCP/IP or SPX communications protocol to confirm that a connection is intact. If the server runs for a liveness timeout period (default 2 minutes) without detecting a liveness packet, the communication is severed. The server drops any connections associated with that client. UNIX non-threaded clients and TDS connections do not do liveness checking.

The **-tl** option on the server sets the liveness timeout for all clients that do not specify a liveness period.

Liveness packets are sent at an interval of the (liveness timeout)/4.

You can disable liveness by specifying:

```
dbsrv8 -tl 0
```

-tmt server option

Function To set a reenlistment timeout for participation in distributed transactions.

Syntax { **dbsrv8** | **dbeng8** } **-tmt** *milliseconds* ...

Applies to Windows NT/2000/XP only.

Description Used during recovery of distributed transactions. The value specifies how long the database server should wait to be reenlisted. By default there is no timeout (the database server waits indefinitely).

-tmf server option

Function For recovery from distributed transactions in unusual circumstances.

Syntax	<code>{ dbsrv8 dbeng8 } -tmf ...</code>
Applies to	Windows NT/2000/XP only.
Description	Used during recovery of distributed transactions when the distributed transaction coordinator is not available. It could also be used if starting a database with distributed transactions in the transaction log, on a platform where the distributed transaction coordinator is not available.

Caution

If you use this option, distributed transactions are not recovered properly. It is not for routine use.

-tq time server option

Function	Shut down the server at a specified time.
Syntax	<code>{ dbsrv8 dbeng8 } -tq { datetime time } ...</code>
Applies to	All operating systems and servers.
Description	This is useful for setting up automatic off-line backup procedures (see "Backup and Data Recovery" on page 299). The format for the time is in <i>hh:mm</i> (24 hour clock), and can be preceded by an optional date. If a date is specified, the date and time must be enclosed in double quotes and be in the format <i>YYYY/MM/DD HH:MM</i> .
See also	"sa_server_option system procedure" on page 716 of the book <i>ASA SQL Reference Manual</i>

-u server option

Function	Open files using the operating system disk cache.
Syntax	<code>{ dbsrv8 dbeng8 } -u ...</code>
Applies to	Windows NT/2000/XP, Windows 95/98/Me, UNIX
Description	<p>Files are opened using the operating system disk cache in addition to the database cache.</p> <p>While the operating system disk cache may improve performance in some cases, in general better performance is obtained without this option, using the database cache only.</p>

If the server is running on a dedicated machine, you should not use the `-u` option, as the database cache itself is generally more efficient. You may want to use the `-u` option if the server is running on a machine with several other applications (so that a large database cache may interfere with other applications) and yet IO-intensive tasks are run intermittently on the server (so that a large cache will improve performance).

-ud server option

Function	Run as a daemon.
Syntax	<code>{ dbsrv8 dbeng8 } -ud ...</code>
Applies to	UNIX
Description	Using this option lets you run the server so that it continues running after the current operating system session ends.

-ut server option

Function	Touch temporary files.
Syntax	<code>{ dbsrv8 dbeng8 } -ut <i>minutes</i> ...</code>
Applies to	UNIX
Description	This option causes the server to touch temporary files at specified intervals.

-v server option

Function	Display the software version.
Syntax	<code>{ dbsrv8 dbeng8 } -v ...</code>
Applies to	All operating systems and servers.
Description	Supplies the database server version in a message box, and then stops.

-x server option

Function	To specify server side network communications protocols.
Syntax 1	<code>{ dbsrv8 } -x { all none { [namedpipes spx tcPIP] <i>parmlist</i>,... } } ...</code>

Syntax 2

```
parmlist:
  { parm=value;...}

{ dbeng8 } -x {
  all
  | none
  | { [namedpipes | tcpip ] parmlist,... }
} ...

parmlist:
  { parm=value;...}
```

Applies to

All operating systems and servers.

Description

Use the `-x` option to specify which communications protocols, in addition to shared memory, you want to use to listen for client connection broadcasts.

If you do not specify the `-x` option, the server attempts to listen for client connection broadcasts using all protocols supported by the database server running on your operating system, including shared memory.

If you specify the `-x` option with one or more protocols, the server attempts to listen for client connection broadcasts using the specified protocol(s) and also using a shared-memory protocol.

If you are running Windows CE and specify the `-x` option, the server only attempts to listen for client connection broadcasts using the TCP/IP protocol unless you explicitly request otherwise.

Regardless of which settings you choose for the `-x` option, the server always listens for connection broadcasts using the shared memory protocol. In addition to the shared memory protocol, you can also specify the following:

- ◆ **ALL** Listen for connection attempts by the client using all communications protocols that are supported by the server on this platform, including shared memory. This is the default.
- ◆ **NamedPipes (NP)** Listen for connection attempts by the client using the NamedPipes protocol. NamedPipes is supported on Windows NT/2000/XP as an alternative means of same-machine communication.
- ◆ **NONE** Listen for connection attempts by the client using only the shared memory protocol.
- ◆ **SPX** Listen for connection attempts by the client using the SPX protocol. The SPX protocol is supported by NetWare, Windows NT/2000/XP, and Windows 95/98/Me network servers.

- ◆ **TCPIP (TCP)** Attempt to connect to the client using the TCP/IP protocol. The TCP/IP protocol is supported by the network server on all operating systems, and by the personal database server for same-machine communications.

By default, the database server listens for broadcasts on port 2638, and redirects them to the appropriate port. This ensures a connection in most cases.

You can override this default and cause the server not to listen on port 2638 by setting the option `-sb 0`, or by turning off the `BroadcastListener` option (`BroadcastListener=0`). Additionally, if the client and server are communicating through a firewall, the client must send the packet to the exact port the server is listening on by specifying `DoBroadcast=None` and `Host=`.

☞ For information, see "ServerPort communication parameter" on page 198.

For some protocols, additional parameters may be provided, in the format

```
-x tcpip(PARM1=value1;PARM2=value2;...)
```

☞ For a description of available parameters, see "Network communications parameters" on page 189.

For UNIX, quotation marks are required if more than one parameter is supplied:

```
-x "tcpip(PARM1=value1;PARM2=value2;...)"
```

Examples

Allow only shared memory, TCP/IP and SPX communications:

```
-x tcpip,spx
```

See also

"CommLinks connection parameter" on page 169

-y server option

Function

Run as a Windows service.

Syntax

```
{ dbsrv8 | dbeng8 } -y ...
```

Applies to

Windows 95/98/Me

Description

If the server registered as a Windows service, it continues to operate whether users log on or off, and shutdown commands are ignored.

–z server option

Function	Display communications operations on startup for troubleshooting purposes.
Syntax	{ dbsrv8 dbeng8 } -z ...
Applies to	All operating systems and servers.
Description	This should only be used when tracking problems. The information appears in the database server window.

–zl server option

Function	Turn on capturing of the most recently-prepared SQL statement for each connection to databases on the server.
Syntax	{ dbsrv8 dbeng8 } -zl ...
Applies to	All operating systems and servers.
Description	<p>This feature can also be turned on using the <i>remember_last_statement</i> server setting. You can obtain the most recently-prepared SQL statement for a connection using the LastStatement property function. The <i>sa_conn_activity</i> stored procedure allows you to obtain the most recently-prepared SQL statement for all current connections to databases on the server.</p> <p>For stored procedure calls, only the outermost procedure call appears, not the statements within the procedure.</p>
See also	"Connection-level properties" on page 618 "sa_conn_activity system procedure" on page 686 of the book <i>ASA SQL Reference Manual</i> "sa_server_option system procedure" on page 716 of the book <i>ASA SQL Reference Manual</i>

–zo server option

Function	Redirect request-level logging information to a file separate from the regular log file.
Syntax	{ dbsrv8 dbeng8 } -z ...
Applies to	All operating systems and servers.
Description	<p>Request-level logging is turned on using the -zx option. You can direct the output from this file to a separate file from that specified on a -o option.</p> <p>This option also prevents request-level logging from appearing in the console.</p>

See also "-zr server option" on page 157
 "-zs server option" on page 157

-zr server option

Function Enable request-level logging of operations.

Syntax { **dbsrv8** | **dbeng8** } **-zr** { **all** | **SQL** | **none** } ...

Applies to All operating systems and servers.

Description This should only be used when tracking problems. The information appears in the database server window or is sent to the logging file.

See also "sa_server_option system procedure" on page 716 of the book *ASA SQL Reference Manual*

-zs server option

Function Limit the size of the request-level logging file.

Syntax { **dbsrv8** | **dbeng8** } **-zs** *size* ...

Applies to All operating systems and servers.

Description Request-level logging is turned on using the **-zr** option, and redirected to a separate file using the **-zo** option. You can limit the size of the file using the **-zs** option.

 By default there is no limit. The value is in kilobytes.

See also "-zo server option" on page 156
 "-zr server option" on page 157

Recovery options

These options are for use in recovery situations only.

-a recovery option

Function Apply the named transaction log.

Syntax { **dbsrv8** | **dbeng8** } **-a** *log-filename* ...

Applies to All operating systems and servers.

Description

This is used to recover from media failure on the database file. When this option is specified, the database server applies the log and then terminates—it will not continue to run.

During recovery, you should use the same server that you use in production. For example, if you use the network database server (*dbsrv8.exe*) during production, then you should use the network server for recovery because you may encounter problems as a result of the personal server's 10 connection limit. Specifying a cache size when starting the server can reduce recovery time.

🔗 For information on recovery, see "Backup and Data Recovery" on page 299.

Example

The following example, entered all on one line, applies the log file *asademo.log* to the sample database.

```
dbeng8 "C:\Program Files\Sybase\SQL Anywhere  
8\asademo.db" -a "C:\backup\asademo.log"
```

-f recovery option

Function

Force the database server to start after the transaction log has been lost.

Syntax

```
{ dbsrv8 | dbeng8 } -f ...
```

Applies to

All operating systems and servers.

Description

If there is no transaction log, the database server carries out a checkpoint recovery of the database and then terminates—it does not continue to run. You can then restart the database server without the `-f` option for normal operation.

If there is a transaction log in the same directory as the database, the database server carries out a checkpoint recovery, and a recovery using the transaction log, and then terminates—it does not continue to run. You can then restart the database server without the `-f` option for normal operation.

During recovery, you should use the same server that you use in production. For example, if you use the network database server (*dbsrv8.exe*) during production, then you should use the network server for recovery because you may encounter problems as a result of the personal server's 10 connection limit. Specifying a cache size when starting the server can reduce recovery time.

🔗 For more information see "Backup and Data Recovery" on page 299.

Example

```
dbeng8 "c:\Program Files\Sybase\SQL Anywhere  
8\asademo.db" -f
```


Database options

These options are entered after the database name, and apply only to that database.

-ek database option

Function	To specify the key for a strongly encrypted database.
Syntax	{ dbsrv8 dbeng8 } [<i>server-options</i>] <i>database-file</i> -ek <i>key</i> ...
Description	The -ek database option must be provided after a database filename on the command line. You must provide the KEY value with the -ek option to start an encrypted database. The KEY is a string, including mixed cases, numbers, letters, and special characters.
Example	<pre>dbsrv8 -x tcpip asademo.db -ek "Akmm9u70y"</pre>
See also	"Encryption connection parameter" on page 177 "Encryption Key connection parameter" on page 179 "Encrypting client/server communications" on page 398 "-ep server option" on page 137

-m database option

Function	Truncate the transaction log when a checkpoint is done.
Syntax	{ dbsrv8 dbeng8 } [<i>server-options</i>] <i>database-file</i> -m ...
Applies to	All operating systems and servers.
Description	<p>Truncate (delete) the transaction log when a checkpoint is done, either at shutdown or as a result of a checkpoint scheduled by the server. This provides a way to automatically limit the growth of the transaction log. Checkpoint frequency is still controlled by the CHECKPOINT_TIME and RECOVERY_TIME options (which you can also define on the command line).</p> <p>The -m option is useful where high volume transactions requiring fast response times are being processed, and the contents of the transaction log are not being relied upon for recovery or replication. When this option is selected, there is no protection against media failure on the device that contains the database files.</p> <p>To avoid database file fragmentation, it is recommended that where this option is used, the transaction log be placed on a separate device or partition from the database itself.</p>

This option is the same as the `-m` server option, but applies only to the current database or the database identified by the *database-file* variable.

Replicated databases

Do not use the `-m` option with databases that are being replicated.
Replication inherently relies on transaction log information.

Example

The following example starts a database server named `silver` and loads the sample database. When a checkpoint is done, the transaction log is truncated.

```
dbsrv8 -n silver "c:\Program Files\Sybase\SQL Anywhere
8\asademo.db" -m
```

`-n` database option

Function

Set the name of the database.

Syntax

```
{ dbsrv8 | dbeng8 } [ server-options ] database-file -n string ...
```

Applies to

All operating systems and servers.

Description

Both database servers and databases can be named. Since a database server can load several databases, the database name is used to distinguish the different databases.

By default, the database receives the name of the database file with the path and extension removed. For example, if the database is started on *c:\asalasademo.db* and no `-n` option is specified, the name of the database is **asademo**.

Example

The following example starts the database server with a cache size of 3Mb, loads the sample database, and names the sample database "test".


```
dbsrv8 -c 3Mb "c:\Program Files\Sybase\SQL Anywhere
8\asademo.db" -n "test"
```

There are two `-n` options

The `-n` option is positional. If it appears before a database file name, it is a server option and names the server. If it appears after a database file name, it is a database option and names the database.

For example, the following command names the server `SERV` and the database `DATA`:

```
dbsrv8 -n SERV asademo.db -n DATA
```

 For more information, see "`-n` server option" on page 145.

-r database option

Function	Starts the named database as read-only. No changes to the database(s) are allowed; the server does not modify the database file(s) and transaction log files. This option is position dependent.
Syntax	<code>{ dbsrv8 dbeng8 } -r ...</code>
Applies to	All operating systems and servers.
Description	<p>Opens all database files as read-only with the exception of the temporary file when the option is specified before any database names on the command line. If the <code>-r</code> server option is specified after a database name, only that specific database is read-only. You can make changes on temporary tables, but ROLLBACK has no effect, since the transaction and rollback logs are disabled.</p> <p>Databases distributed on CD-ROM devices, and compressed databases are examples of database files that cannot be modified. You can either create a write file to allow changes to the database outside the database file, or run in read-only mode.</p> <p>If you attempt to modify the database, for example with an INSERT or DELETE statement, a <code>SQLSTATE_READ_ONLY_DATABASE</code> error is returned.</p> <p>Databases that require recovery cannot be started in read-only mode. For example, database files created using an online backup cannot be started in read-only mode if there were any open transactions when the backup was started, since these transactions would require recovery when the backup copy is started.</p>
Example	<p>To open two databases in read-only mode</p> <pre>dbeng8 -r database1.db database2.db</pre> <p>To open only the first of two databases in read-only mode.</p> <pre>dbeng8 database1.db -r database2.db</pre>

CHAPTER 6

Connection and Communication Parameters




About this chapter This chapter provides a reference for the parameters that establish and describe connections from client applications to a database.

Contents

Topic	Page
Connection parameters	164
Network communications parameters	189

Connection parameters

This section describes each connection parameter. Connection parameters are included in connection strings. They can be entered in the following places:

- ◆ In an application's connection string.
 For more information, see "Assembling a list of connection parameters" on page 75.
- ◆ In an ODBC data source.
 For more information, see "Working with ODBC data sources" on page 53.
- ◆ In the Adaptive Server Anywhere Connect dialog.
 For more information, see "Connecting from Adaptive Server Anywhere utilities" on page 52.

The ODBC configuration dialog and the Adaptive Server Anywhere Connect dialog for Windows operating systems share a common format. Some of the parameters correspond to checkboxes and fields in these dialogs, while others can be entered in the text box at the end on the Advanced tab.

Notes

- ◆ Connection parameters are case insensitive.
- ◆ The Usage for each connection parameter describes the circumstances under which the parameter is to be used. Common usage entries include the following:
 - ◆ **Embedded databases** When Adaptive Server Anywhere is used as an embedded database, the connection starts a personal server and loads the database. When the application disconnects from the database, the database is unloaded and the server stops.
 - ◆ **Running local databases** This refers to the case where an Adaptive Server Anywhere personal server is already running, and the database is already loaded on the server.
 - ◆ **Network servers** When Adaptive Server Anywhere is used as a network server, the client application must locate a server already running somewhere on the network and connect to a database.
- ◆ You can use the dbping utility to test connection strings. For example, if the database file *MyDB.db* is not currently running, the following string returns the message Ping database successful:

```
dbping -d -c "dbf=c:\Databases\MyDB.db;astart=yes"
```

The following command, however, returns the message Ping database failed – Database server not running, because Adaptive Server Anywhere is unable to find a server for *MyDB.db*:

```
dbping -d -c "dbf=c:\Databases\MyDB.db;astart=no"
```

For more information, see "The Ping utility" on page 494.

AppInfo connection parameter [APP]

Function	To assist administrators in identifying the origin of particular client connections from a database server.
Usage	Anywhere
Values	<i>String</i>
Default	<i>Empty string</i>
Description	<p>This connection parameter is sent to the database server from Embedded SQL, ODBC, or OLE DB clients. It is not available from Open Client or jConnect applications such as Interactive SQL or Sybase Central.</p> <p>It consists of a generated string that holds information about the client process, such as the IP address of the client machine, the operating system it is running on, and so on. The string is associated in the database server with the connection, and you can retrieve it using the following statement:</p> <pre>select connection_property('appinfo')</pre> <p>Clients can also specify their own string, which is appended to the generated string. The AppInfo property string is a sequence of semi-colon-delimited key=value pairs. The valid keys are as follows:</p> <ul style="list-style-type: none"> ◆ API Either DBLIB or ODBC. ◆ APPINFO If you specified AppInfo in the connection string, the string entered. ◆ EXE The name of the client executable (Windows and NetWare only) ◆ HOST The host name of the client machine ◆ IP The IP address of the client machine (UNIX and NetWare only) ◆ OS The operating system name and version number (for example, Windows NT 4.0, NetWare 5.1) ◆ PID The process ID of the client (Windows and UNIX only) ◆ THREAD The thread ID of the client (Windows and UNIX only)

- ◆ **TIMEZONEADJUSTMENT** the number of minutes that must be added to the Coordinated Universal Time (UTC) to display time local to the connection.
- ◆ **VERSION** The version of the connection protocol in use, including major and minor values, and a build number (for example 8.0.00.3642)

If you specify a debug log file in your client connection parameters, the APPINFO string is added to the file.

Examples

- ◆ Connect to the default database from the C++ version of Interactive SQL:

```
dbisqlc -c uid=DBA;pwd=SQL
```

View the application information:

```
select connection_property('appinfo')
```

The result is as follows (in a single string):

```
HOST=machine-name;  
OS=Windows NT 4.0;  
PID=0x11b;  
THREAD=0x102;  
VERSION=8.0.00.3642
```

- ◆ Connect to the default database from the C++ version of Interactive SQL, appending your own information to the AppInfo property:

```
dbisqlc -c "uid=DBA;pwd=SQL;app=ISQL connection"
```

View the application information:

```
select connection_property('appinfo')
```

The result is as follows (in a single string):

```
HOST=machine-name;  
OS=Windows NT 4.0;  
PID=0x10e;  
THREAD=0xe1;  
VERSION=8.0.00.3642;  
APPINFO=ISQL connection
```

AutoStart connection parameter [ASTART]


Function	To prevent a local database server from being started if no connection is found.
Usage	Anywhere

Values	YES, NO
Default	YES
Description	By default, if no server is found during a connection attempt, and a database file is specified, then a database server is started on the same machine. You can turn this behavior off by setting the AutoStart (ASTART) connection parameter to NO or OFF in the connection string.

AutoStop connection parameter [ASTOP]

Function	To prevent a database from being stopped as soon as there are no more open connections.
Usage	Embedded databases
Values	YES, NO
Default	YES
Description	<p>By default, any server that is started from a connection string is stopped when there are no more connections to it. Also, any database that is loaded from a connection string is unloaded as soon as there are no more connections to it. This behavior is equivalent to <code>AutoStop=YES</code>.</p> <p>If you supply <code>AutoStop=NO</code>, any database that you start in that connection remains running when there are no more connections to it. As a consequence, the database server remains operational as well.</p> <p>The AutoStop (ASTOP) connection parameter is used only if you are connecting to a database that is not currently running. It is ignored if the database is already started.</p>

CharSet connection parameter [CS]

Function	To specify the character set to be used on this connection.
Usage	Anywhere
Values	<i>String</i>
Default	The local character set.
	 For information on how this is determined, see "Determining locale information" on page 288.
Description	If you supply a value for CharSet, the specified character set is used for the current connection.

☞ For a list of valid character set values, see "Character set labels" on page 265.

CommBufferSize connection parameter [CBSIZE]

Function	To set the maximum size of communication packets, in bytes.
Usage	Anywhere
Values	<i>Integer</i>
Default	If no CommBufferSize value is set, the CommBufferSize is controlled by the setting on the server, which defaults to 1460 bytes.
Description	The CommBufferSize (CBSIZE) connection parameter specifies the size of communications packets, in bytes. The minimum value of CommBufferSize (CBSIZE) is 300, and the maximum is 16000.

The protocol stack sets the maximum size of a packet on a network. If you set CommBufferSize to be larger than that permitted by your network, the largest buffers are broken up by the network software. You should set the buffer size to be somewhat smaller than that allowed by your network because the network software may add information to each buffer before sending it over the network.

A larger packet size improves performance for multi-row fetches and fetches of larger rows. As each connection has its own pool of buffers, a large buffer size increases the memory usage. The application side uses four default-sized buffers per connection, and on the server side, there are two.

If CommBufferSize is not specified on the client, the connection uses the engine's buffer size. If CommBufferSize is specified on the client, the connection uses the minimum of either the engine's buffer size or the CommBufferSize value.

This corresponds to the SQL Anywhere Version 5 *dbclient* -p option. Using the -p option to set the CommBufferSize causes all clients which do not specify their own CommBufferSize, as well as all clients which specify a CommBufferSize larger than the engine's CommBufferSize, to use the size specified in the -p option.

Example ♦ To set the buffer size to 400 bytes:

```
...
CommBufferSize=400
...
```

Alternatively, you can set this parameter by entering its value in the Buffer Size text box of the Network tab of the ODBC Configuration dialog.

CommLinks connection parameter [LINKS]

Function Usage	<p>To specify client side network communications protocols.</p> <p>Anywhere. The CommLinks (LINKS) connection parameter is optional for connections to a personal server, and required for connections to a network server.</p>
Values	<i>String</i>
Default	Use only the shared memory communications protocol to connect.
Description	<p>If you do not specify a CommLinks (LINKS) connection parameter, the client searches for a server on the current machine only, and only using a shared memory connection. This is the default behavior, and is equivalent to <code>CommLinks=ShMem</code>. The shared memory protocol is the fastest communication link between a client and server running on the same machine, as is typical for applications connecting to a personal database server.</p> <p>If you specify <code>CommLinks=ALL</code>, the client searches for a server using all available communication protocols. Since there may be an impact on performance if you specify <code>CommLinks=ALL</code>, use this setting only when you don't know which protocol to use.</p> <p>If you specify one or more protocols in the CommLinks (LINKS) connection parameter, the client uses the named communication protocol(s) to search for a network database server.</p> <p>CommLinks (LINKS) connection parameter values are case insensitive, and include:</p> <ul style="list-style-type: none">◆ SharedMemory (ShMem) Start the shared memory protocol for same-machine communication. This is the default setting. The client tries shared memory first if it appears in a list of protocols, regardless of the order in which protocols appear.◆ ALL Start all available communications protocols. Use this setting if you are unsure of which communication protocol(s) to use.◆ NamedPipes (NP) For C2 security purposes, connect from a Windows NT/2000/XP client to a database server on the same machine that was started with the <code>-sc</code> option.◆ TCP/IP (TCP) Start the TCP/IP communications protocol. TCP/IP is supported on all operating systems.◆ SPX Start the SPX communications protocol. The SPX protocol is supported for Windows and NetWare clients.

Each of these values can have additional network communications parameters supplied.

☞ For a list of parameters, see "Network communications parameters" on page 189.

You may wish to use a specific protocol, as opposed to ALL, for the following reasons:

- ◆ The network library starts slightly faster if the client uses only necessary network protocols.
- ◆ Connecting to the database may be faster.
- ◆ You must specify the protocol explicitly if you wish to tune the broadcast behavior of a particular protocol by providing additional network communications parameters.

The **CommLinks (LINKS)** connection parameter corresponds to the database server `-x` option.

Examples

- ◆ The following connection string fragment starts the TCP/IP protocol only:

```
CommLinks=tcpip
```

- ◆ The following connection string fragment starts the shared memory protocol and searches for the database server over shared memory. If the search fails, it then starts the TCP/IP port and searches for the server on the local network. If that fails, it starts the SPX port and searches for the engine over SPX.

```
CommLinks=tcpip,shmem,spx
```

- ◆ The following connection string fragment starts the SPX port and searches for the engine over SPX. If the search fails, it then starts the TCP port and searches for the server on the local network, as well as the host kangaroo. Note that if the server is found over SPX, the TCP port is NOT started.

```
CommLinks=spx,tcpip(HOST=kangaroo)
```

See also

"Network communications parameters" on page 189

"Client/Server Communications" on page 91

"-x server option" on page 153

Compress connection parameter [COMP]

Function

To turn compression ON or OFF for a connection. Compressing a connection may improve performance under some circumstances.

Usage	Anywhere except with TDS connections. TDS connections (including jConnect) do not support Adaptive Server Anywhere communication compression.
Values	<p>YES, NO</p> <p>In the case of a difference between client and server settings, the client setting applies.</p>
Default	<p>NO</p> <p>If no Compress value is set, the compression status is controlled by the setting on the server, which defaults to no compression.</p>
Description	<p>The packets sent between an Adaptive Server Anywhere client and server can be compressed using the Compress (COMP) connection parameter. Large data transfers with highly compressible data tend to get the best compression rates.</p> <p>Available values of the Compress (COMP) connection parameter are case insensitive, and include:</p> <ul style="list-style-type: none"> ◆ YES Turn communication compression on for this connection. ◆ NO Turn communication compression off for this connection. <p>To save yourself time and possible disappointment, it is wise to conduct a performance analysis on the particular network and using the particular application before using communication compression in a production environment.</p> <p>To enable compression for all remote connections on the server, use the <code>-pc</code> server option.</p> <p>Note that same-machine connections over any communication link will not enable compression, even if the <code>-pc</code> option or <code>COMPRESS=YES</code> parameter is used.</p>
Examples	<ul style="list-style-type: none"> ◆ The following connection string fragment turns packet compression ON: <code>Compress=YES</code> ◆ The following connection string fragment turns packet compression OFF: <code>Compress=NO</code>
See also	<p>"<code>-pc</code> server option" on page 147</p> <p>"Adjusting communication compression settings to improve performance" on page 98</p>

CompressionThreshold connection parameter [COMPTH]

Function	<p>To increase or decrease the size limit at which packets are compressed. Changing the compression threshold can help performance of a compressed connection by allowing you to only compress packets when compression will increase the speed at which the packets are transferred.</p>
Usage	<p>Anywhere except TDS. Only applies to compressed connections.</p>
Values	<p><i>Integer</i>, representing the minimum byte-size of packets to be compressed. Values less than 80 are not recommended.</p> <p>If both client and server specify different compression threshold settings, the client setting applies.</p>
Default	<p>If no CompressionThreshold value is set, the compression threshold value is controlled by the setting on the server, which defaults to 120 bytes.</p>
120	
Description	<p>When compression is enabled, individual packets may or may not be compressed, depending on their size. For example, Adaptive Server Anywhere does not compress packets smaller than the compression threshold, even if communication compression is enabled. As well, small packets (less than about 100 bytes) usually do not compress at all. Since CPU time is required to compress packets, attempting to compress small packets could actually decrease performance.</p> <p>Generally speaking, lowering the compression threshold value may improve performance on very slow networks, while raising the compression threshold may improve performance by reducing CPU. However, since lowering the compression threshold value will increase CPU usage on both the client and server, a performance analysis should be done to determine whether or not changing the compression threshold is beneficial.</p>
Example	<p>◆ Connect, with a compression threshold of 100 bytes.</p> <pre>CompressionThreshold=100</pre>
See also	<p>"-pt server option" on page 148</p> <p>"Adjusting communication compression settings to improve performance" on page 98</p>

ConnectionName connection parameter [CON]

Function	<p>Names a connection, to make switching to it easier in multi-connection applications.</p>
Usage	<p>Not available for ODBC.</p>
Values	<p><i>String</i></p>

Default	No connection name.
Description	<p>An optional parameter, providing a name for the particular connection you are making. You may leave this unspecified unless you are going to establish more than one connection, and switch between them.</p> <p>The connection name is not the same as the data source name.</p>
Example	<p>◆ Connect, naming the connection First_Con:</p> <pre>CON=First_Con</pre>

DatabaseFile connection parameter [DBF]

Function	The database file to which you want to connect.
Usage	Embedded databases
Values	<i>String</i>
Default	There is no default setting.
Description	<p>To load and connect to a specific database file.</p> <ul style="list-style-type: none"> ◆ If a database is loaded with a name that is the same as the DatabaseFile (DBF) connection parameter, but without the <i>.db</i> extension, the connection is made to that database instead. ◆ If the filename does not include an extension, a file of name <i>.db</i> is looked for. ◆ The path of the file is relative to the working directory of the database server. If you start the server from a command prompt, the working directory is the directory that you are in when entering the command. If you start the server from an icon or shortcut, it is the working directory that the icon or shortcut specifies. <p>You can also use UNC filenames and Novell NetWare Directory Services file names.</p>

🔗 For information, see "The database server" on page 120.

Caution

The database file must be on the same machine as the database server. Managing a database file that is located on a network drive can lead to file corruption.

Example

- ◆ To load and connect to the sample database, *asademo.db*, installed in the directory *c:\Program Files\Sybase\SQL Anywhere 8*, use the following **DatabaseFile (DBF)** connection parameter:

```
DBF=c:\Program Files\Sybase\SQL  
Anywhere 8\asademo.db
```

DatabaseName connection parameter [DBN]

Function

Identifies a loaded database to which a connection needs to be made.

Usage

Running local databases or network servers.

Values

String

Default

There is no default setting.

Description

Whenever a database is started on a server, it is assigned a database name. The default database name is the name of the database file with the extension and path removed.

Example

- ◆ Connect to a running database named Kitchener:

```
DBN=Kitchener
```

DatabaseSwitches connection parameter [DBS]

Function

To provide database-specific options (switches) when starting a database.

Usage

Connecting to a server when the database is not loaded. This connection parameter autostarts an engine with the specified database and options if a server is not already running.

Values

String

Default

No options

Description

You should supply DatabaseSwitches only if you are connecting to a database that is not currently running. When the server starts the database specified by DatabaseFile, the server uses the supplied DatabaseSwitches to determine startup options for the database.

Only *database* options (switches)

can be supplied using this parameter. Server switches must be supplied using the StartLine connection parameter.

☞ For information about database options, see "Database options" on page 159.

Example

- ◆ The following command, entered all on one line at a command prompt, connects to the default database server, loads the database file *asademo.db* (**DatabaseFile (DBF)** connection parameter), names it as *my_db* (DBS parameter) and connects to the database of that name (**DatabaseName (DBN)** connection parameter).

```
dbisql -c "uid=DBA;pwd=SQL;dbf=c:\Program
Files\Sybase\SQL Anywhere
8\asademo.db;dbn=my_db;dbs=-n my_db"
```

See also

"The database server" on page 120
 "StartLine connection parameter" on page 187

DataSourceName connection parameter [DSN]

Function

Tells the ODBC driver manager or Embedded SQL library where to look in the *odbc.ini* file or registry to find ODBC data source information.

Usage

Anywhere

Values

String

Default

There is no default data source name.

Description

It is common practice for ODBC applications to send only a data source name to ODBC. The ODBC driver manager and ODBC driver locate the data source, which contains the remainder of the connection parameters.

In Adaptive Server Anywhere, Embedded SQL applications can also use ODBC data sources to store connection parameters.

Example

- ◆ The following parameter uses a data source name:

```
DSN=Dynamo Demo
```

See also

"FileDataSourceName connection parameter" on page 180

DisableMultiRowFetch connection parameter [DMRF]

Function

To turn off multi-row fetches across the network

Usage

Anywhere

Values

YES, NO

Default

NO

Description

By default, when the database server gets a simple fetch request, the application asks for extra rows. You can disable this behavior by setting this parameter to **ON**.

🔗 For more information, see "Using cursors in procedures and triggers" on page 545 of the book *ASA SQL User's Guide*.

🔗 Setting the **DisableMultiRowFetch (DMRF)** connection parameter to **ON** is equivalent to setting the **PREFETCH** database option to **OFF**. For more information, see "Prefetching rows" on page 40 of the book *ASA Programming Guide*.

Example

- ◆ The following connection string fragment prevents prefetching:

```
DMRF=YES
```

EngineName connection parameter [ENG]

Function

Synonym for `ServerName`. The name of a running database server to which you want to connect.

Usage

Network servers or running personal servers.

Values

String

Default

The default local database server.

Description

`EngineName` is not needed if you wish to connect to the default local database server.

You need to supply an `EngineName` only if more than one local database server is running, or if you wish to connect to a network server. In the Sybase Central and Interactive SQL Connect dialog, and in the ODBC Administrator, this is the Server Name field.

The server name is interpreted according to the character set of the client machine. Multi-byte characters are not recommended in server names.

Names must be a valid identifier. Long engine names are truncated to different lengths depending on the protocol.

Protocol	Truncation Length
Unix shared memory	31 bytes
non-Unix shared memory	40 bytes
TCP/IP	40 bytes
SPX	32 bytes
Named Pipes	8 bytes

Example

- ◆ Connect to a server named Guelph:

ENG=Guelph

See also

"Identifiers" on page 7 of the book *ASA SQL Reference Manual*
"-n server option" on page 145

EncryptedPassword connection parameter [ENP]

Function

To provide a password, stored in an encrypted fashion in a data source.

Usage

Anywhere

Values

String

Default

None

Description

Data sources are stored on disk as a file or in the registry. Storing passwords on disk may present a security problem. For this reason, when you enter a password into a data source, it is stored in an encrypted form.

If both the **Password (PWD)** connection parameter and the **EncryptedPassword (ENP)** connection parameter are specified, **Password (PWD)** takes precedence.

Encryption connection parameter [ENC]

Function

To encrypt packets sent between the client application and the server.

Usage

For ECC_TLS (Certicom), RSA_TLS, TCP/IP only.

For *none* or *simple*, anywhere.

Values

String

Default

NONE

Description

If an Encryption value is not set, encryption is controlled by the setting on the server, which defaults to no encryption.

You can use this parameter if you are concerned about the security of network packets. Encryption does affect performance marginally. The **Encryption (ENC)** connection parameter accepts the following arguments:

none accepts communication packets that are not encrypted. This value is equivalent to NO in previous versions of Adaptive Server Anywhere.

simple accepts communication packets that are encrypted with simple encryption supported on all platforms and on previous versions of Adaptive Server Anywhere. This value is equivalent to YES in previous versions of Adaptive Server Anywhere.

ECC_TLS (formerly Certicom) accepts communication packets that are encrypted using Certicom encryption technology. To use this type of encryption, both the server and the client must be operating on Solaris, Linux, NetWare, and all supported Windows operating systems except Windows CE, and the connection must be over the TCP/IP port. UNIX platforms, except for Solaris and Linux, do not recognize the client or server **Certicom** parameter. To authenticate the server, the Certicom software verifies that the server's certificate values match any values you supply about the client using the following arguments:

- ◆ **trusted_certificates** specify the certificate file the client uses to authenticate the server.
- ◆ **certificate_company** specify the value for the organization field. The server's value and the client's value must match.
- ◆ **certificate_unit** specify the value for the organization unit field. The server's value and the client's value must match.
- ◆ **certificate_name** specify the certificate's common name. The server's value and the client's value must match.

RSA_TLS accepts communication packets that are encrypted using RSA encryption technology. To use this type of encryption, both the server and the client must be operating on Solaris, Linux, NetWare, and all supported Windows operating systems except Windows CE, and the connection must be over the TCP/IP port. UNIX platforms, except for Solaris and Linux, do not recognize the client or server **RSA_TLS** parameter. To authenticate the server, the Certicom software verifies that the server's certificate values match any values you supply about the client using the following arguments:

- ◆ **trusted_certificates** specify the certificate file the client uses to authenticate the server.

- ◆ **certificate_company** specify the value for the organization field. The server's value and the client's value must match.
- ◆ **certificate_unit** specify the value for the organization unit field. The server's value and the client's value must match.
- ◆ **certificate_name** specify the certificate's common name. The server's value and the client's value must match.

Caution

The sample certificate should be used for testing purposes only. The sample certificate provides no security in deployed situations because it and the corresponding password are widely distributed with Sybase software. To protect your system, you must create your own certificate.

For information about certificates, see "Self-signed certificates" on page 294 of the book *MobiLink Synchronization User's Guide*.

You can use the *connection_property* system function to retrieve the encryption settings for the current connection. The function returns one of three values: none, simple, or Certicom, depending which type of encryption is being used.

For information about using the *connection_property* system function, see "CONNECTION_PROPERTY function" on page 113 of the book *ASA SQL Reference Manual*.

Example

- ◆ The following connection string fragment connects to a database server *myeng* with a TCP/IP link, using Certicom encryption and the sample trusted certificate:

```
"ENG=myeng; LINKS=tcPIP; Encryption=ECC_TLS
(trusted_certificates=sample.crt) "
```

- ◆ The following connection string fragment connects to a database server *myeng* with a TCP/IP link, using RSA encryption and the sample trusted certificate:

```
"ENG=myeng; LINKS=tcPIP; Encryption=RSA_TLS
(trusted_certificates=sample.crt) "
```

See also

"-ec server option" on page 135
 "Encrypting client/server communications" on page 398

Encryption Key connection parameter [DBKEY]**Function**

To start an encrypted database with a connect request.

Usage	Anywhere
Values	<i>String</i>
Default	None
Description	<p>You must specify this parameter when you start an encrypted database with a connect request. You do not need to specify this parameter if you are connecting to an encrypted database that is already running.</p> <p>The encryption key is a string, including mixed cases, numbers, letters, and special characters.</p>
Example	<p>♦ The following fragment illustrates the use of the Encryption Key (DBKEY) connection parameter:</p> <pre>"UID=dba ; PWD=sql ; ENG=myeng ; DBKEY=V3mo j3952B"</pre>
See also	<p>"-ek database option" on page 159</p> <p>"-ep server option" on page 137</p> <p>"Encrypting client/server communications" on page 398</p>

FileDataSourceName connection parameter [FILEDSN]

Function	<p>The FileDataSourceName (FILEDSN) connection parameter tells the client library that there is an ODBC file data source holding information about the database to which you want to connect.</p> <p>Both ODBC and Embedded SQL applications can use File data sources</p>
Usage	Anywhere
Values	<i>String</i>
Default	There is no default name.
Description	<p>File data sources hold the same information as ODBC data sources stored in the registry. File data sources can be easily distributed to end users so that connection information does not have to be reconstructed on each machine.</p>
See also	"DataSourceName connection parameter" on page 175

ForceStart connection parameter [FORCESTART]

Function	To start a server without attempting to connect to one.
Usage	Only with the <i>db_start_engine</i> function
Values	YES, NO

Default NO

See also "db_start_engine function" on page 240 of the book *ASA Programming Guide*

Idle connection parameter [IDLE]

Function To specify the connection's idle timeout period.

Usage Anywhere except with TDS and Shared Memory connections. Shared Memory and TDS connections (including jConnect) ignore the Adaptive Server Anywhere **Idle (IDLE)** connection parameter.

Values *Integer*

Default *None*

Description The **Idle (IDLE)** connection parameter applies only to the current connection. You can have multiple connections on the same server set to different timeout values.

If no connection idle timeout is specified, the connection timeout value defaults to either the specified server timeout value, or to the server timeout default value of 240 minutes.

In case of a conflict between timeout values, the connection timeout value supersedes any server timeout value whether specified or unspecified.

Example ♦ The following connection string fragment sets the timeout value for this connection to 10 minutes:

```
"ENG=myeng;LINKS=tcip;IDLE=10"
```

See also "-ti server option" on page 150

Integrated connection parameter [INT]

Function To use the integrated login facility.

Usage Anywhere

Values YES, NO

Default NO

Description The **Integrated (INT)** connection parameter has the following settings:

- ◆ **YES** An integrated login is attempted. If the connection attempt fails and the LOGIN_MODE option is set to Mixed, a standard login is attempted.
- ◆ **NO** This is the default setting. No integrated login is attempted.

For a client application to use an integrated login, the server must be running with the LOGIN_MODE database option set to Mixed or Integrated.

Example

- ◆ The following data source fragment uses an integrated login:

```
INT=YES
```

See also

"LOGIN_MODE option" on page 577

LazyClose connection parameter [LCLOSE]

Function	Enabling this option causes the CLOSE <i>cursor-name</i> database request to be queued, and then sent to the server with the next database request. This eliminates a network request each time a cursor is closed.
Usage	Anywhere
Values	YES, NO
Default	NO
Description	<p>When this parameter is enabled, cursors are not actually closed until the next database request. Any isolation level 1 cursor stability locks still apply to the cursor while the CLOSE <i>cursor-name</i> database request is queued.</p> <p>Enabling this option can improve performance if your:</p> <ul style="list-style-type: none">◆ network exhibits poor latency◆ application sends many cursor open and close requests <p>Note that in rare circumstances, canceling the next request after the CLOSE <i>cursor-name</i> database request can leave the cursor in a state where it appears to be closed on the client side, but is not actually closed on the server side. Subsequent attempts to open another cursor with the same name will fail. Using LazyClose is not recommended if your application cancels requests frequently.</p>

LivenessTimeout connection parameter [LTO]

Function	To control the termination of connections when they are no longer intact.
Usage	Network server only.

Windows 95/98/Me and Windows NT/2000/XP, NetWare, and multi-threaded UNIX applications only.

Values *Integer*

Default **120**

If no LivenessTimeout value is set, the liveness timeout is controlled by the setting on the server, which defaults to 120 seconds.

Description A **liveness packet** is sent periodically across a client/server TCP/IP or SPX communications protocol to confirm that a connection is intact. If the client runs for the liveness timeout period without detecting a liveness packet, the communication is severed.

Liveness packets are sent at an interval of one quarter of the LivenessTimeout value when the connection is idle; that is, there have been no request or response packets sent back and forth for at least one quarter of the LivenessTimeout value.

Example ♦ The following sets a Liveness timeout value of 60 seconds

LTO=60

Alternatively, you can set this parameter by entering its value in the Liveness Timeout text box of the Network tab of the ODBC Configuration dialog.

Logfile connection parameter [LOG]

Function To send client error messages and debugging messages to a file.

Usage Anywhere

Values *String*

Default No log file

Description If you want to save client error messages and debugging messages in a file, use the **Logfile (LOG)** connection parameter.

If the file name includes a path, it is relative to the current working directory of the client application.

The **LogFile (LOG)** connection parameter is connection-specific, so from a single application you can set different LogFile arguments for different connections.

Example The following command line starts Interactive SQL connecting to the ASA 8.0 Sample data source with a **LogFile (LOG)** connection parameter:

```
dbisql -c "DSN=ASA 8.0 Sample;LOG=d:\logs\test.txt"
```

Typical log file contents are as follows:

```
Fri Oct 27 2000 11:45
Application information:
"HOST=NAME-PC;OS=Windows NT 4.0 (Service Pack
5);PID=0x14b;THREAD=0x148;EXE=C:\ASA80\WIN32\DBPING.EXE;
VERSION=8.0.0.1271;API=DBLIB;TIMEZONEADJUSTMENT=-300"
Attempting to connect using:
UID=dba;PWD=***;ENG=name;DBG=YES;LOG=c:\temp\cli.out;LIN
KS=shmem,tcpip
Attempting to connect to a running server...
Trying to start SharedMemory link ...
    SharedMemory link started successfully
Attempting SharedMemory connection (no asasrv.ini cached
address)
Failed to connect over SharedMemory
Trying to start TCPIP link ...
Loading wsock32.dll
Loading ws2_32.dll
TCP using Winsock version 2.0
My IP address is 172.31.142.196
My IP address is 127.0.0.1
    TCPIP link started successfully
Attempting TCPIP connection (address 172.31.143.196:2638
found in asasrv.ini cache)
Trying to find server at cached address
172.31.143.196:2638 without broadcasting
    Server not found (no reply received)
Looking for server with name NAME
I am in a class B network
Sending broadcast to find server
Using broadcast address of: 172.31.255.255:2638
I am in a class A network
Sending broadcast to find server
Using broadcast address of: 127.255.255.255:2638
Found database server at address 172.31.142.196:2638
Found database server NAME on TCPIP link
Connected to server over TCPIP at address
172.31.142.196:2638
Writing server address 172.31.142.196:2638 to asasrv.ini
cache
    Liveness timeout 120, liveness retransmit period 30
Connected to the server, attempting to connect to a
running database...
Connected to database successfully
```

Password connection parameter [PWD]

Function	To provide a password for the connection.
-----------------	---

Usage	Anywhere
Values	<i>String</i>
Default	No password provided.
Description	<p>Every user of a database has a password. The password must be supplied for the user to be allowed to connect to the database.</p> <p>The Password (PWD) connection parameter is not encrypted. If you are storing passwords in a data source, you should use the EncryptedPassword (ENP) connection parameter. Sybase Central and the Adaptive Server Anywhere ODBC configuration tool both use encrypted parameters.</p> <p>If both Password (PWD) connection parameter and the EncryptedPassword (ENP) connection parameter are specified, the Password (PWD) connection parameter takes precedence.</p>
Example	<p>◆ The following connection string fragment supplies the user ID DBA and password SQL.</p> <pre>UID=DBA ; PWD=SQL</pre> <p>Alternatively, you can set these parameters in the User ID and Password text boxes in the Connect dialog and ODBC Administrator dialog.</p>
See also	"EncryptedPassword connection parameter" on page 177

PrefetchBuffer connection parameter [PBUF]

Function	Set the maximum amount of memory for buffering rows, in kilobytes.
Usage	Anywhere
Values	<i>Integer</i>
Default	16 (Windows CE) 64 (all other platforms)
Description	<p>The PrefetchBuffer (PBUF) connection parameter controls the memory allocated on the client to store prefetched rows. In some circumstances, increasing the number of rows prefetched from the database server by the client can improve query performance. You can increase the number of rows prefetched using the PrefetchRows (PROWS) and PrefetchBuffer (PBUF) connection parameters.</p> <p>Increasing the PrefetchBuffer (PBUF) connection parameter increases the amount of memory used to buffer GET DATA requests. This may improve performance for some applications that process many GET DATA (SQLGetData) requests.</p>

🔗 For more information, see "PrefetchRows connection parameter" on page 186.

Examples

- ◆ The following connection string fragment could be used to determine if the PrefetchBuffer memory limit is reducing the number of rows prefetched.

```
...PrefetchRows=100;Debug=YES;Logfile=c:\client.txt
```

- ◆ The following string could be used to increase the memory limit to 256K:

```
...PrefetchRows=100;PrefetchBuffer=256
```

PrefetchRows connection parameter [PROWS]

Function	Set the maximum number of rows to prefetch when querying the database.
Usage	Anywhere
Values	<i>Integer</i>
Default	10
Description	Increasing the number of rows prefetched from the database server by the client can improve performance on cursors that only fetch relative 0 or 1, with either single row or wide fetches. Wide fetches include embedded SQL array fetches and ODBC block fetches.

Improvements occur particularly under the following conditions:

- ◆ The application fetches many rows (several hundred or more) with very few absolute fetches.
- ◆ The application fetches rows at a high rate, and the client and server are on the same machine or connected by a fast network.
- ◆ Client/server communication is over a slow network, such as a dial-up link or wide area network.

The number of rows prefetched is limited both by the **PrefetchRows (PROWS)** connection parameter and the **PrefetchBuffer (PBUF)** connection parameter, which limits the memory available for storing prefetched rows.

🔗 For more information, see "PrefetchBuffer connection parameter" on page 185.

Example

- ◆ The following connection string fragment sets the number of prefetched rows to 100:

```
...PrefetchRows=100;...
```

ServerName connection parameter [ENG]

This is a synonym for the **EngineName (ENG)** connection parameter.

🔗 For more information, see "EngineName connection parameter" on page 176.

StartLine connection parameter [START]

Function	To start a database server running from an application.
Usage	Embedded databases
Values	<i>String</i>
Default	No StartLine parameter
Description	You should supply a StartLine (START) connection parameter only if you are connecting to a database server that is not currently running. The StartLine (START) connection parameter is a command line to start a personal database server.

🔗 For a detailed description of available command line options, see "The database server" on page 120.

Example	◆ The following data source fragment starts a personal database server with a cache of 8 Mb.
----------------	--

```
StartLine=dbeng8 -c 8M asademo.db
```

Unconditional connection parameter [UNC]

Function	To stop a server using <i>dbstop</i> even when there are connections to the server.
Usage	Anywhere
Values	YES, NO
Default	NO
Description	The <i>dbstop</i> utility shuts down a database server. If you specify UNC=YES in the connection string, the server is shut down even if there are active connections. If Unconditional is not set to YES, then the server is shut down only if there are no active connections.

Example	◆ The following command line shuts down the server unconditionally:
----------------	---

```
dbstop -c "UID=DBA;PWD=SQL;ENG=server-name;UNC=YES"
```

See also "Stopping a database server using the dbstop command-line utility" on page 505

Userid connection parameter [UID]

Function	The user ID with which you log on to the database.
Usage	Anywhere
Values	<i>String</i>
Default	<i>None</i>
Description	You must always supply a user ID when connecting to a database.
Example	<ul style="list-style-type: none">◆ The following connection string fragment supplies the user ID DBA and password SQL: <pre>uid=DBA;pwd=SQL</pre>

Network communications parameters

Communication parameters (for both the client and the server) enable you to work around peculiarities of different network protocol implementations.

You can supply the network communication parameters in the server command. For example:

```
dbsrv8 -x tcpip(PARM1=value1;PARM2=value2;. . .),SPX
```

From the client side, you enter the communications parameters as the **CommLinks (LINKS)** connection parameter:

```
CommLinks=tcpip(PARM1=value1;PARM2=value2;. . .),SPX
```

If there are spaces in a parameter, the network communication parameters must be enclosed in quotation marks to be parsed properly by the system command interpreter:

```
dbsrv8 -x "tcpip(PARM1=value 1;PARM2=value 2;...),SPX"
CommLinks="tcpip(PARM1=value 1;PARM2=value 2;...),SPX"
```

The quotation marks are required under UNIX if more than one parameter is given because UNIX interprets the semicolon as a command separator.

Boolean parameters are turned on with YES, ON, TRUE, or 1, and are turned off with any of NO, OFF, FALSE, and 0. The parameters are case insensitive.

The examples provided should all be entered on a single line; you can also include them in a configuration file and use the @ server option to invoke the configuration file.

The parameters currently available for TCP/IP and SPX are as follows. PrefetchOnOpen is an ODBC-only communication parameter

TCP/IP and SPX communications parameters

TCP/IP	SPX
Broadcast [BCAST]	BroadcastListener [BLISTENER]
BroadcastListener [BLISTENER]	DLL
ClientPort [CPORT]	DoBroadcast [DOBROAD]
DLL	ExtendedName [ENAME]
DoBroadcast [DOBROAD]	Host [IP]
Host [IP]	RegisterBindery [REGBIN]
MyIP [ME]	SearchBindery [BINSEARCH]
ReceiveBufferSize [RCVBUFSZ]	Timeout [TO]

TCP/IP	SPX
SendBufferSize [SNDBUFSZ]	
ServerPort [PORT]	
TDS	
Timeout [TO]	
LocalOnly [LOCAL]	
VerifyServerName [VERIFY]	

Broadcast communication parameter [BCAST]

Usage	TCP/IP
Values	<i>String</i> (in the form of an IP address)
Default	Broadcasts to all addresses on the same subnet
Description	<p>BROADCAST specifies the IP address used by your TCP/IP protocol implementation to identify a broadcast message.</p> <p>Broadcast addresses consist of the network IP address portion, with 255 as the remaining integers. For example</p> <ul style="list-style-type: none">♦ if the network portion is 95 (class A), then the broadcast address would be 95.255.255.255♦ if the network portion is 132.10 (class B), then the broadcast address would be 132.10.255.255♦ if the network portion is 197.31.175 (class C), then the broadcast address would be 197.31.175.255 <p>The broadcast IP address 255.255.255.255 broadcasts to all subnets.</p>

BroadcastListener communication parameter [BLISTENER]

Usage	SPX, TCP/IP, Server side
Values	YES, NO
Default	YES
Description	This option allows you to turn broadcast listening OFF for this port.

Using `-sb 0` is the same as specifying `BroadcastListener=NO` on both TCP/IP and SPX.

Example

- ◆ Start a server that accepts both TCP/IP and SPX connections, but require that TCP/IP connections use `DoBroadcast=NO`:

```
dbsrv8 -x tcpip(BroadcastListener=NO),spx ...
```

See also

"`-sb server option`" on page 150

ClientPort communication parameter [CPORT]

Usage

TCP/IP. Client side only.

Values

Integer

Default

Assigned dynamically per connection by the networking implementation. If you do not have firewall restrictions, it is recommended that you do not use this parameter.

Description

This option is provided for connections across firewalls, as firewall software filters according to TCP/UDP port. It is recommended that you do not use this parameter unless you need to for firewall reasons.

The `ClientPort` option designates the port number on which the client application communicates using TCP/IP. You may specify a single port number, or a combination of individual port numbers and ranges of port numbers.

It is best to specify a list or a range of port numbers if you want to make multiple connections using a given Data Source or a given connect string. If you specify a single port number, then your application will be able to maintain only one connection at a time. In fact, even after closing the one connection, there is a short timeout period during which no new connection can be made using the specified port. When you specify a list and/or range of port numbers, the application keeps trying port numbers until it finds one to which it can successfully bind.

Examples

- ◆ The following string makes a connection from an application using port 6000 to a server named `my_server` using port 5000:

```
CommLinks=tcpip{ClientPort=6000;ServerPort=5000};
ServerName=my_server
```

- ◆ The following string makes a connection from an application that can use ports 5050 through 5060, as well as ports 5040 and 5070 for communicating with a server named `my_server` using the default server port:

```
CommLinks=tcpip{ClientPort=5040,5050-5060,5070};ServerName=my_server
```

See also

"Host communication parameter" on page 194
"DoBroadcast communication parameter" on page 192
"ServerPort communication parameter" on page 198
"Connecting across a firewall" on page 93

DLL communication parameter

Usage	TCP/IP, SPX (Windows 95/98/Me, Windows NT/2000/XP)
Values	<i>String</i>
Default	<ul style="list-style-type: none">◆ On Windows NT/2000/XP, the default is ws2_32.dll (Winsock 2.0).◆ On Windows 95/98/Me the default is wsock32.dll (Winsock 1.1).
Description	To support untested TCP/IP protocol stacks where the required networking interface functions are in DLLs that differ from the default protocol stack. The client or server looks for its required functionality in the named DLLs.
Example	<ul style="list-style-type: none">◆ The following command starts a server using Winsock 1.1: <pre>dbsrv8 -x tcpip(dll=wsock32.dll) asademo</pre>

DoBroadcast communication parameter [DOBROAD]

Usage	TCP/IP (all platforms) SPX (all platforms except UNIX and CE)
Values	ALL, NONE, DIRECT (Client side) YES, NO (Server side)
Default	ALL
Description	Client usage With DoBroadcast=ALL (formerly DoBroadcast=YES) a broadcast is performed to search for a server. The broadcast goes first to the local subnet. If HOST= is specified, broadcast packets are also sent to each of the hosts. For TCP, all broadcast packets are UDP packets. For SPX, the broadcast is only performed if the server is not found in the bindery. For SPX, all broadcast packets are IPX packets.

With `DoBroadcast=DIRECT` (formerly `DoBroadcast=NO`), no broadcast is performed to the local subnet to search for a database server. Broadcast packets are sent only to the hosts listed in the **HOST (IP)** communication parameter. If you specify `DoBroadcast=DIRECT`, the **HOST (IP)** communication parameter is required.

Specifying `DoBroadcast=NONE` causes no UDP or IPX broadcasts to be used. A TCP/IP or SPX connection is made directly with the **HOST/PORT** specified, and the engine name is verified. With TCP/IP, you can choose not to verify the engine name by setting the **VerifyServerName (VERIFY)** communication parameter to **NO**. The **HOST (IP)** communication parameter is a required parameter, while the **ServerPort (PORT)** communication parameter is optional.

For **DIRECT** and **NONE**, you must specify the server host with the **HOST** option.

Server usage Setting `DoBroadcast=NO` prevents the database server from broadcasting to find other servers with the same name. This is useful in certain rare circumstances, but it is not generally recommended.

Example

- ◆ The following command starts a client without broadcasting to search for a database server. Instead, the server is looked for only on the computer named silver.

```
CommLinks=tcpip(DOBROADCAST=DIRECT;HOST=silver)
asademo
```

ExtendedName communication parameter [ENAME]

Usage	SPX (platforms other than Windows 95/98/Me or Windows NT/2000/XP)
Values	YES, NO
Default	NO
Description	According to the Novell standard for legal SAP names, the following characters are not allowed:

\ / : ; , * ? + -

If you start a server named "asademo-1", the default behavior is to strip out the - and try to start a server asademo1. By turning on **ExtendedName**, the name is left untouched.

This is an SPX port parameter and is only useful when starting a server.

Caution

Users should be wary of using this option as it is contrary to the SAP standard.

Example

- ◆ The following command starts a NetWare server with the name asademo-1.

```
load dbsrv8.nlm -x spx(ExtendedName=YES) asademo-1
```

Host communication parameter [IP]

Usage

TCP/IP, SPX (all platforms)
Server and client sides

Values

String

Default

No additional machines.

Description

HOST specifies additional machines outside the immediate network to be searched by the client library. On the server, the search is carried out to avoid starting a server with a duplicate name.

For TCP/IP, the HOST value the *hostname* or a dot-separated IP address may be used. You may optionally specify a PORT value as well.

For SPX, an address of the form *a:b:c:d:e:f/g:h:i:j* is used, where *a:b:c:d:e:f* is the node number (Ethernet card address) of the server, and *g:h:i:j* is the network number. You may optionally specify a PORT value.

The server prints addressing information to the database server window during startup if the *-z* option is used. In addition, the application writes this information to its logfile if Debug is set to true and LogFile is specified.

You can use a comma-separated list of addresses to search for more than one machine. You can also append a port number to an IP address, using a colon as separator. Alternatively, you can specify the host and server ports explicitly, as in *HOST=a:b:c:d:e:f/g:h:i;j;PORT=k*.

To specify multiple values for a single parameter, use a comma-separated list. When you specify multiple ports and servers, you can associate a particular port with a specific server by specifying the port in the **HOST (IP)** communication parameter instead of the PORT parameter.

IP and **HOST** are synonyms when using TCP/IP. For SPX, you must use **HOST**. **PORT** and **ServerPort** are synonyms for both protocols.

Examples

- ◆ The following connection string fragment instructs the client to look on the machines kangaroo and 197.75.209.222 (port 2369) to find a database server:

```
LINKS=tcip(IP=kangaroo,197.75.209.222:2369)
```

- ◆ The following connection string fragment instructs the client to look on the machines `my_server` and `kangaroo` to find a database server. A connection is attempted to the first host that responds.

```
LINKS=tcip(HOST=my_server,kangaroo;PORT=2639)
```

- ◆ The following connection string fragment instructs the client to look for a server on `host1` running on port 1234 and for a server on `host2` running on port 4567. The client does not look on `host1` on port 4567 or on `host2` on port 1234.

```
LINKS=tcip(HOST=host1:1234,host2:4567)
```

See also "ClientPort communication parameter" on page 191

LocalOnly communication parameter [LOCAL]

Usage TCP/IP, client side only

Values YES, NO

Default NO

Description The **LocalOnly (LOCAL)** communication parameter allows a client to choose to connect only to a server on the local machine, if one exists. If no server with the matching engine name is found on the local machine, a server will not be autostarted. The **LocalOnly (LOCAL)** communication parameter is only useful if `DoBroadcast=ALL` (the default)

`LocalOnly=YES` uses the regular broadcast mechanism, except that broadcast responses from servers on other machines are ignored.

See also "Broadcast communication parameter [BCAST]" on page 190

MyIP communication parameter [ME]

Usage TCP/IP

Values *String*

Description The **MyIP (ME)** communication parameter is provided for machines with more than one network adapter.

Each adapter has an IP address. By default, the database server uses the first network card it finds. If you want your database server to use more than one network card, specify the address of each card in the **MyIP (ME)** communication parameter.

If the keyword **NONE** is supplied as the IP number, no attempt is made to determine the addressing information. The **NONE** keyword is intended for clients on machines where this operation is expensive, such as machines with multiple network cards or remote access (RAS) software and a network card. It is not intended for use on the server.

Under Windows 95/98/Me or Windows NT/2000/XP, this option can be used multiple times for machines with multiple IP addresses.

You can optionally append a port number to the IP address, separated by a colon.

Example

- ◆ The following command line (entered all on one line) instructs the server to use two network cards, one with a specified port number.

```
dbsrv8 -x  
tcpip(MyIP=192.75.209.12:2367,192.75.209.32)  
"c:\Program Files\Sybase\SQL Anywhere 8\asademo.db"
```

- ◆ The following connection string fragment instructs the client to make no attempt to determine addressing information.

```
LINKS= tcpip(MyIP=NONE)
```

PreFetchOnOpen communication parameter

Usage ODBC

Values YES, NO

Default NO

Description Enabling this option sends a prefetch request with a cursor open request, thereby eliminating a network request to fetch rows each time a cursor is opened. Columns must already be bound in order for the prefetch to occur on the open. Rebinding columns between the cursor open and the first fetch when using PrefetchOnOpen will cause reduced performance.

Calling ODBC's **SQLExecute** or **SQLExecDirect** on a query or stored procedure which returns a result set causes a cursor open.

Enabling this option can improve performance if your:

- ◆ network exhibits poor latency
- ◆ application sends many cursor open and close requests

ReceiveBufferSize communication parameter [RCVBUFSZ]

Usage	TCP/IP
Values	<i>Integer</i>
Default	Machine-dependent
Description	Sets the size for a buffer used by the TCP/IP protocol stack. You may want to increase the value if blob performance over the network is important.

RegisterBindery communication parameter [REGBIN]

Usage	SPX (Windows 95/98/Me and Windows NT/2000/XP only). Server side only.
Values	YES, TRUE, 1 NO, FALSE, 0
Default	TRUE
Description	The database server attempts to register its name with any active binderies on the network when loading the SPX link. To disable this name registration, set RegisterBindery to NO, FALSE, or 0. In this case, the client library must be able to locate the database server over SPX by broadcasting packets.

SearchBindery communication parameter [BINSEARCH]

Usage	SPX (Windows 95/98/Me and Windows NT/2000/XP only)
Values	YES, TRUE, 1 NO, OFF, 0
Default	YES
Description	With SEARCHBINDERY=NO, 0, or OFF, no NetWare bindery is searched for a database server.

SendBufferSize communication parameter [SNDBUFSZ]

Usage	TCP/IP
Values	<i>Integer</i>
Default	Machine-dependent

Description Sets the size for a buffer used by the TCP/IP protocol stack. You may want to increase the value if blob performance over the network is important.

ServerPort communication parameter [PORT]

Usage TCP/IP (all platforms)

Values *Integer*

Default **2638**

Description The Internet Assigned Numbers Authority has assigned the Adaptive Server Anywhere database server port number 2638 to use for TCP/IP communications. However, applications are not disallowed from using this reserved port, and this may result in an addressing collision between the database server and another application.

In the case of the database server, the ServerPort option designates the port number on which to communicate using TCP/IP.

In a data source, the ServerPort option informs the client of the port or ports on which database servers are listening for TCP/IP communication. The client broadcasts to every port that is specified on the **ServerPort (PORT)** communication parameter to find the server.

The database server always listens on port 2638, even if you specify a different port using a network communication parameter. Hence, applications can connect to the database server without specifying a port number. An exception is the HP-UX operating system, on which the server does not listen on port 2638 if it is started on another port.

Example

1 Start a network database server:

```
dbsrv8 -x tcpip -n server1
```

Port number 2638 is now taken.

2 Attempt to start another database server:

```
dbsrv8 -x tcpip -n server2
```

The default port is currently allocated, and so the server starts on another port.

TDS communication parameter

Usage TCP/IP, NamedPipes
Server side only

Values	YES, NO
Default	YES
Description	To disallow TDS connections to a database server, set TDS to NO. If you want to ensure that only encrypted connections are made to your server, these port options are the only way to disallow TDS connections.
Example	<ul style="list-style-type: none"> ◆ The following command starts a database server using the TCP/IP protocol, but disallowing connections from Open Client or jConnect applications. <pre>dbsrv8 -x tcpip(TDS=NO) ...</pre>

Timeout communication parameter [TO]

Usage	TCP/IP, SPX (Windows 95/98/Me, Windows NT/2000/XP, NetWare)
Values	<i>Integer</i> , in seconds
Default	5
Description	Timeout specifies the length of time, in seconds, to wait for a response when establishing communications. It also specifies the length of time to wait for a response when disconnecting. You may wish to try longer times if you are having trouble establishing TCP/IP or SPX communications.
Example	<ul style="list-style-type: none"> ◆ The following data source fragment starts a TCP/IP communications link only, with a timeout period of twenty seconds. <pre>... CommLinks=tcpip(TO=20) ...</pre>

VerifyServerName communication parameter [VERIFY]

Usage	TCP/IP Client side only
Values	YES, NO
Default	YES

Description

When connecting over TCP using the `DoBroadcast=NONE` parameter, the client makes a TCP connection, then verifies that the name of the server found is the same as the one it's looking for. Specifying `VerifyServerName=NO` skips the verification of the server name. This allows Adaptive Server Anywhere clients to connect to an Adaptive Server Anywhere server if they know only an IP address/port.

The server name must still be specified in the connect string, but it is ignored. The **VerifyServerName (VERIFY)** communication parameter is used only if `DoBroadcast=NONE` is specified.

Note

We recommend that you only use this parameter in the rare circumstances when it is not possible to give each server a unique server name, and use that unique name to connect. Giving each server a unique server name, and connecting to the server using that name is still the best way to connect.

See also

"DoBroadcast communication parameter" on page 192

PART TWO

Working with Database Files

This section describes how to install, use, and backup database files. It also describes how to use international language and character sets.

C H A P T E R 7

File Locations and Installation Settings

About this chapter This chapter describes the installation and operating system settings used by Adaptive Server Anywhere. Depending on the operating system, these settings may be stored as environment variables, initialization file entries, or registry settings.

Contents

Topic	Page
Installation directory structure	204
How Adaptive Server Anywhere locates files	206
Environment variables	208
Registry and INI files	213

Installation directory structure

When you install Adaptive Server Anywhere, several directories may be created. Some of the files in these directories are essential, and others are not. This section describes the directory structure.

Adaptive Server Anywhere software, whether you receive it as a product or bundled as part of another product, is installed under a single installation directory. The tools provided with the Adaptive Server Anywhere product, however, are installed in other directories. This section describes only the installation directory structure for Adaptive Server Anywhere itself.

The Adaptive
Server Anywhere
installation
directory

The Adaptive Server Anywhere installation directory itself holds several items, including the following:

- ◆ **The sample database** The sample database is held in the file *asdemo.db*.
- ◆ **Read Me First** A Read Me First file named *readme.txt* holds late-breaking information.

For platforms other than Novell NetWare and Windows CE, there are several directories under the installation directory:

- ◆ **Executable directories** There is a separate directory for each operating system, which holds executables, dynamic link libraries, and help files.

On Windows except Windows CE, these files are installed in the *win32* directory. If you are using UNIX, they are installed in the *bin* directory. On NetWare, the executables are stored in the installation directory itself.

You will not have all these directories on your machine; you will have only the ones required for the operating system version you installed.

- ◆ **Java directory** Java base classes are stored in this directory.
- ◆ **ProcDebug directory** The stored procedure debugger is stored in this directory.
- ◆ **Scripts directory** The scripts directory contains SQL scripts that are used by the database administration utilities and as examples. With the exception the *custom.SQL* script, *do not edit these scripts*. If the scripts directory is not present, the administration utilities will not work.
- ◆ **Examples directories** There are separate directories for C (*cxmp*) and Java (*jxmp*) examples.

- ◆ **h directory** The h directory contains header files for ESQL and ODBC database development.

Novell NetWare file locations

On Novell NetWare, all files are installed to a single directory on the server. Throughout this documentation, when reference is made to files in subdirectories of the installation directory, the file on NetWare is in the installation directory itself.

Windows CE file locations

On Windows CE, all files are installed to the installation directory, and no subdirectories are created. The exception is that all DLLs are installed into the *Windows* directory.

How Adaptive Server Anywhere locates files

The client library and the database server need to locate files for two main purposes:

- ◆ DLLs and initialization files are required to run Adaptive Server Anywhere. If an incorrect DLL is located, there is the possibility of version mismatch errors.
- ◆ Some files are specified in SQL statements and need to be located at run time, such as **INSTALL** or **LOAD TABLE**.

Examples of SQL statements that use file names include the following:

- ◆ **INSTALL statement** The name of the file that holds Java classes.
- ◆ **LOAD TABLE and UNLOAD TABLE statements** The name of the file from which data should be loaded or to which the data should be unloaded.
- ◆ **CREATE DATABASE statement** A file name is needed for this statement and similar statements that can create files (such as **CREATE WRITEFILE**).

In some cases, Adaptive Server Anywhere uses a simple algorithm to locate files. In other cases, a more extensive search is carried out.

Simple file
searching

In many SQL statements (such as **LOAD TABLE**, or **CREATE DATABASE**), the file name is interpreted as relative to the current working directory of the database server.

Also, when a database server is started and a database file name (**DatabaseFile (DBF)** parameter) is supplied, the path is interpreted as relative to the current working directory.

Extensive file
searching

Adaptive Server Anywhere programs, including the database server and administration utilities, carry out a more extensive search for required files, such as DLLs or shared libraries. In these cases, Adaptive Server Anywhere programs look for files in the following order:

- 1 **Executable directory** Holds the program executable file.
- 2 **Related directories** Holds directories with the following paths relative to the program executable directory:
 - ◆ Parent of the executable directory.
 - ◆ A child of the parent directory named *scripts*. The UNIX server does not search in this location.

- 3 **Current working directory** When a program is started, it has a current working directory (the directory from which it is started). This directory is searched for required files.
- 4 **Location registry entry** When installing onto Windows, Adaptive Server Anywhere adds a LOCATION registry entry. The indicated directory is searched, followed by:
 - ◆ A child named *scripts*
 - ◆ A child with the operating system name (*win32*, *win*, and so on)
- 5 **System specific directories** This includes directories where common operating system files are held, such as the *Windows* directory and the *Windows\system* directory on Windows operating systems.
- 6 **CLASSPATH directories** For Java files, directories listed in the CLASSPATH environment variable are searched to locate files.
- 7 **PATH directories** Directories in the system path and the user's path are searched to locate files.

Environment variables

Adaptive Server Anywhere uses a set of environment variables to store various types of information. These environment variables are listed in this section. Not all variables need to be set in all circumstances.

Setting environment variables

The way you set an environment variable depends on the operating system you are using.

❖ To set an environment variable (Windows NT):

- 1 Right-click My Computer and choose Properties from the popup menu.
- 2 Click the Environment tab. If the environment variable does not already exist, type variable and its value in the spaces provided, and click Set.

If the variable does exist, select it from the list of System Variables or User Variables, and make any modifications in the Value field. Click Set to make the setting.

❖ To set an environment variable (Windows 2000):

- 1 Right-click My Computer and choose Properties from the popup menu.
- 2 Click the Advanced tab.
- 3 Click the Environment Variables button. The Environment Variables dialog opens. If the environment variable does not already exist, click New and type the variable name and its value in the spaces provided, and then click OK.

If the variable does exist, select it from the list of System Variables or User Variables, click Edit and make any modifications in the Variable Value field. Click OK to make the setting.

❖ To set an environment variable (UNIX):

- ◆ In one of your startup files (*.cshrc*, *.shrc*, *.login*), add a line that sets the variable.

In some shells (such as sh, bash, or ksh) the line is as follows:

```
export VARIABLE=value
```

In other shells (such as csh, or tsch) the line is as follows:

```
setenv VARIABLE value
```

ASANY8 environment variable

Syntax	ASANY8=directory-name
Default	<i>C:\Program Files\Sybase\SQL Anywhere 8</i>
Description	<p>On installation, the ASANY8 environment variable is set to your SQL Anywhere directory.</p> <p>This environment variable must be set for several reasons. Among the list, samples require this environment variable in order to locate SQL Anywhere applications.</p>

ASANYSH8 environment variable

Syntax	ASANYSH8=directory-name
Default	Interactive SQL, Sybase Central, the debugger, and the Console utility use this variable to determine the location of the shared components directory. ASANYSH8 is set to the location of the shared directory. You set the shared
Description	directory location during installation. The default shared components directory is <i>C:\Program Files\Sybase\Shared</i> .

ASTMP environment variable

Syntax	ASTMP=directory-name
Default	None.
Description	<p>The database server checks the value of the ASTMP environment variable to determine where to hold the temporary file. If the ASTMP environment variable does not exist, then the first of the TMP, TMPDIR, and TEMP environment variables to exist is used on Windows. If the ASTMP environment variable does not exist on UNIX, <i>/tmp</i> is used.</p> <p>In many circumstances, ASTMP is not needed. It can be of use in security-conscious environments when running the database server as a service to enable you to hold the temporary file in a directory that cannot be accessed by other programs.</p>


LD_LIBRARY_PATH environment variable [UNIX]

Syntax	LD_LIBRARY_PATH = <i>installation_path/lib</i>
Description	<p>The LD_LIBRARY_PATH environment variable is used on UNIX only. It is modified by the installation program to include the directories where Adaptive Server Anywhere libraries are located.</p> <p>The executables are located in the <i>lib</i> subdirectory of the installation directory (for example, <i>/opt/SYBSasa8/lib</i>).</p>


PATH environment variable

Syntax	PATH = <i>installation_path</i>
Description	<p>The PATH environment variable is modified by the installation program to include the directories where Adaptive Server Anywhere executables are located.</p> <p>The executables are located in a subdirectory of the installation directory.</p> <p>In addition, if you are using other Sybase applications, the <i>SYBASE\bin</i> and <i>SYBASE\dll</i> directories are added to your path.</p> <p>On UNIX, each user must have the directory holding the executables (<i>/opt/SYBSasa8/bin</i>) added to their path.</p>

SQLCONNECT environment variable

Syntax	SQLCONNECT = <i>parameter#value; ...</i>
Description	<p>The SQLCONNECT environment variable is optional, and is not set by the installation program.</p> <p>SQLCONNECT specifies connection parameters that are used by several of the database administration utilities when connecting to a database server. This string is a list of parameter settings, of the form parameter=value, delimited by semicolons.</p> <p>The number sign "#" is an alternative to the equal sign, and should be used if you are setting the connection parameters string in the SQLCONNECT environment variable. Using "=" inside an environment variable setting is a syntax error. The = sign is allowed only in Windows NT/2000/XP.</p> <p> For a description of the connection parameters, see "Connection parameters" on page 70.</p>

SQLLOCALE environment variable

Syntax	SQLLOCALE =Charset= <i>cslabel</i> ;Language= <i>langlabel</i> ;CollationLabel= <i>colabel</i>
See also	"Setting the SQLLOCALE environment variable" on page 268
Description	<p>The SQLLOCALE environment variable is not set by the installation program, and is required only in multi-character-set environments.</p> <p>The SQLLOCALE environment variable is a single string that consists of three semi-colon-separated assignments. The assignments set out the character set, language, and collation of the environment.</p> <p> For a list of supported character set labels, see "Setting the SQLLOCALE environment variable" on page 268.</p>

SQLPATH environment variable

Syntax	SQLPATH = <i>path</i> ; ...
Description	<p>The SQLPATH environment variable is optional, and is not set by the installation program.</p> <p>Interactive SQL searches along SQLPATH for command files and Help files before searching the system path.</p>

SQLREMOTE environment variable

Syntax	SQLREMOTE = <i>path</i>
Description	<p>The SQLREMOTE environment variable is optional, and is not set by the installation program.</p> <p>Addresses for the FILE message link in SQL Remote are subdirectories of the SQLREMOTE environment variable. This variable should point to a shared directory.</p> <p>On 32-bit Windows operating systems, an alternative to setting the SQLREMOTE environment variable is to set the <i>SQL Remote\Directory</i> registry entry to the proper root directory.</p>

SYBASE environment variable

Syntax	SYBASE = <i>path</i>
---------------	-----------------------------

Description The SYBASE variable marks the home directory for installation of some Sybase applications, including Adaptive Server Enterprise and utilities such as *DSEdit*. You need this variable only if you are using Adaptive Server Anywhere together with other members of the Adaptive Server family.

TEMP environment variable


Syntax **ASTMP=***path*
 TMP=*path*
 TMPDIR=*path*
 TEMP=*path*

Description The database server creates a temporary file for various operations such as sorting and performing unions. Temporary files are placed in the directory specified by the ASTMP, TMP, TMPDIR, or TEMP environment variable. Adaptive Server Anywhere takes the first one of the three that it finds, in that order.

If none of the environment variables is defined, temporary files are placed in the current working directory of the server.

On UNIX, only the ASTMP variable is used. If the ASTMP variable does not exist, files are placed in the /tmp/.SQLAnywhere directory.

On Windows CE, you can specify which directory you want to use as the server's temporary directory, in the registry.

 For more information about setting the temporary directory value, see "Registry settings on Windows CE" on page 214

Registry and INI files

On Windows operating systems (except Windows CE), Adaptive Server Anywhere uses several registry settings. On UNIX, and NetWare, these settings are held in initialization files instead.

The software makes these settings for you, and in general operation you should not need to access the registry. The settings are provided here for those people who make modifications to their operating environment.

Current user and local machine settings

Some operating systems, such as Windows NT, hold two levels of system settings. Some settings are specific to an individual user and are used only when that user is logged on; these settings are called **current user** settings. Some settings are global to the machine, and are available no matter which user is logged on; these are called **local machine** settings. You must have administrator permissions on your machine to make local machine settings.

Adaptive Server Anywhere permits both current user and local machine settings. On Windows NT, for example, these are held in the HKEY_CURRENT_USER registry and the HKEY_LOCAL_MACHINE registry, respectively.

Current user takes precedence

If a setting is made in both the current user and local machine registries, the current user setting takes precedence over the local machine setting.

When local machine settings are needed

If you are running an Adaptive Server Anywhere program as a **service**, you should ensure that the settings are made at the *local machine* level.

Services can continue to run under a special account when you log off a machine as long as you do not shut the machine down entirely. They can be made independent of individual accounts, and therefore need access to local machine settings.

In addition to Adaptive Server Anywhere programs, some Web servers run as services. You must set local machine settings in order for PowerDynamo to work with such a Web server.

In general, the use of local machine settings is recommended.

Registry structure

On Windows (except for Windows CE), you can access the registry directly with the registry editor. The Adaptive Server Anywhere registry entries are held in either the HKEY_CURRENT_USER or HKEY_LOCAL_MACHINE registries, in the following location:

```
Software
  Sybase
    Adaptive Server Anywhere
      8.0
        Sybase Central
          4.0
            Profiles
            Providers
```

Registry settings on installation

The installation program makes the following registry settings in the Sybase registry:

- ◆ **Location** In the *Adaptive Server Anywhere\8.0* registry, this entry holds the installation directory location. For example:

```
Location "c:\Program Files\Sybase\SQL Anywhere 8"
```

- ◆ **Language** In the *Adaptive Server Anywhere\8.0* registry, this entry holds a two-letter code indicating the current language for messages and errors. For example:

```
Language "EN"
```

The default setting is English (EN). The installation program sets this entry only if the software is installed for a language other than English.

- ◆ **Providers** In the *Sybase Central\Providers* registry, this entry stores the file names of installed plug-ins for Sybase Central. Adaptive Server Anywhere has its own Sybase Central plug-in:

```
Adaptive Server Anywhere 8.0
"c:\sybase\asa8\win32\scasany8.dll"
```

Registry settings on Windows CE

You can specify which directory you want to use as the server's temporary directory on Windows CE by setting the

```
HKEY_CURRENT_USER\Software\Sybase\Adaptive Server
Anywhere\8.0\TempFolder
```


value in the registry, where *TempFolder* is the name of the temporary directory you want to use. The server will either:

- ◆ use the specified directory if it exists, or
- ◆ attempt to create the specified directory if it does not already exist, as long as the parent directory already exists.

If the specified directory does not exist and cannot be created, the server will:

- ◆ use the \Temp directory if it exists, or
- ◆ attempt to create a \Temp directory if it does not already exist.

If the \Temp directory does not exist and cannot be created, the server will:

- ◆ use the use the current directory.

C H A P T E R 8

Working with Database Files

About this Chapter This chapter describes how to create, and work with database and associated files.

Contents

Topic	Page
Overview of database files	218
Using additional dbspaces	220
Working with write files	224
Using the utility database	226

Overview of database files

Basic database files

Each database has the following files associated with it.

- ◆ **The database file** This file holds the database information. It typically has the extension *.db*.

🔗 For information on creating databases, see "Working with databases" on page 29 of the book *ASA SQL User's Guide*.

- ◆ **The transaction log** This file holds a record of the changes made to the database, and is necessary for recovery and replication. It typically has the extension *.log*.

🔗 For information on the transaction log, see "Backup and Data Recovery" on page 299.

- ◆ **The temporary file** The database server uses the temporary file to hold information needed during a database session. The database server disregards this file once the database shuts down—even if the server remains running. The file has a server-generated name with the extension *.tmp*.

The temporary file is held in a location determined by an environment variable.

On Windows, the following environment variables are checked, in order:

- ◆ ASTMP
- ◆ TMP
- ◆ TMPDIR
- ◆ TEMP

If none of these is defined, Adaptive Server Anywhere places its temporary file in the current directory.




On UNIX, the ASTMP environment variable is checked.

If this environment variable does not exist, the temporary file is held in the */tmp/SQLAnywhere* directory.

The server creates, maintains, and removes the temporary file. You only need to ensure that there is enough free space available for the temporary file.

Additional files

Other files can also become part of a database system, including:

- ◆ **Additional database files** You can spread your data over several separate files. These additional files are called dbspaces.
 For information on dbspaces, see "CREATE DBSPACE statement" on page 278 of the book *ASA SQL Reference Manual*.
- ◆ **Transaction log mirror files** For additional security, you can create a mirror copy of the transaction log. This file typically has the extension *.mlg*.
 For information on mirrored transaction logs, see "Transaction log mirrors" on page 306.
- ◆ **Write files** If the database file you are working with is designated read-only (for example, because it is distributed on CD-ROM), you can use an additional write file to hold changes you make to the data. However, note that you can only use write files with read-only database files, not with a server running in read-only mode. Running the server in read-only mode allows no changes to the database file whatsoever.
- ◆ **Compressed database files** You can compress a database file. The resulting file is read only, but can be used in conjunction with a write file. Compressed database files are used in place of the actual database file.
 For information on compression, see "The Compression utility" on page 448.

Using additional dbspaces

This section describes how to use additional database files, known as dbspaces.

Typically needed for large databases

For most databases, a single database file is sufficient. However, for users of large databases, additional database files are often necessary. Additional database files are also convenient tools for clustering related information in separate files.

When you initialize a database, it contains one database file. This first database file is called the **main file**. All database objects and all data are placed, by default, in the main file.

Each database file has a maximum allowable size of 256M database pages. For example, a database file created with a database page size of 4 kb can grow to a maximum size of one terabyte (256M*4 kb). However, in practice, the maximum file size allowed by the physical file system in which the file is created affects the maximum allowable size significantly.

While many commonly-employed file systems restrict file size to a maximum of 2 Gb, some, such as the Windows NT/2000/XP file system, allow you to exploit the full database file size. In scenarios where the amount of data placed in the database exceeds the maximum file size, it is necessary to divide the data into more than one database file. As well, you may wish to create multiple dbspaces for reasons other than size limitations, for example to cluster related objects.

Splitting existing databases

If you wish to split existing database objects among several dbspaces, you need to unload your database and modify the generated command file for rebuilding the database. To do so, add IN clauses to specify the dspace for each table you do not wish to place in the main file.

✍ For more information, see "UNLOAD TABLE statement" on page 573 of the book *ASA SQL Reference Manual* and "Setting properties for database objects" on page 34 of the book *ASA SQL User's Guide*.

Creating a dspace

You create a new database file, or **dspace**, either from Sybase Central, or using the CREATE DBSPACE statement. The database file for a new dspace may be on the same disk drive as the main file or on another disk drive. You must have DBA authority to create dbspaces.

Placing tables in dbspaces

For each database, you can create up to twelve dbspaces in addition to the main dspace.

A newly created dspace is empty. When you create a new table you can place it in a specific dspace with an IN clause in the CREATE TABLE statement. If you don't specify an IN clause, the table appears in the main dspace.

Each table is entirely contained in the dspace it is created in. By default, indexes appear in the same dspace as their table, but you can place them in a separate dspace by supplying an IN clause.

❖ To create a dspace (Sybase Central):

- 1 Connect to the database.
- 2 Open the Dbspaces folder for that database.
- 3 Double-click Add Dspace.
- 4 Follow the instructions in the wizard.

The new dspace then appears in the Dbspaces folder.

❖ To create a dspace (SQL):

- 1 Connect to the database.
- 2 Execute a CREATE DBSPACE statement.

Examples

The following command creates a new dspace called **library** in the file *library.db* in the same directory as the main file:

```
CREATE DBSPACE library
AS 'library.db'
```

The following command creates a table *LibraryBooks* and places it in the *library* dspace.

```
CREATE TABLE LibraryBooks (
  title char(100),
  author char(50),
  isbn char(30)
) IN library
```

 See also

- ◆ "CREATE DBSPACE statement" on page 278 of the book *ASA SQL Reference Manual*
- ◆ "Creating tables" on page 40 of the book *ASA SQL User's Guide*
- ◆ "CREATE INDEX statement" on page 300 of the book *ASA SQL Reference Manual*

Deleting a dbspace

You can delete a dbspace using either Sybase Central or Interactive SQL. Before you can delete a space, you must delete all tables that use the space. You must have DBA authority to delete a dbspace.

❖ To delete a dbspace (Sybase Central):

- 1 Open the Dbspaces folder.
- 2 Right-click the desired dbspace and choose Delete from the popup menu.

❖ To delete a dbspace (SQL):

- 1 Connect to a database.
- 2 Execute a DROP DBSPACE statement.

 See also

- ◆ "Deleting tables" on page 43 of the book *ASA SQL User's Guide*
- ◆ "DROP statement" on page 397 of the book *ASA SQL Reference Manual*

Pre-allocating space for database files

Adaptive Server Anywhere automatically grows database files as needed. Rapidly changing database files can lead to excessive file fragmentation on the disk, resulting in potential performance problems. Unless you are working with a database with a high rate of change, you do not need to worry about explicitly allocating space for database files. If you are working with a database with a high rate of change, you may pre-allocate disk space for dbspaces or for transaction logs using either Sybase Central or the ALTER DBSPACE statement.

You must have DBA authority to alter the properties of a database file.

Performance Tip

Running a disk defragmentation utility after pre-allocating disk space helps ensure that the database file is not fragmented over many disjoint areas of the disk drive. Performance can suffer if there is excessive fragmentation of database files.

❖ To pre-allocate space (Sybase Central):

- 1 Open the Dbspaces folder.

- 2 Right-click the desired dbspace and choose Add Pages from the popup menu.
- 3 Enter the number of pages to add to the dbspace and click OK.

❖ **To pre-allocate space (SQL):**

- 1 Connect to a database.
- 2 Execute an ALTER DBSPACE statement.

Example

Increase the size of the SYSTEM dbspace by 200 pages.

```
ALTER DBSPACE system
ADD 200
```

Example

Increase the size of the SYSTEM dbspace by 400 megabytes.

```
ALTER DBSPACE system
ADD 400 MB
```

 See also

- ◆ "Creating a dbspace" on page 220
- ◆ "ALTER DBSPACE statement" on page 209 of the book *ASA SQL Reference Manual*

Working with write files

If you have a read-only database file (for example, if you distribute a database on a CD-ROM), you can use a write file to make local changes to the database.

You create a write file using the Create Write File utility or using the CREATE WRITEFILE statement. In this section, the examples use the utility.

☞ For a description of the CREATE WRITEFILE statement, see "CREATE WRITEFILE statement" on page 373 of the book *ASA SQL Reference Manual*.

☞ For more information about opening a database as read-only to prevent local changes to the database, see "-r server option" on page 149.

❖ To use a write file:

- 1 Create the write file for your database.

For example, to create a write file for the demo database, execute the following command in a directory containing a copy of the demo database file *asademo.db*:

```
dbwrite -c asademo.db
```

This command creates a write file named *asademo.wrt*, with a transaction log named *asademo.wlg*.

- 2 Start a database server and load the write file. By default, the server locates files with the extension *.wrt* first, so the following command starts the personal server running the demo database write file:

```
dbeng8 asademo
```

Messages on the server window indicate which file starts.

- 3 Connect to the database using Interactive SQL. You can use the user ID **DBA** and the password **SQL**, as the demo database is the default.
- 4 Execute queries as usual. For example, the following query lists the contents of the *department* table.

```
SELECT *  
FROM department
```

The data for the department table is obtained from the database file *asademo.db*.

- 5 Try inserting a row. The following statement inserts a row into the department table:

```
INSERT  
INTO department (dept_id, dept_name)  
VALUES (202, 'Eastern Sales')
```

If you committed this change, it would be written to the *asademo.wlg* transaction log, and when the database checkpoints, the changes are written to the *asademo.wrt* write file.

If you now query the department table, the results come from the write file and the database file.

- 6 Try deleting a row. Set the `WAIT_FOR_COMMIT` option to avoid referential integrity complaints here:

```
SET TEMPORARY OPTION wait_for_commit = 'on' ;  
  
DELETE  
FROM department  
WHERE dept_id = 100
```

If you committed this change, the deletion would be marked in the write file. No changes occur to the database file.

For some purposes, it is useful to use a write file with a shared database. If, for example, you have a read-only database on a network server, each user could have their own write file. In this way, they could add local information, which would be stored on their own machine, without affecting the shared database. This approach can be useful for application development also.

Deleting a write file

You can use the *dberase* utility to delete a write file and its associated transaction log.

Using the utility database

The **utility database** is a phantom database with no physical representation. The utility database has no database file, and therefore it cannot contain data.

The utility database feature allows you to execute database file administration statements such as CREATE DATABASE, or ALTER WRITEFILE without first connecting to an existing physical database.

For example, executing the following statement after having connected to the utility database creates a database named *new.db* in the directory *C:\temp*.

```
CREATE DATABASE 'C:\\temp\\new.db'
```

☞ For more information on the syntax of these statements, see "CREATE DATABASE statement" on page 273 of the book *ASA SQL Reference Manual*.

You can also retrieve values of connection properties and database properties using the utility database.

For example, executing the following statement against the utility database returns the default collation sequence, which will be used when creating database:

```
SELECT property( 'DefaultCollation' )
```

☞ For a list of database properties and connection properties, see "Database properties" on page 618.

Allowed
statements for the
utility database

The following statements are the only ones allowed when connected to the utility database:


- ◆ CREATE DATABASE
- ◆ CREATE WRITEFILE
- ◆ ALTER WRITEFILE
- ◆ CREATE COMPRESSED DATABASE
- ◆ CREATE EXPANDED DATABASE
- ◆ CREATE DECRYPTED FILE
- ◆ DROP DATABASE
- ◆ RESTORE DATABASE
- ◆ START DATABASE
- ◆ STOP DATABASE
- ◆ STOP ENGINE

- ◆ SELECT (with no FROM or WHERE clauses)

Connecting to the utility database

You can start the utility database on a database server by specifying **utility_db** as the database name when connecting to the server. The user ID and password requirements are different for the personal server and the network server.

For the personal database server, there are no security restrictions for connecting to the utility database. It is assumed that anybody who can connect to the personal database server can access the file system directly, and so no attempt is made to screen users based on passwords.

 For more information, see "Utility database passwords" on page 228.

❖ To connect to the utility database on the personal server (Interactive SQL):

- 1 Start a database server with the following command:

```
dbeng8.exe -n TestEng
```
- 2 Start Interactive SQL.
- 3 In the Connect dialog, enter **DBA** as the user ID, and enter any non-blank password. The password itself is not checked, but it must be non-empty.
- 4 On the Database tab, enter **utility_db** as the database name.
- 5 Click OK to connect.

Interactive SQL connects to the utility database on the personal server named **TestEng**, without loading a real database.

For the network server, there are security restrictions on connections. A password is held in the file *util_db.ini* in the same directory as the database server executable.

❖ To connect to the utility database on the network server (Interactive SQL):

- 1 Add a password to the file *util_db.ini* in the same directory as the database server executable. For example, the following *util_db.ini* file has the password ASA.

```
[UTILITY_DB]
PWD=ASA
```

- 2 Start a database server with the following command:

```
dbsrv8.exe -n TestEng
```

- 3 Start Interactive SQL.
- 4 In the Connect dialog, enter **DBA** as the user ID, and enter the password held in the file *util_db.ini*.
- 5 On the Database tab, enter **utility_db** as the database name.
- 6 Click OK to connect.

Interactive SQL connects to the utility database on the network server named **TestEng**, without loading a real database.

 For more information, see "Connecting to a Database" on page 37.

Utility database server security

There are two aspects to utility database server security:

- ◆ Who can connect to the utility database? This is controlled by the use of passwords.
- ◆ Who can execute file administration statements? This is controlled by database server options.

Utility database passwords

The personal server and the network server have different security models for connections.

For the personal server, you must specify the user ID **DBA**. You must also specify a password, but it can be any password. Since the personal server is for single machine use, security restrictions (for example passwords) are unnecessary.

For the network server, you must specify the user ID **DBA**, and the password that is held in a file named *util_db.ini*, stored in the same directory as the database server executable file. As this directory is on the server, you can control access to the file, and thereby control who has access to the password.

The *util_db.ini* file

The *util_db.ini* file has the following contents:

```
[UTILITY_DB]  
PWD=password
```

Use of the **utility_db** security level relies on the physical security of the computer hosting the database server, since the *util_db.ini* file can be easily read using a text editor.

Permission to execute file administration statements

A level of security is provided for the ability to execute certain administration tasks. The `-gu` database server option controls who can execute the file administration statements.

There are four levels of permission for the use of file administration statements. These levels include: **all**, **none**, **DBA**, and **utility_db**. The **utility_db** level permits only a person with authority to connect to the utility database to use the file administration statements.

-gu option	Effect	Applies to
All	Anyone can execute file administration statements	Any database including utility database
None	No one can execute file administration statements	Any database including utility database
DbA	Only DBA-authority users can execute file administration statements	Any database including utility database
Utility_db	Only the user who can connect to utility database can execute file administration statements	Only the utility database

For more information on the database server `-gu` option, see "`-gu` server option" on page 143.

Examples

- ◆ To prevent the use of the file administration statements, start the database server using the **none** permission level of the `-gu` option. The following command starts a database server and names it **TestSrv**. It loads the sample database, but prevents anyone from using that server to create or delete a database, or execute any other file administration statement regardless of their resource creation rights, or whether or not they can load and connect to the utility database.

```
dbsrv8.exe -n TestSrv -gu none asademo.db
```

- ◆ To permit only the users knowing the utility database password to execute file administration statements, start the server at the command prompt with the following command.

```
dbsrv8 -n TestSrv -gu utility_db
```

Assuming the utility database password has been set during installation to **asa**, the following command starts the Interactive SQL utility as a client application, connects to the server named **TestSrv**, loads the utility database, and connects the user.

```
dbisql -c  
"uid=DBA;pwd=asa;dbn=utility_db;eng=TestSrv"
```

Having executed the above statement successfully, the user connects to the utility database, and can execute file administration statements.

C H A P T E R 9

Automating Tasks Using Schedules and Events

About this chapter This chapter describes how to use scheduling and event handling features of Adaptive Server Anywhere to automate database administration and other tasks.

Contents

Topic	Page
Introduction	232
Understanding schedules	234
Understanding events	236
Understanding event handlers	240
Schedule and event internals	242
Scheduling and event handling tasks	244

Introduction

Many database administration tasks are best carried out systematically. For example, a regular backup procedure is an important part of proper database administration procedures.

You can automate routine tasks in Adaptive Server Anywhere by adding an **event** to a database, and providing a schedule for the event. Whenever one of the times in the schedule passes, the database server executes a sequence of actions called an **event handler**.

Database administration also requires taking action when certain conditions occur. For example, it may be appropriate to e-mail a notification to a system administrator when a disk containing the transaction log is filling up so that the administrator can handle the situation. These tasks too can be automated by defining event handlers for one of a set of **system events**.

Chapter contents

This chapter contains the following material:

- ◆ An introduction to scheduling and event handling (this section).
- ◆ Concepts and background information to help you design and use schedules and event handlers:
 - ◆ "Understanding schedules" on page 234.
 - ◆ "Understanding events" on page 236.
- ◆ A discussion of techniques for developing event handlers:
 - ◆ "Developing event handlers" on page 240.
- ◆ Internals information:
 - ◆ "Schedule and event internals" on page 242.
- ◆ Step by step instructions for how to carry out automation tasks.
 - ◆ "Scheduling and event handling tasks" on page 244.

Questions and answers

To answer the question...	Consider reading...
What is a schedule?	"Understanding schedules" on page 234.
What is a system event?	"Understanding events" on page 236
What is an event handler?	"Understanding event handlers" on page 240
How do I debug event handlers?	"Developing event handlers" on page 240

To answer the question...	Consider reading...
How does the database server use schedules to trigger event handlers?	"How the database server checks for scheduled times" on page 242
How can I schedule regular backups?	For an example, see "Understanding schedules" on page 234.
What kind of system events can the database server use to trigger event handlers?	"Understanding events" on page 236. "CREATE EVENT statement" on page 285 of the book <i>ASA SQL Reference Manual</i> .
What connection do event handlers get executed on?	"How event handlers are executed" on page 243.
How do event handlers get information about what triggered them?	"Developing event handlers" on page 240 "EVENT_PARAMETER function" on page 134 of the book <i>ASA SQL Reference Manual</i>

Understanding schedules

By scheduling activities you can ensure that a set of actions is executed at a set of preset times. The scheduling information and the event handler are both stored in the database itself.

You can define complex schedules by associating more than one schedule with a named event.

The following examples give some ideas for scheduled actions that may be useful.

When scheduling events, you can use either full-length English day names (Monday, Tuesday, and so on) or the abbreviated forms of the day (Mon, Tue, and so on). Note that you must use the full-length English day names if you want the day names to be recognized by a server running in a language other than English.

☞ For more information, see "CREATE EVENT statement" on page 285 of the book *ASA SQL Reference Manual*.

Examples

Carry out an incremental backup daily at 1:00 am:

```
create event IncrementalBackup
schedule
    start time '1:00 AM' every 24 hours
handler
begin
    backup database directory 'c:\\backup'
    transaction log only
    transaction log rename match
end
```

Summarize orders at the end of each business day:

```
create event Summarize
schedule
    start time '6:00 pm'
    on ( 'Monday', 'Tuesday', 'Wednesday', 'Thursday',
        'Friday' )
handler
begin
    insert into DBA.OrderSummary
    select max( date_ordered ),
           count( * ),
           sum( amount )
    from DBA.Orders
    where date_ordered = current date
end
```

Defining schedules

Schedule definitions have several components to them, to permit flexibility:

- ◆ **Name** Each schedule definition has a name. You can assign more than one schedule to a particular event, which can be useful in designing complex schedules.
- ◆ **Start time** You can define a start time for the event, which is the time that it is first executed.
- ◆ **Range** As an alternative to a start time, you can specify a range of times for which the event is active.
- ◆ **Recurrence** Each schedule can have a recurrence. The event is triggered on a frequency that can be given in hours, minutes, or seconds on a set of days that can be specified as days of the week or days of the month.

Understanding events

The database server tracks several kinds of system events. Event handlers are triggered when the system event is checked by the database server, and satisfies a provided **trigger condition**.

By defining event handlers to execute when a chosen system event occurs and satisfies a trigger condition that you define, you can improve the security and safety of your data, and help to ease administration.

🔗 For more information on the available system events, see "Choosing a system event" on page 236.

🔗 For more information on trigger conditions, see "Defining trigger conditions for events" on page 237.

Choosing a system event

Adaptive Server Anywhere tracks several system events. Each system event provides a hook on which you can hang a set of actions. The database server tracks the events for you, and executes the actions (as defined in the event handler) when needed.

The available system events include the following:

- ◆ **Backup** You can use the **BackupEnd** event type to take actions at the end of a backup.
- ◆ **DatabaseStart** The database is started.
- ◆ **Connection events** When a connection is made (**Connect**) or when a connection attempt fails (**ConnectFailed**). You may want to use these events for security purposes.
- ◆ **Free disk space** Tracks the available disk space on the device holding the database file (**DBDiskSpace**), the log file (**LogDiskSpace**), or temporary file (**TempDiskSpace**). This system event is not available on the following operating systems:
 - ◆ Windows 95 before OSR2
 - ◆ Windows CE

You may want to use disk space events to alert administrators in case of a disk space shortage.

- ◆ **File size** The file reaches a specified size. This can be used for the database file (**GrowDB**), the transaction log (**GrowLog**), or the temporary file (**GrowTemp**).

You may want to use file size events to track unusual actions on the database, or monitor bulk operations.

- ◆ **SQL errors** When an error is triggered, you can use the **RAISERROR** event type to take actions.
- ◆ **Idle time** The database server has been idle for a specified time. You may want to use this event type to carry out routine maintenance operations at quiet times.

Defining trigger conditions for events

Each event definition has a system event associated with it. It also has one or more trigger conditions. The event handler is triggered when the trigger conditions for the system event are satisfied.

The trigger conditions are included in the **WHERE** clause of the **CREATE EVENT** statement, and can be combined using the **AND** keyword. Each trigger condition is of the following form:

event_condition(*condition-name*) *comparison-operator value*

The *condition-name* argument is one of a set of preset strings, which are appropriate for different event types. For example, you can use **DBSize** (the database file size in Megabytes) to build a trigger condition suitable for the **GrowDB** system event. The database server does not check that the condition-name matches the event type: it is your responsibility to ensure that the condition is meaningful in the context of the event type.

Examples

- ◆ Limit the transaction log size to 10 Mb:

```
create event LogLimit
type GrowLog
where event_condition( 'LogSize' ) > 10
handler
begin
    backup database
    directory 'c:\\logs'
    transaction log only
    transaction log rename match
end
```

- ◆ Notify an administrator when free disk space on the device containing the database file falls below 10%, but do not execute the handler more than once every five minutes (300 seconds):

```
create event LowDBSpace
type DBDiskSpace
where event_condition( 'DBFreePercent' ) < 10
and event_condition( 'Interval' ) >= 300
handler
begin
    call xp_sendmail( recipient='DBAdmin',
                     subject='Low disk space',
                     "message"='Database free disk space '
                     || event_parameter( 'DBFreeSpace' ) );
end
```

- ◆ Notify an administrator of a possible attempt to break into the database:

```
create event SecurityCheck
type ConnectFailed
handler
begin
    declare num_failures int;
    declare mins int;

    insert into FailedConnections( log_time )
    values ( current timestamp );

    select count( * ) into num_failures
    from FailedConnections
    where log_time >= dateadd( minute, -5,
                             current timestamp );

    if( num_failures >= 3 ) then
        select datediff( minute, last_notification,
                        current timestamp ) into mins
        from Notification;

        if( mins > 30 ) then
            update Notification
            set last_notification = current timestamp;

            call xp_sendmail( recipient='DBAdmin',
                             subject='Security Check',
                             "message"=
                             'over 3 failed connections in last 5 minutes' )
        end if
    end if
end
```

- ◆ Run a process when the server has been idle for ten minutes. Do not execute more frequently than once per hour:


```
create event Soak
type ServerIdle
where event_condition( 'IdleTime' ) >= 600
and event_condition( 'Interval' ) >= 3600
handler
begin
    message ' Insert your code here ... '
end
```

Understanding event handlers

Event handlers execute on a separate connection from the action that triggered the event, and so do not interact with client applications. They execute with the permissions of the creator of the event.

Developing event handlers

Event handlers, whether for scheduled events or for system event handling, contain compound statements, and are similar in many ways to stored procedures. You can add loops, conditional execution, and so on, and you can use the Adaptive Server Anywhere debugger to debug event handlers.

Context information for event handlers

One difference between event handlers and stored procedures is that event handlers do not take any arguments. Certain information about the context in which an event was triggered is available through the **event_parameter** function, which supplies information about the connection that caused an event to be triggered (connection ID, user ID), as well as the event name and the number of times it has been executed.

🔗 For more information, see "EVENT_PARAMETER function" on page 134 of the book *ASA SQL Reference Manual*.

Testing event handlers

During development, you want event handlers to be triggered at convenient times. You can use the TRIGGER EVENT statement to explicitly cause an event to execute, even when the trigger condition or scheduled time has not occurred. However, TRIGGER EVENT does not cause disabled event handlers to be executed.

🔗 For more information, see "TRIGGER EVENT statement" on page 566 of the book *ASA SQL Reference Manual*.

While it is not good practice to develop event handlers on a production database, you can disable event handlers from Sybase Central or explicitly using the ALTER EVENT statement.

It can be useful to use a single set of actions to handle multiple events. For example, you may want to take a notification action if disk space is limited on any of the devices holding the database or log files. To do this, create a stored procedure and call it in the body of each event handler.

Debugging event handlers

Debugging event handlers is very similar to debugging stored procedures. The event handlers appear in the procedures list.

One difference is that, because each event handler runs on its own connection, you must be sure to select **All connections** before setting a breakpoint in an event handler.

You can also determine how many instances of a particular event handler are currently active using the *NumActive* event parameter. This function is useful if you want to limit an event handler so that only one instance executes at any given time.

🔗 For more information about the *NumActive* event parameter, see "EVENT_PARAMETER function" on page 134 of the book *ASA SQL Reference Manual*.

🔗 For more information about step-by-step instructions, see "Debugging an event handler" on page 246.

Schedule and event internals

This section describes how the database server processes schedules and event definitions.

How the database server checks for events

Events are classified according to their **event type**, as specified directly in the CREATE EVENT statement or using Sybase Central. There are two kinds of event types:

- ◆ **Active event types** Some event types are the result of action by the database server itself. These active event types include growing database files, or the start and end of different database actions (**BackupEnd** and so on) or RAISERROR.

When the database server takes the action, it checks to see whether the trigger conditions defined in the WHERE clause are satisfied, and if so, triggers any events defined for that event type.

- ◆ **Polled event types** Some event types are not triggered solely by database actions. The free disk space types (**DBDiskSpace** and so on) as well as the **IdleTime** types are of this kind.

For these types of events, the database server polls every thirty seconds, starting approximately thirty seconds after the database server is started.

For the **IdleTime** event type, the database server checks whether the server has been idle for the entire thirty seconds. If no requests have started and none are currently active, it adds the idle check interval time in seconds to the idle time total; otherwise, the idle time total is reset to 0. The value for **IdleTime** is therefore always a multiple of thirty seconds. When **IdleTime** is greater than the interval specified in the trigger condition, event handlers associated with **IdleTime** are fired.

How the database server checks for scheduled times

The calculation of scheduled event times is done when the database server starts, and each time a scheduled event handler completes.

The calculation of the next scheduled time is based on the increment specified in the schedule definition, with the increment being added to the previous start time. If the event handler takes longer to execute than the specified increment, so that the next time is earlier than the current time, the database server increments until the next scheduled time is in the future.

An event handler that takes sixty-five minutes to execute and is requested to run every hour between 9:00 and 5:00 will run every two hours, at 9:00, 11:00, 1:00, and so on.

To run a process such that it operates between 9:00 and 5:00 and delays for some period before the next execution, you could define a handler to loop until its completion time has passed, with a sleep instruction, perhaps using `xp_cmdshell`, between each iteration.

If you are running a database server intermittently, and it is not running at a scheduled time, the event handler does not run at startup. Instead, the next scheduled time is computed at startup. If, for example, you schedule a backup to take place every night at one o'clock, but regularly shut down the database server at the end of each work day, the backup never takes place.

How event handlers are executed

When an event handler is triggered, a temporary internal connection is made on which the event handler is executed. The handler is *not* executed on the connection that caused the handler to be triggered, and consequently statements such as `MESSAGE .. TO CLIENT`, which interact with the client application, are not meaningful within event handlers.

The temporary connection on which the handler is executed does not count towards the connection limit for licensing purposes.

Event creation requires DBA authority, and events execute with the permissions of their creator. If you wish event handlers to execute with non-DBA authority, you can call a procedure from within the handler, as stored procedures run with the permissions of their creator.


Any event errors are logged to the server console.

Scheduling and event handling tasks

This section collects together instructions for tasks related to automating tasks with schedules and events.

Adding a schedule or event to a database

Schedules and events are handled in a similar fashion, both from Sybase Central and in SQL.

 For more information, see "Understanding schedules" on page 234, and "Understanding events" on page 236.

❖ To add a schedule or event to a database (Sybase Central):


- 1 Connect to the database as a user with DBA authority.
- 2 Open the Events folder for your database.
- 3 Double-click Add Event. The Event Creation wizard appears.
- 4 Follow the instructions in the wizard.

The wizard contains many options, depending on the schedule or event you wish to create. These are explained in detail in other tasks.

❖ To add a schedule or event to a database (SQL):

- 1 Connect to the database as a user with DBA authority.
- 2 Execute a CREATE EVENT statement.

The CREATE EVENT statement contains many options, depending on the schedule or event you wish to create. These are explained in detail in other tasks.

 For more information, see "CREATE EVENT statement" on page 285 of the book *ASA SQL Reference Manual*.

Adding a manually-triggered event to a database

If you create an event handler without a schedule or system event to trigger it, it is executed only when manually triggered.

❖ To add a manually-triggered event to a database (Sybase Central):

- 1 Connect to the database as a user with DBA authority.

- 2 Open the Events folder for your database.
- 3 Double-click Add Event. The Event Creation wizard appears.
- 4 Type a name for the event, and click Next.
- 5 Select Manually, and click Next.
- 6 Type the SQL statements for your event handler, and click Next.
- 7 Select Enable this event, and select Execute at all Locations, and click Next.
- 8 Type a comment describing the event, and click Finish to add the event to the database.

If you wish to accept the default values for all remaining options, you can click Finish at an earlier stage in the wizard.

❖ **To add a manually-triggered event to a database (SQL):**

- 1 Connect to the database as a user with DBA authority.
- 2 Execute a CREATE EVENT statement with no schedule or WHERE clause. The restricted syntax of the CREATE EVENT is as follows:

```
CREATE EVENT event-name  
HANDLER  
BEGIN  
    ... event handler  
END
```

If you are developing event handlers, you can add schedules or system events to control the triggering of an event later, either using Sybase Central or the ALTER EVENT statement.

 See also:

- ◆ For information on triggering events, see "Triggering an event handler" on page 246.
- ◆ For information on altering events, see "ALTER EVENT statement" on page 211 of the book *ASA SQL Reference Manual*.

Triggering an event handler

Any event handler can be triggered manually, in addition to those occasions when it executes because of a schedule or system event. Triggering events manually can be useful during development of event handlers, and also, for certain events, in production environments. For example, you may have a monthly sales report scheduled, but from time to time you may want to obtain a sales report for a reason other than the end of the month.

☞ For more information on developing event handlers, see "Developing event handlers" on page 240.

❖ To trigger an event handler (Sybase Central):

- 1 Connect to the database as a user with DBA authority.
- 2 Open the Events folder for your database.
- 3 Right-click the event you wish to trigger and choose Trigger from the popup menu. The Trigger Event dialog appears.

The event must be enabled before you can select Trigger from the popup menu. You can enable the event on the General tab of the Event property sheet.

- 4 Supply any parameters the event handler requires, in the following form:

parameter=value;parameter=value

Click OK to trigger the event handler.

❖ To trigger an event handler (SQL):

- 1 Connect to the database as a user with DBA authority.
- 2 Execute the TRIGGER EVENT statement, supplying the name of the event. For example:

```
TRIGGER EVENT sales_report_event
```

☞ For more information, see "TRIGGER EVENT statement" on page 566 of the book *ASA SQL Reference Manual*.

Debugging an event handler

Debugging is a regular part of any software development. Event handlers can be debugged during the development process.

☞ For more information on developing event handlers, see "Developing event handlers" on page 240.

☞ For more information on using the debugger, see "Debugging Logic in the Database" on page 571 of the book *ASA SQL User's Guide*.

❖ **To debug an event handler:**

- 1 Start the Adaptive Server Anywhere debugger.
From the Start menu, choose Programs►Sybase
SQL Anywhere 8►Adaptive Server Anywhere 8►Debug Database
Objects.
- 2 In the Connections window, double-click All Connections.
- 3 In the Procedures window, double-click the event you wish to debug.
The event definition appears in the Source window.
- 4 In the Source window, set a breakpoint.
- 5 From Interactive SQL or another application, trigger the event handler
using the TRIGGER EVENT statement.
- 6 The execution stops at the breakpoint you have set. You can now use the
debugger features to trace execution, local variables, and so on.

CHAPTER 10

International Languages and Character Sets

About this chapter

This chapter describes how to configure your Adaptive Server Anywhere installation to handle international language issues.

Contents

Topic	Page
Introduction to international languages and character sets	250
Understanding character sets in software	253
Understanding locales	262
Understanding collations	269
Understanding character set translation	278
Collation internals	282
International language and character set tasks	287

Introduction to international languages and character sets

This section provides an introduction to the issues you may face when working in an environment that uses more than one character set, or when using languages other than English.

When you create a database, you specify a collating sequence, or collation, for the database to use. A collation is a combination of a character set and a sort order for characters in the database.

Adaptive Server Anywhere international features

Adaptive Server Anywhere provides two sets of features that are of particular interest when setting up databases for languages.

- ◆ **Collations** You can choose from a wide selection of supplied collations when you create a database. By creating your database with the proper collation, you ensure proper sorting of data.

Whenever the database compares strings, sorts strings, or carries out other string operations such as case conversion, it does so using the collation sequence. The database carries out sorting and string comparison when statements such as the following are executed:

- ◆ Queries with an ORDER BY clause.
- ◆ Expressions that use string functions, such as LOCATE, SIMILAR, SOUNDEX.
- ◆ Conditions using the LIKE keyword.
- ◆ Queries that use index lookups on character data.


The database also uses collations to identify valid or unique identifiers (column names and so on).

- ◆ **Character set translation** You can set up Adaptive Server Anywhere to convert data between the character set encoding on your server and client systems, thus maintaining the integrity of your data even in mixed character set environments.

Character set translation is provided between client and server, as well as by the ODBC driver. The Adaptive Server Anywhere ODBC driver provides OEM to ANSI character set translation and Unicode support.

Using the default collation

If you use the default actions when creating a database, the Initialization utility infers a collation from the character set used by the operating system on the machine at which you create the database.

 For more information on how to find the default collation in your environment, see "Finding the default collation" on page 287.

If all the machines in your environment share the same character set, then this choice of collation ensures that all the character data in the database and in client applications is represented in the same manner. As long as the collation provides a proper character set and sort order for your data, using this default setting is a simple way of ensuring that characters are represented consistently throughout the system.

If it is not possible to set up your system in this default manner, you need to decide which collation to use in your database, and whether to use character set translation to ensure that data is exchanged consistently between the pieces of your database system. This chapter provides the information you need to make and implement these decisions.

Character set questions and answers

The following table identifies where you can find answers to questions.

To answer the question...	Consider reading...
How do I set up my computing environment to treat character sets properly?	"Configuring your character set environment" on page 287
How do I decide which collation to use for my database?	"Understanding collations" on page 269
How are characters represented in software, and Adaptive Server Anywhere in particular?	"Understanding character sets in software" on page 253
What collations does Adaptive Server Anywhere provide?	"Supplied collations" on page 269
How do I ensure that error and informational messages sent from the database server to client applications are sent in the proper language and character set for my application?	"Character translation for database messages" on page 278

To answer the question...	Consider reading...
I have a different character set on client machines from that in use in the database. How can I get characters to be exchanged properly between client and server?	"Starting a database server using character set translation" on page 291
What character sets can I use for connection strings?	"Connection strings and character sets" on page 279
How do I create a collation that is different from the supplied ones?	"Creating a database with a custom collation" on page 293
How do I change the collation sequence of an existing database?	"Changing a database from one collation to another" on page 294.
How do I create a database for Windows CE?	"Creating databases for Windows CE" on page 273

Understanding character sets in software

This section provides general information about software issues related to international languages and character sets.

Pieces in the character set puzzle

There are several distinct aspects to character storage and display by computer software:

- ◆ Each piece of software works with a character set. A **character set** is a set of symbols, including letters, digits, spaces, and other symbols.
- ◆ To handle these characters, each piece of software employs a character set **encoding**, in which each character is mapped onto one or more bytes of information, typically represented as hexadecimal numbers. This encoding is also called a **code page**.
- ◆ Database servers, which sort characters (for example, list names alphabetically), use a collation. A **collation** is a combination of a character encoding (a map between characters and hexadecimal numbers) and a **sort order** for the characters. There may be more than one sort order for each character set; for example, a case-sensitive order and a case-insensitive order, or two languages may sort characters in a different order.
- ◆ Characters are printed or displayed on a screen using a **font**, which is a mapping between characters in the character set and its appearance. Fonts are handled by the operating system.
- ◆ Operating systems also use a **keyboard mapping** to map keys or key combinations on the keyboard to characters in the character set.

Language issues in client/server computing

Database users working at client applications may see or access strings from the following sources:

- ◆ **Data in the database** Strings and other text data are stored in the database. The database server processes these strings when responding to requests.

For example, the database server may be asked to supply all the last names beginning with a letter ordered less than N in a table. This request requires string comparisons to be carried out, and assumes a character set ordering.

The database server receives strings from client applications as streams of bytes. It associates these bytes with characters according to the database character set. If the data is held in an indexed column, the index is sorted according to the sort order of the collation.

- ◆ **Database server software messages** Applications can cause database errors to be generated. For example, an application may submit a query that references a column that does not exist. In this case, the database server returns a warning or error message. This message is held in a **language resource library**, which is a DLL or shared library called by Adaptive Server Anywhere.
- ◆ **Client application** The client application interface displays text, and internally the client application may process text.
- ◆ **Client software messages** The client library uses the same language library as the database server to provide messages to the client application.
- ◆ **Operating system** The client operating system has text displayed on its interface, and may also process text.

For a satisfactory working environment, all these sources of text must work together. Loosely speaking, they must all be working in the user's language and/or character set.

Code pages

Many languages have few enough characters to be represented in a single-byte character set. In such a character set, each character is represented by a single byte: a two-digit hexadecimal number.

At most, 256 characters can be represented in a single byte. No single-byte character set can hold all of the characters used internationally, including accented characters. This problem was addressed by the development of a set of code pages, each of which describes a set of characters appropriate for one or more national languages. For example, code page 869 contains the Greek character set, and code page 850 contains an international character set suitable for representing many characters in a variety of languages.

Upper and lower
pages

With few exceptions, characters 0 to 127 are the same for all the single-byte code pages. The mapping for this range of characters is called the **ASCII** character set. It includes the English language alphabet in upper and lower case, as well as common punctuation symbols and the digits. This range is often called the **seven-bit** range (because only seven bits are needed to represent the numbers up to 127) or the **lower** page. The characters from 128 to 255 are called **extended characters**, or **upper** code-page characters, and vary from code page to code page.

Problems with code page compatibility are rare if the only characters used are from the English alphabet, as these are represented in the ASCII portion of each code page (0 to 127). However, if other characters are used, as is generally the case in any non-English environment, there can be problems if the database and the application use different code pages.

Example

Suppose a database holding French language strings uses code page 850, and the client operating system uses code page 437. The character Å (upper case A grave) is held in the database as character \xB7 (decimal value 183). In code page 437, character \xB7 is a graphical character. The client application receives this byte and the operating system displays it on the screen, the user sees a graphical character instead of an A grave.

ANSI and OEM code pages in Windows

For PC users, the issue is complicated because there are at least two code pages in use on most PCs. Character-mode applications (those using the console or command prompt window) in Windows 95/98/Me and Windows NT/200/XP, use code pages taken from the IBM set. These are called **OEM code pages** (Original Equipment Manufacturer) for historical reasons.

Windows operating systems do not require the line drawing characters that were held in the extended characters of the OEM code pages, so they use a different set of code pages. These pages are based on the ANSI standard and are therefore commonly called **ANSI code pages**.

Adaptive Server Anywhere supports collations based on both OEM and ANSI code pages.

Example

Consider the following situation:

- ◆ A PC is running the Windows 95 operating system with ANSI code page 1252.
- ◆ The code page for character-mode applications is OEM code page 437.
- ◆ Text is held in a database created using the collation corresponding to OEM code page 850.

An upper case A grave in the database is stored as character 183. This value appears as a graphical character in a character-mode application. The same character appears as a dot in a Windows application.

☞ For more information about choosing a single-byte collation for your database, see "Understanding collations" on page 269.

Supported code pages

The following table lists supported code pages.

Code Page	Description
1252	Windows Latin-1
037	USA, Canada (Bilingual, French), Netherlands, Portugal, Brazil, Australia
273	IBM Austria, Germany
277	IBM Denmark, Norway
278	IBM Finland, Sweden
280	IBM Italy
284	IBM Catalan/Spain, Spanish Latin America
285	IBM United Kingdom, Ireland
297	IBM France
420	IBM Arabic
424	IBM Hebrew
437	MS-DOS United States, Australia, New Zealand, South Africa
500	EBCDIC 500V1
737	PC Greek
775	PC Baltic
838	IBM Thailand extended SBCS
850	MS-DOS Latin-1
852	MS-DOS Latin-2
855	IBM Cyrillic
856	IBM Hebrew
857	IBM Turkish

Code Page	Description
858	Variant of Cp850 with Euro character
860	MS-DOS Portuguese
861	MS-DOS Icelandic
862	PC Hebrew
863	MS-DOS Canadian French
864	PC Arabic
865	MS-DOS Nordic
866	MS-DOS Russian
868	MS-DOS Pakistan
869	IBM Modern Greek
870	IBM Multilingual Latin-2
871	IBM Iceland
874	IBM Thai
875	IBM Greek
918	IBM Pakistan (Urdu)
921	IBM Latvia, Lithuania (AIX, DOS)
922	IBM Estonia (AIX, DOS)
930	Japanese Katakana-Kanji mixed with 4370 UDC, superset of 5026
933	Korean Mixed with 1880 UDC, superset of 5029
935	Simplified Chinese Host mixed with 1880 UDC, superset of 5031
937	Traditional Chinese Host mixed with 6204 UDC, superset of 5033
939	Japanese Latin Kanji mixed with 4370 UDC, superset of 5035
942	IBM OS/2 Japanese, superset of Cp932
942	C Variant of Cp942
943	IBM OS/2 Japanese, superset of Cp932 and Shift-JIS
943	C Variant of Cp943


Code Page	Description
948	OS/2 Chinese (Taiwan) superset of 938
949	PC Korean
949	C Variant of Cp949
950	PC Chinese (Hong Kong, Taiwan)
964	AIX Chinese (Taiwan)
970	AIX Korean
1006	IBM AIX Pakistan (Urdu)
1025	IBM Multilingual Cyrillic: Bulgaria, Bosnia, Herzegovina, Macedonia (FYR)
1026	IBM Latin-5, Turkey
1046	IBM Arabic - Windows
1097	IBM Iran (Farsi)/Persian
1098	IBM Iran (Farsi)/Persian (PC)
1112	IBM Latvia, Lithuania
1122	IBM Estonia
1123	IBM Ukraine
1124	IBM AIX Ukraine
1140	Variant of Cp037 with Euro character
1141	Variant of Cp273 with Euro character
1142	Variant of Cp277 with Euro character
1143	Variant of Cp278 with Euro character
1144	Variant of Cp280 with Euro character
1145	Variant of Cp284 with Euro character
1146	Variant of Cp285 with Euro character
1147	Variant of Cp297 with Euro character
1148	Variant of Cp500 with Euro character
1149	Variant of Cp871 with Euro character
1250	Windows Eastern European
1251	Windows Cyrillic

Code Page	Description
1253	Windows Greek
1254	Windows Turkish
1255	Windows Hebrew
1256	Windows Arabic
1257	Windows Baltic
1258	Windows Vietnamese
1381	IBM OS/2, DOS People's Republic of China (PRC)
1383	IBM AIX People's Republic of China (PRC)
33722	IBM-eucJP - Japanese (superset of 5050)

Multibyte character sets

Some languages, such as Japanese and Chinese, have many more than 256 characters. These characters cannot all be represented using a single byte, but can be represented in multibyte character sets. In addition, some character sets use the much larger number of characters available in a multibyte representation to represent characters from many languages in a single, more comprehensive, character set.

Multibyte character sets are of two types. Some are **variable width**, in which some characters are single-byte characters, others are double-byte, and so on. Other sets are **fixed width**, in which all characters in the set have the same number of bytes. Adaptive Server Anywhere supports only variable-width character sets.

 For more information on the multibyte character sets, see "Using multibyte collations" on page 277.

Example

As an example, characters in the Shift-JIS character set are either one or two bytes in length. If the value of the first byte is in the range of hexadecimal values from \x81 to \x9F or from \xE0 to \xEF (decimal values 129-159 or 224-239) the character is a two-byte character and the subsequent byte (called a follow byte) completes the character. A **follow byte** is any byte(s) other than the first byte.

If the first byte is outside this range, the character is a single-byte character and the next byte is the first byte of the following character.

- ◆ The properties of any Shift-JIS character can be read from its first byte also. Characters with a first byte in the range \x09 to \x0D, or \x20 are space characters.
- ◆ Characters in the ranges \x41 to \x5A, \x61 to \x7A, \x81 to \x9F or \xE0 to \xEF are considered to be alphabetic (letters).
- ◆ Characters in the range \x30 to \x39 are digits.

Sorting characters using collations

The database collation sequence includes the notion of alphabetic ordering of letters, and extends it to include all characters in the character set, including digits and space characters.

Associating more than one character with each sort position

More than one character can be associated with each sort position. This is useful if you wish, for example, to treat an accented character the same as the character without an accent.

Two characters with the same sort position are considered identical in all ways by the database. Therefore, if a collation assigned the characters *a* and *e* to the same sort position, then a query with the following search condition:

```
WHERE coll = 'want'
```

is satisfied by a row for which **coll** contains the entry **went**.

At each sort position, lowercase and uppercase forms of a character can be indicated. For case-sensitive databases, the lowercase and uppercase characters are not treated as equivalent. For case-insensitive databases, the lowercase and uppercase versions of the character are considered equivalent.

Tip

Any code that selects a default collation for a German system should select 1252LATIN1, *not* 1252DEU. 1252DEU differentiates between characters with and without an umlaut, while 1252LATIN1 does not. 1252LATIN1 considers Muller and Müller equal, but 1252DEU does not consider them equal. Because 1252DEU views characters with umlauts as separate characters, it has the following alphabetic ordering: ob, öa.

First-byte collation orderings for multibyte character sets

A sorting order for characters in a multibyte character set can be specified only for the first byte. Characters that have the same first byte are sorted according to the hexadecimal value of the following bytes.

International aspects of case sensitivity

Adaptive Server Anywhere is always **case preserving** and **case insensitive** for identifiers, such as table names and column names. This means that the names are stored in the case in which they are created, but any access to the identifiers is done in a case-insensitive manner.

For example, the names of the system tables are held in upper case (SYSDOMAIN, SYSTABLE, and so on), but access is case insensitive, so that the two following statements are equivalent:

```
SELECT *  
FROM systable  
  
SELECT *  
FROM SYSTABLE
```

The equivalence of upper and lower case characters is enforced in the collation. There are some collations where particular care may be needed when assuming case insensitivity of identifiers.

Example

In the Turkish 857TRK collation, the lower case **i** does not have the character **I** as its upper case equivalent. Therefore, despite the case insensitivity of identifiers, the following two statements are *not* equivalent in this collation:

```
SELECT *  
FROM sysdomain  
  
SELECT *  
FROM SYSDOMAIN
```

Understanding locales

Both the database server and the client library recognize their language and character set environment using a **locale definition**.

Introduction to locales

The application locale, or client locale, is used by the client library when making requests to the database server, to determine the character set in which results should be returned. If character set translation is enabled (the default), the database server compares its own locale with the application locale to determine whether character set translation is needed. Different databases on a server may have different locale definitions.

The locale consists of the following components:

- ◆ **Language** The language is a two-character string using the ISO-639 standard values: DE for German, FR for French, and so on. Both the database server and the client have language values for their locale.

The database server uses the locale language to determine the following behavior:

- ◆ Which language library to load.
- ◆ The language is used together with the character set to determine which collation to use when creating databases, if no collation is explicitly specified.

The client library uses the locale language to determine the following behavior:

- ◆ Which language library to load.
- ◆ Which language to request from the database.

🔗 For more information, see "Understanding the locale language" on page 263.

- ◆ **Character set** The character set is the code page in use. The client and server both have character set values, and they may differ. If they differ, character set translation may be required to enable interoperability.

For machines that use both OEM and ANSI code pages, the ANSI code page is the value used here.

🔗 For more information, see "Understanding the locale character set" on page 264.

- ◆ **Collation label** The collation label is the Adaptive Server Anywhere collation. The client side does not use a collation label. Different databases on a database server may have different collation labels.

🔗 For more information, see "Understanding the locale collation label" on page 267.

Understanding the locale language

The locale language is an indicator of the language being used by the user of the client application, or expected to be used by users of the database server.

🔗 For more information about how to find locale settings, see "Determining locale information" on page 288.

The client library determines the language component of the locale as follows:

- 1 On Windows, it checks the Adaptive Server Anywhere language registry entry, as described in "Registry settings on installation" on page 214.
- 2 On other operating systems, or if the registry setting is not present, it checks the operating system language setting.

The database server determines the language component of the locale as follows:

- 1 It checks the SQLLOCALE environment variable, if it exists.

🔗 For more information, see "Setting the SQLLOCALE environment variable" on page 268.

- 2 On Windows, it checks the Adaptive Server Anywhere language registry entry, as described in "Registry settings on installation" on page 214.
- 3 On other operating systems, or if the registry setting is not present, it checks the operating system language setting.

Language label values


The following table displays the valid language label values, together with the equivalent ISO 639 labels:

Language label	Alternative label	ISO_639 language code
chinese	simpchin	ZH
danish	N/A	DA
french	N/A	FR
german	N/A	DE
italian	N/A	IT


Language label	Alternative label	ISO_639 language code
japanese	N/A	JA
korean	N/A	KO
norwegian	norweg	NO
polish	N/A	PL
portuguese	portugue	PT
russian	N/A	RU
spanish	N/A	ES
swedish	N/A	SV
tchinese	tradchin	TW
ukrainian	N/A	UK
us_english	english	EN

Understanding the locale character set


Both application and server locale definitions have a character set. The application uses its character set when requesting character strings from the server. If character set translation is enabled (the default), the database server compares its character set with that of the application to determine whether character set translation is needed.

 For more information about how to find locale settings, see "Determining locale information" on page 288.

The client library determines the character set as follows:

- 1 If the connection string specifies a character set, it is used.
 For more information, see "CharSet connection parameter" on page 167.
- 2 Open Client applications check the *locales.dat* file in the Sybase *locales* directory.
- 3 Character set information from the operating system is used to determine the locale:
 - ◆ On Windows operating systems, use the GetACP system call. This returns the ANSI character set, not the OEM character set.
 - ◆ On UNIX, default to ISO8859-1.
 - ◆ On other platforms, use code page 1252.

The database server determines the character set for a connection as follows:

- 1 The character set specified by the client is used if it is supported.
 For more information, see "CharSet connection parameter" on page 167.
- 2 The database's character set is used if the client specifies a character set that is not supported.

When a new database is created, the database server determines the character set for the new database as follows.

- 1 A collation is specified in the CREATE DATABASE statement or on with the *dbinit* utility.
- 2 The SQLLOCALE environment variable is used if it exists.
- 3 Character set information from the operating system is used to determine the locale.
 - ◆ On Windows operating systems, use the GetACP system call. This returns the ANSI character set, not the OEM character set.
 - ◆ On UNIX, default to ISO8859-1.
 - ◆ On other platforms, use code page 1252.

The locale character set and language are used to determine which collation to use when creating a database if none is explicitly specified.

Character set labels

The following table displays the valid character set label values, together with the equivalent IANA labels and a description:


Character set label	IANA label	Description
big5	<N/A>	Traditional Chinese (cf. CP950)
cp437	<N/A>	IBM CP437 - U.S. code set
cp850	<N/A>	IBM CP850 - European code set
cp852	<N/A>	PC Eastern Europe
cp855	<N/A>	IBM PC Cyrillic
cp856	<N/A>	Alternate Hebrew
cp857	<N/A>	IBM PC Turkish
cp860	<N/A>	PC Portuguese
cp861	<N/A>	PC Icelandic
cp862	<N/A>	PC Hebrew
cp863	<N/A>	IBM PC Canadian French code page

Character set label	IANA label	Description
cp864	<N/A>	PC Arabic
cp865	<N/A>	PC Nordic
cp866	<N/A>	PC Russian
cp869	<N/A>	IBM PC Greek
cp874	<N/A>	Microsoft Thai SB code page
cp932	windows-31j	Microsoft CP932 = Win31J-DBCS
cp936	<N/A>	Simplified Chinese
cp949	<N/A>	Korean
cp950	<N/A>	PC (MS) Traditional Chinese
cp1250	<N/A>	MS Windows Eastern European
cp1251	<N/A>	MS Windows Cyrillic
cp1252	<N/A>	MS Windows US (ANSI)
cp1253	<N/A>	MS Windows Greek
cp1254	<N/A>	MS Windows Turkish
cp1255	<N/A>	MS Windows Hebrew
cp1256	<N/A>	MS Windows Arabic
cp1257	<N/A>	MS Windows Baltic
cp1258	<N/A>	MS Windows Vietnamese
deckanji	<N/A>	DEC UNIX JIS encoding
euccns	<N/A>	EUC CNS encoding: Traditional Chinese with extensions
eucgb	<N/A>	EUC GB encoding = Simplified Chinese
eucjis	euc-jp	Sun EUC JIS encoding
eucksc	<N/A>	EUC KSC Korean encoding (cf. CP949)
greek8	<N/A>	HP Greek-8
iso_1	iso_8859-1:1987	ISO 8859-1 Latin-1
iso15	<N/A>	ISO 8859-15 Latin1 with Euro, etc.
iso88592	iso_8859-2:1987	ISO 8859-2 Latin-2 Eastern Europe
iso88595	iso_8859-5:1988	ISO 8859-5 Latin/Cyrillic
iso88596	iso_8859-6:1987	ISO 8859-6 Latin/Arabic

Character set label	IANA label	Description
iso88597	iso_8859-7:1987	ISO 8859-7 Latin/Greek
iso88598	iso_8859-8:1988	ISO 8859-8 Latin/Hebrew
iso88599	iso_8859-9:1989	ISO 8859-9 Latin-5 Turkish
koi8	<N/A>	KOI-8 Cyrillic
mac	macintosh	Standard Mac coding
mac_cyr	<N/A>	Macintosh Cyrillic
mac_ee	<N/A>	Macintosh Eastern European
macgrk2	<N/A>	Macintosh Greek
macturk	<N/A>	Macintosh Turkish
roman8	hp-rpman8	HP Roman-8
sjis	shift_jis	Shift JIS (no extensions)
tis620	<N/A>	TIS-620 Thai standard
turkish8	<N/A>	HP Turkish-8
utf8	utf-8	UTF-8 treated as a character set


Understanding the locale collation label

Each database has its own collation. The database server uses the collation label taken from the locale definition to determine which code page to use when initializing a database.

 For more information about how to find locale settings, see "Determining locale information" on page 288.

The database server determines the collation label as follows:

- 1 It checks the `SQLLOCALE` environment variable, if it exists.

 For more information, see "Setting the `SQLLOCALE` environment variable" on page 268.

- 2 It uses an internal table to find a collation label corresponding to the language and character set.

Collation label values


The collation label is a label for one of the supplied Adaptive Server Anywhere collations, as listed in "Understanding collations" on page 269.

Setting the SQLLOCALE environment variable

The SQLLOCALE environment variable is a single string that consists of three semi-colon-separated assignments. It has the following form:

```
CS=cslabel;LANG=langlabel;LABEL=colabel
```

where *cslabel* is a character set label from the list in "Character set labels" on page 265, *langlabel* is a language label from the list in "Language label values" on page 263, and *colabel* is taken from the list in "Understanding collations" on page 269, or is a custom collation label.

 For more information about how to set environment variables, see "Setting environment variables" on page 208.

Understanding collations

This section describes the supplied collations, and provides suggestions as to which collations to use under certain circumstances.

For more information about how to create a database with a specific collation, see "Creating a database with a named collation" on page 290 and "The Initialization utility" on page 465.

For information on changing a database from one collation to another, see "Changing a database from one collation to another" on page 294.

Supplied collations

The following collations are supplied with Adaptive Server Anywhere. You can obtain this list by typing the following command at the command prompt:

```
dbinit -l+
```

Collation label	Type	Description
437LATIN1	OEM	Code Page 437, Latin 1, Western
437ESP	OEM	Code Page 437, Spanish
437SVE	OEM	Code Page 437, Swedish/Finnish
819CYR	ANSI	Code Page 819, Cyrillic
819DAN	ANSI	Code Page 819, Danish
819ELL	ANSI	Code Page 819, Greek
819ESP	ANSI	Code Page 819, Spanish
819ISL	ANSI	Code Page 819, Icelandic
819LATIN1	ANSI	Code Page 819, Latin 1, Western
819LATIN2	ANSI	Code Page 819, Latin 2, Central/Eastern European
819NOR	ANSI	Code Page 819, Norwegian
819RUS	ANSI	Code Page 819, Russian
819SVE	ANSI	Code Page 819, Swedish/Finnish
819TRK	ANSI	Code Page 819, Turkish
850CYR	OEM	Code Page 850, Cyrillic, Western
850DAN	OEM	Code Page 850, Danish

Collation label	Type	Description
850ELL	OEM	Code Page 850, Greek
850ESP	OEM	Code Page 850, Spanish
850ISL	OEM	Code Page 850, Icelandic
850LATIN1	OEM	Code Page 850, Latin 1, Western
850LATIN2	OEM	Code Page 850, Latin 2, Central/Eastern European
850NOR	OEM	Code Page 850, Norwegian
850RUS	OEM	Code Page 850, Russian
850SVE	OEM	Code Page 850, Swedish/Finnish
850TRK	OEM	Code Page 850, Turkish
852LATIN2	OEM	Code Page 852, Latin 2, Central/Eastern European
852CYR	OEM	Code Page 852, Cyrillic
852POL	OEM	Code Page 852, Polish
855CYR	OEM	Code Page 855, Cyrillic
856HEB	OEM	Code Page 856, Hebrew
857TRK	OEM	Code Page 857, Turkish
860LATIN1	OEM	Code Page 860, Latin 1, Western
861ISL	OEM	Code Page 861, Icelandic
862HEB	OEM	Code Page 862, Hebrew
863LATIN1	OEM	Code Page 863, Latin 1, Western
865NOR	OEM	Code Page 865, Norwegian
866RUS	OEM	Code Page 866, Russian
869ELL	OEM	Code Page 869, Greek
920TRK	ANSI	Code Page 920, Turkish, ISO-8859-9
932JPN	Multibyte	Code Page 932, Japanese Shift-JIS with Microsoft extensions
936ZHO	Multibyte	Code Page 936, Simplified Chinese, PRC GBK 2312-80 8-bit encoding
949KOR	Multibyte	Code Page 949, Korean KS C 5601-1987 encoding, Wansung
950TWN	Multibyte	Code Page 950, Traditional Chinese, Big 5 Encoding
1250LATIN2	ANSI	Code Page 1250, Windows Latin 2,

Collation label	Type	Description
		Central/Eastern European
1250POL	ANSI	Code Page 1250, Windows Latin 2, Polish
1251CYR	ANSI	Code Page 1251, Cyrillic
1252DEU	ANSI	Code Page 1251, Windows Specialty German, Umlaut chars not equal
1252LATIN1	ANSI	Code Page 1252, Windows Latin 1, Western
1254TRK	ANSI	Code Page 1254, Windows Latin 5, Turkish, ISO 8859-9 with extensions
EUC_JAPAN	Multibyte	Japanese EUC JIS X 0208-1990 and JIS X 0212-1990 Encoding
EUC_CHINA	Multibyte	Simplified Chinese GB 2312-80 Encoding
EUC_TAIWAN	Multibyte	Taiwanese Big 5 Encoding
EUC_KOREA	Multibyte	Korean KS C 5601-1992 Encoding, Johab
ISO_1	ANSI	ISO8859-1, Latin 1, Western
ISO_BINENG	ANSI	Binary ordering, English ISO/ASCII 7-bit letter case mappings
ISO1LATIN1	ANSI	ISO8859-1, ISO Latin 1, Western, Latin 1 Ordering
ISO9LATIN1	ANSI	ISO8859-15, ISO Latin 9, Western, Latin 1 Ordering
UTF8	Multibyte	Unicode Transformation Format-8, 8-bit multibyte encoding for UCS-2

Recommended collations

When you create a database, Adaptive Server Anywhere chooses a default collation based on operating system language and character set settings. In most cases, the default collation is a suitable choice, but you can also explicitly choose a collation to match your needs from the wide selection of supplied collations.

In some cases, Adaptive Server Anywhere supports more than one collation for a particular language. Different collations vary in the characters they include and in the way they sort special characters, including ligatures (characters consisting of two or more characters joined together) and accented characters. You should choose a collation that uses a character set and sort order that are appropriate for the data in your database.

☞ For more information about sorting of data and international features, see "Adaptive Server Anywhere international features" on page 250.

Another consideration when you choose a collation is whether to use an ANSI or an OEM collation. OEM collations are suited to character-mode applications (those using the console or command prompt window) in Windows 95/98/Me and Windows NT/2000/XP. ANSI collations are designed for Windows environments, and are recommended in the majority of cases.

☞ For more information about the differences between ANSI and OEM collations, see "ANSI or OEM?" on page 274 and "ANSI and OEM code pages in Windows" on page 255.

The following collations are recommended because they provide the most likely match to the character set being used by the application: they have an appropriate sort order and support for the symbols and accented characters required for each specific language. The recommended collations are italicized in cases where alternative collations are also provided. Some languages do not have a recommended UNIX collation.

Last resort collation

If Adaptive Server Anywhere cannot locate a specific collation based on your operating system's character set and language, the Adaptive Server Anywhere registry setting, or the `SQLLOCALE` environment variable, `ISO_BINENG` becomes the default, regardless of which operating system you are using.

If this is not suitable for your requirements, you can rebuild the database and choose a different collation.

☞ For more information about changing collations, see "Changing a database from one collation to another" on page 294.

Language	Windows collations	UNIX collations
Western European (including English, French, German, Italian, Portuguese, and Spanish)	1252LATIN1, ISO9LATIN1, ISO1LATIN1, ISO_1	ISO9LATIN1, ISO1LATIN1, ISO_1
Eastern European (including Croatian, Czech, Hungarian, Romanian, Serbian, Slovakian, and Slovenian)	1250LATIN2	
Japanese	932JPN, SJIS, SJIS2	EUC_JAPAN
Korean	949KOR	EUC_KOREA
Polish	1250POL	
Russian and Ukrainian	1251CYR	
Simplified Chinese	936ZHO	EUC_CHINA
Traditional Chinese	950TWN	EUC_TAIWAN

For information about creating a database with a specific collation, see "Creating a database with a named collation" on page 290.

Creating databases for Windows CE

Windows CE is a Unicode-based operating system. The Adaptive Server Anywhere ODBC driver supports either ASCII (8-bit) or Unicode strings, and carries out character set translation as needed. If developing Embedded SQL applications, you can use Windows API functions to get the Unicode versions of strings from the database.

When creating databases for Windows CE, you should use a collation based on the same single- or multi-byte character set that Windows would use for the language of interest. For example, if you are using English, French, or German, use the 1252Latin1 collation. If you are using Japanese, use the 932JPN collation, and if you are using Korean, use the 949KOR collation.

ANSI or OEM?

Adaptive Server Anywhere collations are based on code pages that are designated as either ANSI or OEM. In most cases, using an ANSI code page is recommended. If you use OEM, choose a code page that matches the OEM code pages on your users' client machines.

You should not use a separate translation driver under any circumstance. Translation drivers interfere with the server's character set translation. Using a separate translation driver will likely result in data corruption.

For Interactive SQL and Sybase Central, the jConnect driver or the JDBC/ODBC bridge (depending on which method you use) handles character set translation.

Notes on ANSI collations

The ISO_1 collation

ISO_1 is provided for compatibility with the Adaptive Server Enterprise default ISO_1 collation. The differences are as follows:

- ◆ The lower case letter sharp s (\xDF) sorts with the lower case s in Adaptive Server Anywhere, but after ss in Adaptive Server Enterprise.
- ◆ The ligatures corresponding to **Æ** and **æ** (\xC6 and \xE6) sort after **A** and **a** respectively in Adaptive Server Anywhere, but after **Æ** and **æ** in Adaptive Server Enterprise.

The 1252LATIN1 collation

This collation includes the euro currency symbol and several other characters (Z-with-caron and z-with-caron). For single-byte Windows operating systems, this is the recommended collation in most cases. This collation is recommended for Windows users using English or Western European languages.

Windows NT service patch 4 changes the default character set in many locales to a new 1252 character set on which 1252LATIN1 is based. If you have this service patch, you should use this collation.

The euro symbol sorts with the other currency symbols.

The ISO1LATIN1 collation

This collation is the same as ISO_1, but with sorting for values in the range A0-BF. For compatibility with Adaptive Server Enterprise, the ISO_1 collation has no characters for 0xA0-0xBF. However the ISO Latin 1 character set on which it is based does have characters in these positions. The ISO1LATIN1 collation reflects the characters in these positions.

If you are not concerned with Adaptive Server Enterprise compatibility, ISO1LATIN1 is generally recommended instead of ISO_1.

ISO1LATIN1 is recommended for Unix users using English or Western European languages.

The ISO9LATIN1 collation

This collation is the same as ISO1LATIN1, but it includes the euro currency symbol and the other new characters included in the 1252 LATIN1 collation.

If your machine uses the ISO Latin 9 character set, then you should use this collation.

Notes on OEM collations

The following table displays the built-in collations that correspond to OEM code pages. The table and the corresponding collations were derived from several manuals from IBM concerning National Language Support, subject to the restrictions mentioned above. (This table represents the best information available at the time of writing.)

Country	Language	Primary Code Page	Primary Collation	Secondary Code Page	Secondary Collation
Argentina	Spanish	850	850ESP	437	437ESP
Australia	English	437	437LATIN1	850	850LATIN1
Austria	German	850	850LATIN1	437	437LATIN1
Belgium	Belgian Dutch	850	850LATIN1	437	437LATIN1
Belgium	Belgian French	850	850LATIN1	437	437LATIN1
Belarus	Belarussian	855	855CYR		
Brazil	Portuguese	850	850LATIN1	437	437LATIN1
Bulgaria	Bulgarian	855	855CYR	850	850CYR
Canada	Cdn French	850	850LATIN1	863	863LATIN1
Canada	English	437	437LATIN1	850	850LATIN1
Croatia	Croatian	852	852LATIN2	850	850LATIN2
Czech Republic	Czech	852	852LATIN2	850	850LATIN2
Denmark	Danish	850	850DAN		
Finland	Finnish	850	850SVE	437	437SVE
France	French	850	850LATIN1	437	437LATIN1
Germany	German	850	850LATIN1	437	437LATIN1


Country	Language	Primary Code Page	Primary Collation	Secondary Code Page	Secondary Collation
Greece	Greek	869	869ELL	850	850ELL
Hungary	Hungarian	852	852LATIN2	850	850LATIN2
Iceland	Icelandic	850	850ISL	861	861ISL
Ireland	English	850	850LATIN1	437	437LATIN1
Israel	Hebrew	862	862HEB	856	856HEB
Italy	Italian	850	850LATIN1	437	437LATIN1
Mexico	Spanish	850	850ESP	437	437ESP
Netherlands	Dutch	850	850LATIN1	437	437LATIN1
New Zealand	English	437	437LATIN1	850	850LATIN1
Norway	Norwegian	865	865NOR	850	850NOR
Peru	Spanish	850	850ESP	437	437ESP
Poland	Polish	852	852LATIN2	850	850LATIN2
Portugal	Portuguese	850	850LATIN1	860	860LATIN1
Romania	Romanian	852	852LATIN2	850	850LATIN2
Russia	Russian	866	866RUS	850	850RUS
S. Africa	Afrikaans	437	437LATIN1	850	850LATIN1
S. Africa	English	437	437LATIN1	850	850LATIN1
Slovak Republic	Slovakian	852	852LATIN2	850	850LATIN2
Slovenia	Slovenian	852	852LATIN2	850	850LATIN2
Spain	Spanish	850	850ESP	437	437ESP
Sweden	Swedish	850	850SVE	437	437SVE
Switzerland	French	850	850LATIN1	437	437LATIN1
Switzerland	German	850	850LATIN1	437	437LATIN1
Switzerland	Italian	850	850LATIN1	437	437LATIN1
Turkey	Turkish	857	857TRK	850	850TRK
UK	English	850	850LATIN1	437	437LATIN1

Country	Language	Primary Code Page	Primary Collation	Secondary Code Page	Secondary Collation
USA	English	437	437LATIN1	850	850LATIN1
Venezuela	Spanish	850	850ESP	437	437ESP
Yugoslavia	Macedonian	852	852LATIN2	850	850LATIN2
Yugoslavia	Serbian Cyrillic	855	855CYR	852	852CYR
Yugoslavia	Serbian Latin	852	852LATIN2	850	850LATIN2

Using multibyte collations

This section describes how multibyte character sets are handled. The description applies to the supported collations and to any multibyte custom collations you may create.

Adaptive Server Anywhere provides collations using several multibyte character sets.

 For more information, see "Understanding collations" on page 269.

Adaptive Server Anywhere supports variable width character sets. In these sets, some characters are represented by one byte, and some by more than one, to a maximum of four bytes. The value of the first byte in any character indicates the number of bytes used for that character, and also indicates whether the character is a space character, a digit, or an alphabetic (alpha) character.

Adaptive Server Anywhere does not support fixed-length multibyte character sets such as 2-byte Unicode (UCS-2) or 4-byte Unicode (UCS-4).

Understanding character set translation

Adaptive Server Anywhere can carry out character set translation among character sets that represent the same characters, but at different positions in the character set or code page. There needs to be a degree of compatibility between the character sets for this to be possible. For example, character set translation is possible between EUC-JIS and Shift-JIS character sets, but not between EUC-JIS and OEM code page 850.

This section describes how Adaptive Server Anywhere carries out character set translation. This information is provided for advanced users, such as those who may be deploying applications or databases in a multi-character-set environment.

Character translation for database messages

Error and other messages from the database software are held in a language resource library. Localized versions of this library are provided with localized versions of Adaptive Server Anywhere.

Client application users may see messages from the database as well as data from the database. Some database messages, which are strings from the language library, may include placeholders that are filled by characters from the database. For example, if you execute a query with a column that does not exist, the returned error message is:


Column *column-name* not found

where *column-name* is filled in from the database.

To present these kinds of information to the client application in a consistent manner, even if the database is in a different character set from the language library, the database server automatically translates the characters of the messages so that they match the character set used in the database collation.

To use character translation for database messages, ensure that the collation for your database is compatible with the character set used on your computer, and with the character set used in the Adaptive Server Anywhere language resource library. The language resource library differs among different localized versions of Adaptive Server Anywhere. You must check that the characters of interest to you exist in each character set.

Messages are always translated into the database collation character set, regardless of whether character set translation is turned on or off.



A further character set translation is carried out if character set is turned on (the default) for the database server, and if the client character set is different from that used in the database collation.

Connection strings and character sets

Connection strings present a special case for character set translation. The connection string is parsed by the client library, in order to locate or start a database server. This parsing is done with no knowledge of the server character set or language.

The interface library parses the connection string as follows:

- 1 It is broken down into its *keyword = value* components. This can be done independently of the character set, as long as you do not use the curly braces {} around **CommLinks (LINKS)** connection parameters. Instead, use the recommended parentheses (). Curly braces are valid follow bytes (bytes other than the first byte) in some multi-byte character sets.
- 2 The server is located. The server name is interpreted according to the character set of the client machine. In the case of Windows operating systems, the ANSI character set is used. Extended chars can be used unless they cause character set conversion issues between the client and server machines.

For maximum compatibility among different machines, you should use server names built from alphabetic or numeric ASCII characters 0 to 127 (or 33 to 126) and the underscore, using no punctuation characters. Server names are truncated at 40 characters.


- 3 The **DatabaseName (DBN)** or **DatabaseFile (DBF)** connection parameters are interpreted in the database server character set.

- 4 Once the database is located, the remaining connection parameters are interpreted according to its character set.

Avoiding character-set translation

There is a performance cost associated with character set translation. If you can set up an environment such that no character set translation is required, then you do not have to pay this cost, and your setup is simpler to maintain.

If you work with a single-byte character set and are concerned only with seven-bit ASCII characters (values 0 through 127), then you do not need character set translation. Even if the code pages are different in the database and on the client operating system, they are compatible over this range of characters. Many English-language installations meet these requirements. In Adaptive Server Anywhere 8.0 and higher, character set translation is turned on by default. You can turn it off using the `-cst-` option.

 For more information, see "Turning off character set translation on a database server" on page 292.

If you do require the use of extended characters, there are other steps you may be able to take:

- ◆ If the code page on your client machine operating system matches that used in the database, no character set translation is needed for data in the database.

For example, in many environments it is appropriate to use the 1252LATIN1 collation in your database, which corresponds to the ANSI code page in many single-byte environments.

- ◆ If you are able to use a version of Adaptive Server Anywhere built for your language, and if you use the code page on your operating system, no character set translation is needed for database messages. The character set used in the Adaptive Server Anywhere message strings is as follows:


Language	Character set
English	cp1252
French	cp1252
German	cp1252
Japanese	cp932 (Shift-JIS)

Also, recall that client/server character set translation takes place by default. You can turn it off using the `-cst-` option.

☞ For more information, see "-ct server option" on page 131.

Collation internals

This section describes internal technical details of collations, including the file format of collation files.

 This section is of particular use if you want to create a database using a custom collation. For information on the steps involved, see "Creating a custom collation" on page 292, and "Creating a database with a custom collation" on page 293.

You can create a database using a collation different from the supplied collations. This section describes how to build databases using such custom collations.

In building multibyte custom collations, you can specify which ranges of values for the first byte signify single- and double-byte (or more) characters, and which specify space, alpha, and digit characters. However, all first bytes of value less than \x40 must be single-byte characters, and no follow bytes may have values less than \x40. This restriction is satisfied by all supported encodings.

Collation files may include the following elements:

- ◆ Comment lines, which are ignored by the database.
- ◆ A title line.
- ◆ A collation sequence section.
- ◆ An Encodings section.
- ◆ A Properties section.

Comment lines

In the collation file, spaces are generally ignored. Comment lines start with either the percent sign (%) or two dashes (--).

The title line

The first non-comment line must be of the form:

Collation *collation_label (collation_name) (ase_charset) (ase_so_sensitive) (ase_so_insensitive) (java_charset)*

Descriptions of arguments

Argument	Description
<i>Collation</i>	A required keyword.
<i>collation_label</i>	The collation label, which appears in the system tables as SYS.SYSCOLLATION.collation_label and SYS.SYSINFO.default_collation. The label must contain no more than 10 characters.
<i>collation_name</i>	A text description of the collation, usually describing the character set and ordering of the collation.
<i>ase_charset</i>	The ASE name of a character set matching the character set of this collation. This name is used to make character set mappings when server character set translation is enabled (the default).
<i>ase_so_sensitive</i>	The name of an Open Client or ASE case sensitive collation matching this collation.
<i>ase_so_insensitive</i>	The name of an Open Client or ASE case sensitive collation matching this collation.
<i>java_charset</i>	The name of a Java character set matching the character set of this collation. This name is used when converting character data between the Java virtual machine and the database.

For example, the 932JPN collation file contains the following collation line, with label 932JPN and name (Code Page 932, Japanese Shift-JIS with Microsoft extensions):

```
Collation 932JPN (Code Page 932, Japanese Shift-JIS with Microsoft extensions) (cp932)
(bin_cp932) (bin_cp932) (SJIS)
```

The collation sequence section

After the title line, each non-comment line describes one position in the collation. The ordering of the lines determines the sort ordering used by the database, and determines the result of comparisons. Characters on lines appearing higher in the file (closer to the beginning) sort before characters that appear later.

The form of each line in the sequence is:

```
[sort-position] : character [ [, character ] ...]
```

or

```
[sort-position] : character [lowercase uppercase]
```

Descriptions of arguments

Argument	Description
<i>sort-position</i>	Optional. Specifies the position at which the characters on that line will sort. Smaller numbers represent a lesser value, so will sort closer to the beginning of the sorted set. Typically, the sort-position is omitted, and the characters sort immediately following the characters from the previous sort position.
<i>character</i>	The character whose sort-position is being specified.
<i>lowercase</i>	Optional. Specifies the lowercase equivalent of the character. If not specified, the character has no lowercase equivalent.
<i>uppercase</i>	Optional. Specifies the uppercase equivalent of the character. If not specified, the character has no uppercase equivalent.

Multiple characters may appear on one line, separated by commas (.). In this case, these characters are sorted and compared as if they were the same character.

Specifying character and sort-position

Each character and sort position is specified in one of the following ways:

Specification	Description
<code>\dnnn</code>	Decimal number, using digits 0-9 (such as <code>\d001</code>)
<code>\xhh</code>	Hexadecimal number, using digits 0-9 and letters a-f or A-F (such as <code>\xB4</code>)
<code>'c'</code>	Any character in place of c (such as <code>','</code>)
<code>c</code>	Any character other than quote (<code>'</code>), back-slash (<code>\</code>), colon (<code>:</code>) or comma (<code>,</code>). These characters must use one of the previous forms.

The following are some sample lines for a collation:

```
% Sort some special characters at the beginning:
: ' '
: -
: \xF2
: \xEE
: \xF0
: -
: ' , '
: ;
: ' : '
: !
% Sort some letters in alphabetical order
: A a A
: a a A
: B b B
: b b B
% Sort some E's from code page 850,
% including some accented extended characters:
: e e E, \x82 \x82 \x90, \x8A \x8A \xD4
: E e E, \x90 \x82 \x90, \xD4 \x8A \xD4
```

Other syntax notes

For databases using case-insensitive sorting and comparison (no `-c` specified in the `dbinit` command), the lower case and upper case mappings are used to find the lower case and upper case characters that are sorted together.

For multibyte character sets, the first byte of a character is listed in the collation sequence, and all characters with the same first byte are sorted together, and ordered according to the value of the following bytes. For example, the following is part of the Shift-JIS collation file:

```
: \xfb
: \xfc
: \xfd
```

In this collation, all characters with first byte `\xfc` come after all characters with first byte `\xfb` and before all characters with first byte `\xfd`. The two-byte character `\xfc \x01` would be ordered before the two-byte character `\xfc \x02`.

Any characters omitted from the collation are added to the end of the collation. The tool that processes the collation file issues a warning.

The encodings section

The encodings section is optional, and follows the collation sequence. It is not useful for single-byte character sets.

The encodings section lists which characters are lead-bytes, for multi-byte character sets, and what are valid follow-bytes.

For example, the Shift-JIS encodings section is as follows:

```
Encodings:  
[\x00-\x80,\xa0-\xdf,\xf0-\xff]  
[\x81-\x9f,\xe0-\xef][\x00-\xff]
```

The first line following the section title lists valid single-byte characters. The square brackets enclose a comma-separated list of ranges. Each range is listed as a hyphen-separated pair of values. In the Shift-JIS collation, values `\x00` to `\x80` are valid single-byte characters, but `\x81` is not a valid single-byte character.

The second line following the section title lists valid multibyte characters. Any combination of one byte from the second line followed by one byte from the first is a valid character. Therefore `\x81\x00` is a valid double-byte character, but `\xd0 \x00` is not.

The properties section

The properties section is optional, and follows the encodings section.

If a properties section is supplied, an encodings section must be supplied also.

The properties section lists values for the first-byte of each character that represent alphabetic characters, digits, or spaces.

The Shift-JIS properties section is as follows:

```
Properties:  
space: [\x09-\x0d,\x20]  
digit: [\x30-\x39]  
alpha: [\x41-\x5a,\x61-\x7a,\x81-\x9f,\xe0-\xef]
```

This indicates that characters with first bytes `\x09` to `\x0d`, as well as `\x20`, are to be treated as space characters, digits are found in the range `\x30` to `\x39` inclusive, and alphabetic characters in the four ranges `\x41-\x5a`, `\x61-\x7a`, `\x81-\x9f`, and `\xe0-\xef`.

International language and character set tasks

This section groups together the tasks associated with international language and character set issues.

Finding the default collation


If you do not explicitly specify a collation when creating a database, a default collation is used. The default collation depends on the operating system you are working on.

❖ To find the default collation for your machine:

- 1 Start Interactive SQL. Connect to the sample database.
- 2 Enter the following query:


```
SELECT PROPERTY( 'DefaultCollation' )
```

The default collation is returned.

 For more information about this collation, see "Supplied collations" on page 269.


Configuring your character set environment

This section describes how to set up your computing environment so that character set issues are handled properly. If you set your locale environments properly, then you do not need to explicitly choose collations for your databases, and you can turn off character set translation between client and server.

 For more information about turning off character set translation, see "Turning off character set translation on a database server" on page 292.

❖ To configure your character set environment:

- 1 Determine the default locale of each computing platform in your environment. The default locale is the character set and language of each computer. On Windows operating systems, the character set is the ANSI code page.

 For more information about how to find locale information, see "Determining locale information" on page 288.

- 2 Decide whether the locale settings are appropriate for your environment.

☞ For more information, see "Understanding collations" on page 269.

- 3 If the default settings are inappropriate, decide on a character set, language, and database collation that matches your data and avoids character set translation.

☞ For more information, see "Avoiding character-set translation" on page 280.

- 4 Set locales on each of the machines in the environment to these values.

☞ For more information, see "Setting locales" on page 289.

- 5 Create your database using the default collation. If the default collation does not match your needs, create a database using a named collation.

☞ For more information, see "Creating a database with a named collation" on page 290, and "Changing a database from one collation to another" on page 294.

When choosing the collation for your database,

- ◆ Choose a collation that uses a character set and sort order appropriate for the data in the database. It is often the case that there are several alternative collations that meet this requirement, including some that are OEM collations and some that are ANSI collations.
- ◆ There is a performance cost, as well as extra complexity in system configuration, when you use character set translation. Choose a collation that avoids the need for character set translation.

You can avoid character set translation by using a collation sequence in the database that matches the character set in use on your client machine operating system. In the case of Windows operating systems on the client machine, choose the ANSI character set. Character set translation is enabled by default for database servers that are version 8.0 or higher of Adaptive Server Anywhere. You can turn off character set translation using the `-ct-` option.

☞ For more information, see "Avoiding character-set translation" on page 280.

Determining locale information

You can determine locale information using system functions.

☞ For more information, see "System functions" on page 101 of the book *ASA SQL Reference Manual*.

❖ **To determine the locale of a database server:**

- 1 Start Interactive SQL, and connect to a database server.
- 2 Execute the following statement to determine the database server character set:

```
SELECT PROPERTY( 'CharSet' )
```

The query returns one of the supported character sets listed in "Character set labels" on page 265.

- 3 Execute the following statement to determine the database server language:

```
SELECT PROPERTY( 'Language' )
```

The query returns one of the supported languages listed in "Language label values" on page 263.

- 4 Execute the following statement to determine the database server default collation:

```
SELECT PROPERTY( 'DefaultCollation' )
```

The query returns one of the collations listed in "Supplied collations" on page 269.

Notes

To obtain client locale information, connect to a database server running on your current machine.

To obtain the character set for an individual database, execute the following statement:

```
SELECT DB_PROPERTY ( 'CharSet' )
```

Setting locales

You can use the default locale on your operating system, or explicitly set a locale for use by the Adaptive Server Anywhere components on your machine.

❖ **To set the Adaptive Server Anywhere locale on a computer:**

- 1 If the default locale is appropriate for your needs, you do not need to take any action.

✍ For more information about how to find out the default locale of your operating system, see "Determining locale information" on page 288.

- 2 If you need to change the locale, create a `SQLLOCALE` environment variable with the following value:

```
Charset=cslabel;Language=langlabel;CollationLabel=collabel
```


where *cslabel* is a character set label from the list in "Character set labels" on page 265, *langlabel* is a language label from the list in "Language label values" on page 263, and *CollationLabel* is taken from the list in "Understanding collations" on page 269, or is a custom collation label.

☞ For more information on how to set environment variables on different operating systems, see "Setting environment variables" on page 208.

Creating a database with a named collation

You may specify the collation for each database when you create the database. The default collation is inferred from the code page and sort order of the database server's computer's operating system.

- ❖ **To specify a database collation when creating a database (Sybase Central):**
 - ◆ You can use the Create Database wizard in Sybase Central to create a database. In the wizard, there is a page where you choose a collation from a list.



❖ **To specify a database collation when creating a database (command prompt):**

- 1 List the recommended collation sequences:

```
dbinit -l
```

The first column of the list is the collation label, which you supply when creating the database.

```
437LATIN1  Code Page 437, Latin 1, Western
437ESP     Code Page 437, Spanish
437SVE     Code Page 437, Swedish/Finnish
819CYR     Code Page 819, Cyrillic
819DAN     Code Page 819, Danish
819ELL     Code Page 819, Greek
...
```

- 2 Create a database using the *dbinit* utility, specifying a collation sequence using the *-z* option. The following command creates a database with a Greek collation.

```
dbinit -z 869ELL mydb.db
```

❖ **To specify a database collation when creating a database (SQL):**

- ◆ You can use the CREATE DATABASE statement to create a database. The following statement creates a database with a Greek collation:

```
CREATE DATABASE 'mydb.db'
COLLATION '819ELL'
```

Starting a database server using character set translation

Character set translation takes place if the client and server locales are different. Character set translation is turned on by default for database servers that are Adaptive Server Anywhere version 8.0 or higher. For database servers that are version 7.x or earlier, you must explicitly turn on character set conversion in the database server command.

❖ **To enable character set translation on a database server:**

- ◆ Start the database server using the *-ct+* option. For example:

```
dbsrv8 -ct+ asademo.db
```

Turning off character set translation on a database server

Character set translation is turned on by default. If you do not require character set translation, you can disable it in the database server command using the `-ct-` option.

❖ To disable character set translation on a database server:

- ◆ Start the database server using the `-ct-` option. For example:

```
dbsrv8 -ct- asademo.db
```

🔗 For more information, see "`-ct` server option" on page 131.

Creating a custom collation

If none of the supplied collations meet your needs, you can modify a supplied collation to create a **custom collation**. You can then use this custom collation when creating a database.

🔗 For more information about supplied collations, see "Supplied collations" on page 269.

❖ To create a custom collation:

- 1 **Decide on a starting collation** You should choose a collation as close as possible to the one you want to create as a starting point for your custom collation.

For a listing of recommended collations, see "Understanding collations" on page 269. Alternatively, run *dbinit* with the `-l` (lower case L) option:

```
dbinit -l
```

- 2 **Create a custom collation file** You do this using the Collation utility. The output is a collation file.

For example, the following statement extracts the 1252LATIN1 collation into a file named *mycustomcol.col*:

```
dbcollat -z 1252LATIN1 mycustomcol.col
```

- 3 **Edit the custom collation file** Open the collation file (in this case *mycustomcol.col*) in a text editor.
- 4 **Change the collation label** The collation label is specified on a line near the top of the file, starting with **Collation**.

The other entries on this line relate the Adaptive Server Anywhere collation label to the names that Java and the Sybase TDS interface give to the same collation information. If you are not using these interfaces you do not need to alter these entries.


The Collation line takes the following form:

```
Collation collation_label (collation_name)  
      (ase_charset) (ase_so_sensitive)  
      (ase_so_insensitive) (java_charset)
```

where character set label (*ase_charset*) and the two sort-order labels (*ase_so_sensitive* and *ase_so_insensitive*) state which Open Client character set and sort order is the closest to the current collation. The *java_charset* label is the closest character set known to Java.

You should edit this line to provide a new label. The label you need to change is the *collation_label*: in the example in step two, it is 1252LATIN1.

- 5 **Change the collation definition** Make the changes you wish in the custom collation file to define your new collation.

 For more information about the collation file contents and format, see "Collation internals" on page 282.

- 6 **Convert the file to a SQL script** You do this using the *dbcollat* utility using the *-d* option.

For example, the following command creates the *mycustom.SQL* file from the *mycustomcol.col* collation file:

```
dbcollat -d mycustomcol.col mycustom.SQL
```

- 7 **Add the SQL script to the script in your installation** The script used when creating databases is held in the scripts subdirectory of your Adaptive Server Anywhere installation directory. Append the contents of *mycustom.SQL* to the end of *custom.SQL*.

The new collation is now in place, and can be used when creating a database.

Creating a database with a custom collation

If none of the supplied collations meet your needs, you can create a database using a custom collation. The custom collation is used in indexes and any string comparisons.

❖ **To create a database with a custom collation:**

- 1 Create a custom collation.

You must have a custom collation in place to use when creating a database.

✍ For more information about how to create custom collations, see "Creating a custom collation" on page 292.

- 2 Create the new database.

Use the Initialization utility, specifying the name of your custom collation.

For example, the following command creates a database named *temp.db* using the custom collation sequence **newcol**.

```
dbinit -z newcol temp.db
```

You can also use the Initialization utility from Sybase Central.

Changing a database from one collation to another

Changing your database from one collation to another may be a good idea for any number of reasons. It can be especially useful, for example, if you want to:

- ◆ avoid the need for character set translation across your setup.
- ◆ unify the character set you are using with the collation in your database. Using the same character set defined in your database is especially important for sorting purposes.
- ◆ use a different character. You may, for example, want to move from an OEM character set to a Windows character set.

Simply modifying the collation in an existing database is not permitted, since it would invalidate all the indexes for that database. In order to change the collation for a database, you must rebuild the database. Rebuilding a database creates a new database with new settings (including collation settings), using the old database's data.

When you change the collation for a database, there are two main scenarios to consider. The difference between the two lies in whether the character set of the data needs to be changed.

Example 1

In the first example, only the collation needs to be changed. The data should not change character sets. To resolve the collation issue, you need to rebuild the database with new collation settings using the old data.

Consider an old database using the 850LATIN1 collation. If the database contains data inserted from a Windows 'windowed' application, it is likely that the data is actually from the CP1252 character set, which does not match CP850 used by the 850LATIN1 collation. This situation will often be discovered when an ORDER BY clause seems to sort accented characters incorrectly. To correct this problem, you would create a new database using the 1252LATIN1 collation, and move the data from the old database to the new database without translation, since the data is already in the character set (CP1252) that matches the new database's collation.

The simplest way to ensure that translation does not occur is to start the server with the `-ct-` option.

☞ For more information about rebuilding a database, see "Rebuilding databases" on page 440 of the book *ASA SQL User's Guide*.

☞ For more information about specifying collations when creating databases, see "Creating a database with a named collation" on page 290.

Example 2

In the second situation, both the collation and the character set need to be changed. To resolve the collation and character set issues, you need to rebuild the database with the new collation settings, and change the character set of the data.

Suppose that the 850LATIN1 database had been used properly such that it contains characters from the CP850 character set. However, you want to update both the collation and the character set, perhaps to avoid extra translation. You would create a new database using 1252LATIN1, and move the data from the old database to the new database with translation, thus converting the CP850 characters to CP1252.

The translation of the database data from one character set to another occurs using the client-server translation feature of the server. This feature, which is turned on by default and can also be invoked using the `-ct+` option, translates the data during the communication between the client application and the server. The database's collation determines the server's character set. The locale setting of the operating system determines the client's default character set, however, the client's character set can be overridden by the **CharSet (CS)** connection parameter.

☞ For more information about the **CharSet (CS)** connection parameter, see "CharSet connection parameter" on page 167.

Since character set translation takes place during the communication between the client application and the server, an external unload or reload is necessary. If you use internal unload and reload, you avoid character set translation altogether, and end up where you began. Similarly, if character set translation occurs in both the unload and the reload steps, you perform the translation and then immediately undo the translation and still end up where you began. Character set translation can occur in either the unload or the reload steps, but not in both.

❖ **To convert a database from one collation to another, and translate the data's character set (using translation on reload):**

- 1 Unload the data from the source database.

You can use the Unload utility to produce a *reload.SQL* file and a set of data files in the character set of the source database. Since we do not want any translation during this phase, ensure that character set translation is not enabled on the server running the source database. For servers before version 8, ensure that `-ct` is not specified. If you are using a server that is version 8 or higher, ensure that `-ct-` (no character set translation) is specified when the server is started.

If the unload/reload is occurring on a single machine, use the `-ix` option to do an internal unload and an external reload. If the unload/reload occurs across machines, use the Unload utility with the `-xx` option to force an external unload and an external reload.

Remember that an "external" unload or reload means that an application (dbunload and DBISQL) opens a cursor on the database and either reads or writes the data to disk. An "internal" unload or reload means that an UNLOAD TABLE or LOAD TABLE is used so that the server reads or writes the data itself.

If you want to unload data from specific tables, use the Interactive SQL OUTPUT statement.

🔗 For more information on the Unload utility, see "The Unload utility" on page 513.

- 2 Create a target database with the appropriate collation using the Initialization utility including the `-z` option to specify the collation sequence for the target database.

The database should be created and reloaded using the version of the server and tools corresponding to the server that they will use to run it.

🔗 For more information on specifying collations when creating databases, see "Creating a database with a named collation" on page 290.

- 3 Start the target database using the server running with character set translation enabled by specifying the `-ct+` option. You can also specify the `-z` option.

The `-ct+` option enables the database server to carry out character set translation into the character set of the target database. In version 8 and higher of Adaptive Server Anywhere, character set translation is enabled by default.

The `-z` option, while not required, allows for verification of the character sets that will be used, and that translation will occur. However, using the `-z` server option can cause slower performance during the reload.

☞ For more information, see "Starting a database server using character set translation" on page 291.

- 4 If you generated a *reload.SQL* file in step 1, you can run the file in a command shell using a command such as the following.

```
dbisql -c "uid=dba;pwd=sql" -codepage cp950
reload.sql
```

You must supply the appropriate connection parameters for your database, as well as the code page (cp950 in the example above) that was used to generate the *reload.SQL* file.

☞ For more information about the `-codepage` option, see "The Interactive SQL utility" on page 472.

If you exported only table data in step 1, you can import the data using the INPUT statement in Interactive SQL.

☞ For more information about the INPUT statement, see "INPUT statement [Interactive SQL]" on page 459 of the book *ASA SQL Reference Manual*.

C H A P T E R 1 1

Backup and Data Recovery

About this chapter

This chapter describes how to protect your data against operating system crashes, file corruption, disk failures, and total machine failure.

The chapter describes how to make backups of your database, how to restore data from a backup, and how to run your server so that performance and data protection concerns are addressed.

Contents

Topic	Page
Introduction to backup and recovery	300
Understanding backups	305
Designing backup procedures	308
Configuring your database for data protection	317
Backup and recovery internals	321
Backup and recovery tasks	329

Introduction to backup and recovery

A **backup** is a copy of the information in a database, held in some physically separate location from your database. If the database becomes unavailable, perhaps because of damage to a disk drive, you can **restore** it from the backup. Depending on the nature of the damage, it is often possible to restore from backups all committed changes to the database up to the time it became unavailable.

Restoring databases from a backup is one aspect of database **recovery**. The other aspect is recovery from operating system or database server crashes, and improper shutdowns. The database server checks on database startup whether the database was shut down cleanly at the end of the previous session. If it was not, the server executes an automatic recovery process to restore information. This mechanism recovers all changes up to the most recently committed transaction.

Chapter contents

This chapter contains the following material:

- ◆ An introduction to backup and recovery (this section).
- ◆ Concepts and background information to help you design and use an appropriate backup strategy:
 - ☞ For more information, see "Understanding backups" on page 305.
- ◆ Information to help you decide the type and frequency of backups you use, and the way you run your database server so that your data is well protected:
 - ☞ For more information, see "Designing backup procedures" on page 308 and "Configuring your database for data protection" on page 317.
- ◆ Information for advanced users, describing Adaptive Server Anywhere internal operations related to backup and recovery:
 - ☞ For more information, see "Backup and recovery internals" on page 321.
- ◆ Step by step instructions for how to carry out backup and recovery tasks.
 - ☞ For more information, see "Backup and recovery tasks" on page 329

Questions and answers

To answer the question...	Consider reading...
What is a backup?	"Introduction to backup and recovery" on page 300
What is recovery?	"Introduction to backup and recovery" on page 300
What is a transaction log?	"The transaction log" on page 305
What are media and system failure?	"Protecting your data against failure" on page 302
From what kinds of failure do backups protect my data?	"Protecting your data against failure" on page 302
What tools are available for backups?	"Ways of making backups" on page 303
What types of backup are available?	"Types of backup" on page 308
What type of backup should I use?	"Designing backup procedures" on page 308
If my database file or transaction log becomes corrupt, what data may be lost?	"Protecting your data against media failure" on page 306
How are backups executed?	"Understanding backups" on page 305
How often do I carry out backups?	"Scheduling backups" on page 309
Can I schedule automatic backups?	"Scheduling backups" on page 309
My database is involved in replication. How does this affect my backup strategy?	"A backup scheme for databases involved in replication" on page 312 "Backup methods for remote databases in replication installations" on page 314
How can I backup to tape?	"Backing up a database directly to tape" on page 338
How do I plan a backup schedule?	"Designing a backup and recovery plan" on page 315
Can I automate backups?	"Automating Tasks Using Schedules and Events" on page 231
How can I be sure that my database file is not corrupt?	"Ensuring your database is valid" on page 315 "Validating a database" on page 331

To answer the question...	Consider reading...
How can I be sure that my transaction log is not corrupt?	"Validating the transaction log on database startup" on page 327 "Validating a transaction log" on page 332
How can I run my database for maximum protection against failures?	"Configuring your database for data protection" on page 317
How can I ensure high availability and machine redundancy?	"Protecting against total machine failure" on page 318 "Making a live backup" on page 340
How do I carry out a backup?	"Making a full backup" on page 329
How do I restore data from backups when a failure occurs?	"Recovering from media failure on the database file" on page 340 "Recovering from media failure on an unmirrored transaction log" on page 341 "Recovering from media failure on a mirrored transaction log" on page 342

Protecting your data against failure


If your database has become unusable, you have experienced a database **failure**. Adaptive Server Anywhere provides protection against the following categories of failure:

Media failure The database file and/or the transaction log become unusable. This may occur because the file system or the device storing the database file becomes unusable, or it may be because of file corruption.

For example:

- ◆ The disk drive holding the database file or the transaction log file becomes unusable.
- ◆ The database file or the transaction log file becomes corrupted. This can happen because of hardware problems or software problems.

Backups protect your data against media failure.


 For more information, see "Understanding backups" on page 305.


System failure A system failure occurs when the computer or operating system goes down while there are partially completed transactions. This could occur when the computer is inappropriately turned off or rebooted, when another application causes the operating system to crash, or because of a power failure.

For example:

- ◆ The computer or operating system becomes temporarily unavailable while there are partially completed transactions, perhaps because of a power failure or operating system crash, or because the computer is inappropriately rebooted.

After a system failure occurs, the database server recovers automatically when you next start the database. The results of each transaction committed before the system error are intact. All changes by transactions that were not committed before the system failure are canceled.

 For more information about the recovery mechanism, see "Backup and recovery internals" on page 321.

 It is possible to recover uncommitted changes manually. For information, see "Recovering uncommitted operations" on page 344.

Ways of making backups

There are several distinct ways of making backups. This section introduces each of the major approaches, but does not address any issues of appropriate options.

You can make backups in the following ways:

- ◆ **Sybase Central** You can use the Backup Database wizard in Sybase Central to make a backup. You can access the wizard by selecting a database and choosing Backup Database from the File menu (or from the popup menu).
- ◆ **utility** The *dbbackup* command-line utility makes backups. For example, executing the following command at the command prompt makes backup copies of the database and transaction log in the directory *c:\backup* on the client machine:

```
dbbackup -c "connection-string" c:\backup
```

- ◆ **SQL Statement** You can use a SQL statement to make the database server execute a backup operation. For example, the following statement places backup copies of the database file and transaction log into the directory *c:\backup* on the server machine.

```
BACKUP DATABASE  
DIRECTORY 'c:\\backup'
```

- ◆ **Offline backup** The above examples are all online backups, executed against a running database. You can make offline backups by copying the database files when the database is not running.

Notes

You must have DBA authority or REMOTE DBA authority to make backups of a database.

Understanding backups

To understand what files you need to back up, and how you restore databases from backups, you need to understand how the changes made to the database are stored on disk.

The database file


When a database is shut down, the database file holds a complete and current copy of all the data in the database. When a database is running, however, the database file is generally not current.

The only time a database file is guaranteed to hold a complete and current copy of all data is when a checkpoint takes place. At a checkpoint, all the contents of the database cache are written out to the disk.

The database server checkpoints a database under the following conditions:

- ◆ As part of the database shutdown operations
- ◆ When the amount of time since the last checkpoint exceeds the database option `CHECKPOINT_TIME`
- ◆ When the estimated time to do a recovery operation exceeds the database option `RECOVERY_TIME`
- ◆ When the database server is idle long enough to write all dirty pages
- ◆ When a connection issues a `CHECKPOINT` statement
- ◆ When the database server is running without a transaction log and a transaction is committed

Between checkpoints, you need both the database file and another file, called the transaction log, to ensure that you have a complete copy of all committed transactions.

 For more information about checkpoints, see "Checkpoints and the checkpoint log" on page 322, and "How the database server decides when to checkpoint" on page 325.

The transaction log

The **transaction log** is a separate file from the database file. It stores all changes to the database. Inserts, updates, deletes, commits, rollbacks, and database schema changes are all logged. The transaction log is also called the **forward log** or the **redo log**.

The transaction log is a key component of backup and recovery, and is also essential for data replication using SQL Remote or the Replication Agent.

By default, all databases use transaction logs. Using a transaction log is optional, but you should always use a transaction log unless you have a specific reason not to. Running a database with a transaction log provides much greater protection against failure, better performance, and the ability to replicate data.

For more information on how to use a transaction log to protect against media failure, see "Protecting against media failure on the database file" on page 317.

When changes are forced to disk

Like the database file, the transaction log is organized into **pages**: fixed size areas of memory. When a change is recorded in the transaction log, it is made to a page in memory. The change is forced to disk when the earlier of the following happens:

- ◆ The page is full.
- ◆ A COMMIT is executed.

In this way, completed transactions are guaranteed to be stored on disk, while performance is improved by avoiding a write to the disk on every operation.

Configuration options are available to allow advanced users to tune the precise behavior of the transaction log. For more information, see "COOPERATIVE_COMMITS option" on page 561, and "DELAYED_COMMITS option" on page 565.

Transaction log mirrors

A **transaction log mirror** is an identical copy of the transaction log, maintained at the same time as the transaction log. If a database has a mirrored transaction log, every database change is written to both the transaction log and the transaction log mirror. By default, databases do not have transaction log mirrors.

A transaction log mirror provides extra protection for critical data. It enables complete data recovery in the case of media failure on the transaction log. A mirrored transaction log also enables a database server to carry out automatic validation of the transaction log on database startup.

For more information, see "Protecting against media failure on the transaction log" on page 317.

Protecting your data against media failure

Backups protect your data against media failure as part of the data protection mechanisms Adaptive Server Anywhere provides.

☞ For an overview of data protection mechanisms, see "Protecting your data against failure" on page 302.

The practical aspects of recovery from media failure depend on whether the media failure is on the database file or the transaction log file.

Media failure on the database file If your database file is not usable, but your transaction log is still usable, you can recover all committed changes to the database as long as you have a proper backup procedure in place. All information since the last backed up copy of the database file is held in backed up transaction logs, or in the online transaction log.

☞ For more information on how to configure your database system, see "Protecting against media failure on the database file" on page 317.

Media failure on the transaction log file Unless you use a mirrored transaction log, you cannot recover information entered between the last database checkpoint and a media failure on the transaction log. For this reason, it is recommended that you use a mirrored transaction log in setups such as SQL Remote consolidated databases, where loss of the transaction log can lead to loss of key information, or the breakdown of a replication system.

☞ For more information, see "Protecting against media failure on the transaction log" on page 317.

Designing backup procedures

When you make a backup, you have a set of choices to make about how to manage transaction logs. The choices you make depend on a set of factors including the following:

- ◆ Is the database involved in replication?

In this chapter, replication means SQL Remote replication, or MobiLink synchronization where *dbmlsync.exe* is running, or a database using the Replication Agent. Each of these replication methods requires access to the transaction log, and potentially to old transaction logs.

- ◆ How fast is the transaction log file growing relative to your available disk space? If the transaction log is growing quickly, you may not be able to afford to keep transaction logs available.

Types of backup

This section assumes that you are familiar with basic concepts related to backups.

☞ For more information about concepts related to backups, see "Introduction to backup and recovery" on page 300, and "Understanding backups" on page 305.

Backups can be categorized in several ways:

- ◆ **Full backup and incremental backup** A **full backup** is a backup of both the database file and of the transaction log. An **incremental backup** is a backup of the transaction log only. Typically, full backups are interspersed with several incremental backups.

☞ For more information on making backups, see "Making a full backup" on page 329, and "Making an incremental backup" on page 330.

- ◆ **Server-side backup and client-side backup** You can execute an online backup from a client machine using the Backup utility. To execute a server side backup, you execute the BACKUP statement; the database server then carries out the backup.

You can easily build server side backup into applications because it is a SQL statement. Also, server-side backup is generally faster because the data does not have to be transported across the client/server communications system.

Instructions for server-side and client-side backups are given together for each backup procedure.

- ◆ **Archive backup and image backup** An **archive backup** copies the database file and the transaction log into a single archive file, typically on a tape drive. An **image backup** makes a copy of the database file and/or the transaction log, each as separate files. You can only carry out archive backups as server-side backups, and you can only make full backups.

You should use an archive backup if you are backing up directly to tape. Otherwise, an image backup has more flexibility for transaction log file management.

Archive backups are supported on Windows NT/2000/XP and UNIX platforms only. On Windows CE, only image backups are permitted.

☞ For more information on archive backups, see "Backing up a database directly to tape" on page 338.

- ◆ **Online and offline backup** Backing up a running database provides a snapshot of a consistent database, even though other users are modifying the database. An offline backup consists simply of copying the files. You should only carry out an offline backup when the database is not running, and when the database server was shut down properly.

The information in this chapter focuses on online backups.

- ◆ **Live backup** A live backup is a *continuous* backup of the database that helps protect against total machine failure.

☞ For more information on when to use live backups, see "Protecting against total machine failure" on page 318.

☞ For more information on how to make a live backup, see "Making a live backup" on page 340.

Scheduling backups

Most backup schedules involve periodic full backups interspersed with incremental backups of the transaction log. There is no simple rule for deciding how often to make backups of your data. The frequency with which you make backups depends on the importance of your data, how often it changes, and other factors.

Most backup strategies involve occasional full backups, interspersed by several incremental backups. A common starting point for backups is to carry out a weekly full backup, with daily incremental backups of the transaction log. Both full and incremental backups can be carried out online (while the database is running) or offline, on the server side or the client side. Archive backups are always full backups.

The kinds of failure against which a backup schedule protects you depends not only on how often you make backups, but also on how you operate your database server.

🔗 For more information, see "Configuring your database for data protection" on page 317.

You should always keep more than one full backup. If you make a backup on top of a previous backup, a media failure in the middle of the backup leaves you with no backup at all. You should also keep some of your full backups offsite to protect against fire, flood, earthquake, theft, or vandalism.

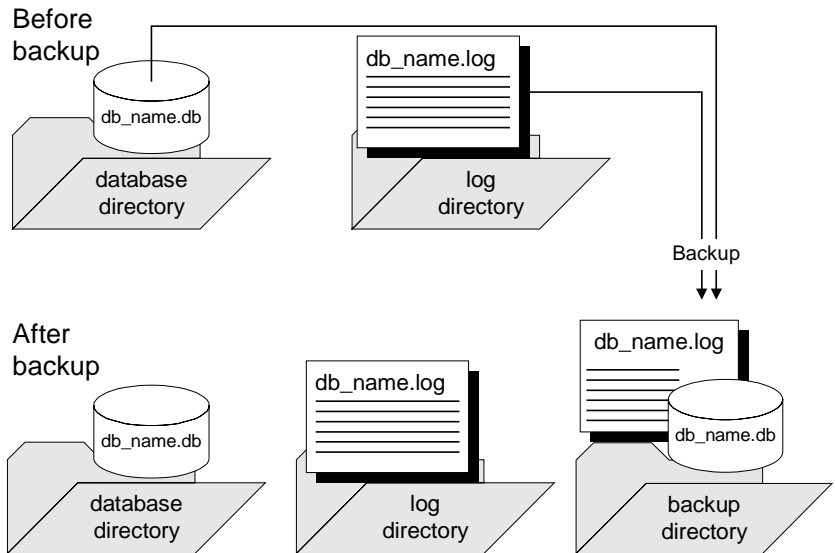
You can use the event scheduling features of Adaptive Server Anywhere to perform online backups automatically at scheduled times.

🔗 For more information on scheduling operations such as backups, see "Automating Tasks Using Schedules and Events" on page 231.

A backup scheme for when disk space is plentiful

If disk space is not a problem on your production machine (where the database server is running) then you do not need to worry about choosing special options to manage the transaction log file. In this case, you can use a simple form of backup that makes copies of the database file and transaction log, and leaves the transaction log in place. All backups leave the database file in place.

A full backup of this kind is illustrated in the figure below. In an incremental backup, only the transaction log is backed up.

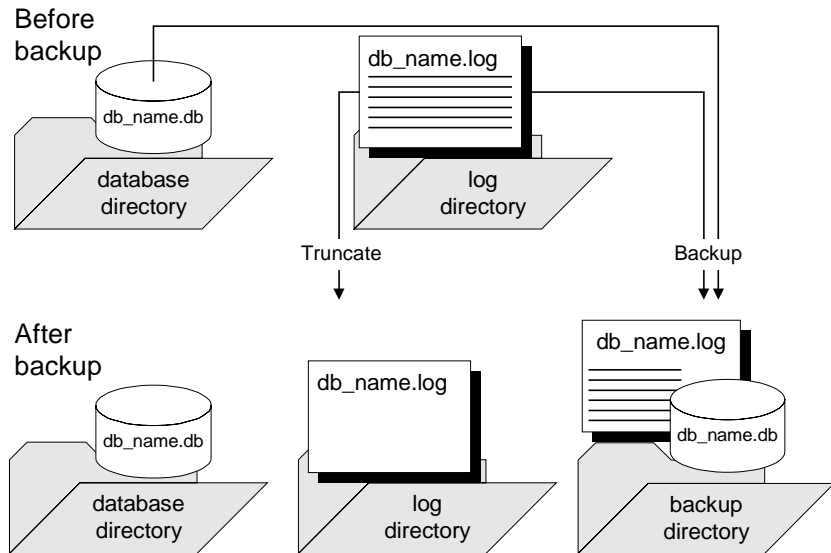


For more information on how to carry out backups of this type, see "Making a backup, continuing to use the original transaction log" on page 333.

A backup scheme for databases not involved in replication

In many circumstances, disk space limitations make it impractical to let the transaction log grow indefinitely. In this case, you can choose to delete the contents of the transaction log when the backup is complete, freeing the disk space. You should not choose this option if the database is involved in replication because replication requires access to the transaction log.

A full backup, which truncates the log file, is illustrated in the figure below. In an incremental backup, only the transaction log is backed up.



Deleting the transaction log after each incremental backup makes recovery from a media failure on the database file a more complex task as there may then be several different transaction logs since the last full backup. Each transaction log needs to be applied in sequence to bring the database up to date.

You can use this kind of backup at a database that is operating as a MobiLink consolidated database because MobiLink does not rely on the transaction log. If you are running SQL Remote or the MobiLink *dbm/sync.exe* application, you must use a scheme suitable for preserving old transaction logs, as described in the following section.

☞ For more information on how to carry out a backup of this type, see "Making a backup, deleting the original transaction log" on page 334.

A backup scheme for databases involved in replication


If your database is part of a SQL Remote installation, the Message Agent needs access to old transactions. If it is a consolidated database, it holds the master copy of the entire SQL Remote installation, and thorough backup procedures are essential.


If your database is a primary site in a Replication Server installation, the Replication Agent requires access to old transactions. However, disk space limitations often make it impractical to let the transaction log grow indefinitely.

If your database is participating in a MobiLink setup, using the *dbmlsync.exe* executable, the same considerations apply. However, if your database is a MobiLink consolidated database, you do not need old transaction logs and can use a scheme for databases not involved in replication, as described in the previous section.

In these cases, you can choose backup options to rename and restart the transaction log. This kind of backup prevents open-ended growth of the transaction log, while maintaining information about the old transactions for the Message Agent and the Replication Agent.

This kind of backup is illustrated in the figure below.



 For more information on how to carry out a backup of this kind, see "Making a backup, renaming the original transaction log" on page 335.

Offline transaction logs

In addition to backing up the transaction log, the backup operation renames the online transaction log to a filename of the form *YYYYMMDDxx.log*. This file is no longer used by the database server, but is available for the Message Agent and the Replication Agent. It is called an **offline** transaction log. A new online transaction log is started with the same name as the old online transaction log.

There is no Year 2000 issue with the two-digit year in the *YYYYMMDDxx.log* filenames. The names are used for distinguishability only, not for ordering. For example, the renamed log file from the first backup on December 10, 2000, is named *001210AA.log*. The first two digits indicate the year, the second two digits indicate the month, the third two digits indicate the day of the month, and the final two characters distinguish among different backups made on the same day.

The Message Agent and the Replication Agent can use the offline copies to provide the old transactions as needed. If you set the `DELETE_OLD_LOGS` database option to ON, then the Message Agent and Replication Agent delete the offline files when they are no longer needed, saving disk space.

Backup methods for remote databases in replication installations

Backup procedures are not as crucial at remote databases as at the consolidated database. You may choose to rely on replication to the consolidated database as a data backup method. In the event of a media failure, the remote database would have to be re-extracted from the consolidated database, and any operations that have not been replicated would be lost. (You could use the Log Translation utility to attempt to recover lost operations.)

Even if you do choose to rely on replication to protect remote database data, backups may still need to be done periodically at remote databases to prevent the transaction log from growing too large. You should use the same option (rename and restart the log) as at the consolidated database, running the Message Agent so that it has access to the renamed log files. If you set the DELETE_OLD_LOGS option to ON at the remote database, the old log files will be deleted automatically by the Message Agent when they are no longer needed.

Automatic
transaction log
renaming

You can use the `-x` Message Agent option to eliminate the need to rename the transaction log on the remote computer when the database server is shut down. The `-x` option renames the transaction log after it has been scanned for outgoing messages.

Backup a database to a tape drive

All the types of backup described above are image backups. The backup copy of each file is also a file. To make backups to tape using an image backup, you have to take each backup copy and put it on tape using a disk backup utility.

You can carry out direct backup to a tape drive using an archive backup. Archive backups are always full backups. An archive backup makes copies of both the database file and the transaction log, but these copies are placed into a single file.

☞ You can make archive backups using the BACKUP statement. For information, see "Backing up a database directly to tape" on page 338, and "BACKUP statement" on page 245 of the book *ASA SQL Reference Manual*.


☞ You can restore the backup using the RESTORE statement. For information, see "Restoring an archive backup" on page 343, and "RESTORE DATABASE statement" on page 511 of the book *ASA SQL Reference Manual*.

Designing a backup and recovery plan

For reliable protection of your data, you should develop and implement a backup schedule. You should also ensure that you have a set of tested recovery instructions.


Typical schedules call for occasional full backups interspersed with several incremental backups. The frequency of each depends on the nature of the data that you are protecting.

If you use internal backups, you can use the scheduling features in Adaptive Server Anywhere to automate the task. Once you specify a schedule, the backups are carried out automatically by the database server.

 For more information on automating backups, see "Automating Tasks Using Schedules and Events" on page 231.

The length of time your organization can function without access to the data in your database imposes a maximum recovery time, and you should develop and test a backup and recovery plan that meets this requirement.


You should verify that you have the protection you need against media failure on the database file and on the transaction log file. If you are running in a replication environment, you should consider using a mirrored transaction log.

 For more information about media failure, see "Protecting your data against media failure" on page 306.

Factors that affect recovery time

External factors such as available hardware, the size of database files, recovery medium, disk space, and unexpected errors can affect your recovery time. When planning a backup strategy, you should allow additional recovery time for miscellaneous tasks that must be performed, such as entering recovery commands or retrieving and loading tapes.

Adding more files into the recovery scenario increases the places where recovery can fail. As the backup and recovery strategy develops, you should consider checking your recovery plan.

 For more information about how to implement a backup and recovery plan, see "Implementing a backup and recovery plan" on page 329.

Ensuring your database is valid

Database file corruption may not be apparent until applications try to access the affected part of the database. As part of your data protection plan, you should periodically check that your database has no errors. You can do this by **validating** the database. This task requires DBA authority.

Database validation includes a scan of every row in every table and a look-up of each row in each index on the table. Validation requires exclusive access to each table in turn. For this reason, it is best to validate when there is no other activity on the database. Database validation does not validate data, continued row references, or foreign key relationships unless you perform a full validation using the `-f` option.

Backup copies of the database and transaction log must not be changed in any way. If there were no transactions in progress during the backup, you can check the validity of the backup database using read-only mode. However, if transactions were in progress, the database server must carry out recovery on the database when you start it. Recovery modifies the backup copy, which is not desirable.

If you can be sure that no transactions are in progress when the backup is being made, the database server does not need to carry out the recovery steps. In this case, you can carry out a validity check on the backup using the read-only database option.

Tip

Using the `BACKUP` statement with the `WAIT BEFORE START` clause ensures that no transactions are in progress when you make a backup.

If a base table in the database file is corrupt, you should treat the situation as a media failure, and recover from your previous backup. If an index is corrupt, you may want to unload the database without indexes, and reload.

🔗 For more information, see "Validating a database" on page 331, and "Validating a transaction log" on page 332.

🔗 For more information on read-only databases, see "`-r` server option" on page 149.

Configuring your database for data protection

There are several ways in which you can configure your database and the database server to provide protection against media failure while maintaining performance.

Protecting against media failure on the database file

When you create a database, by default the transaction log is put on the same device and in the same directory as the database. This arrangement does not protect against all kinds of media failure, and you should consider placing the transaction log in another location for production use.

For comprehensive protection against media failure, you should keep the transaction log on a different device from the database file. Some machines with two or more hard drives have only one physical disk drive with several logical drives or partitions: if you want reliable protection against media failure, make sure that you have a machine with at least two storage devices.

Placing the transaction log on a separate device can also result in improved performance by eliminating the need for disk head movement between the transaction log and the main database file.

You should not place the transaction log on a network directory. Reading and writing pages over a network gives poor performance and may result in file corruption.

☞ For more information on creating databases, see "Creating a database" on page 29 of the book *ASA SQL User's Guide*.

☞ For more information on how to change the location of a transaction log, see "Changing the location of a transaction log" on page 345.

Protecting against media failure on the transaction log

It is recommended that you use a transaction log mirror when running high-volume or extremely critical applications. For example, at a consolidated database in a SQL Remote setup, replication relies on the transaction log, and if the transaction log is damaged or becomes corrupt, data replication can fail.

If you are using a mirrored transaction log, and an error occurs while trying to write to one of the logs (for example, if the disk is full), the database server stops. The purpose of a transaction log mirror is to ensure complete recoverability in the case of media failure on either log device; this purpose would be lost if the server continued with a single log.

Where to store the transaction log mirror

There is a performance penalty for using a mirrored log as each database log write operation must be carried out twice. The performance penalty depends on the nature and volume of database traffic and on the physical configuration of the database and logs.

A transaction log mirror should be kept on a separate device from the transaction log. This improves performance. Also, if either device fails, the other copy of the log keeps the data safe for recovery.

Alternatives to a transaction log mirror

Alternatives to a mirrored transaction log are to use a disk controller that provides hardware mirroring, or operating-system level software mirroring, as provided by Windows NT/2000/XP and NetWare. Generally, hardware mirroring is more expensive, but provides better performance.

For more information

Live backups provide additional protection that has some similarities to transaction log mirroring. For more information, see "Differences between live backups and transaction log mirrors" on page 319.

🔗 For information on creating a database with a mirrored transaction log, see "The Initialization utility" on page 465.

🔗 For information on changing an existing database to use a mirrored transaction log, see "The Transaction Log utility" on page 507.

Protecting against total machine failure

You can use a **live backup** to provide a redundant copy of the transaction log that is available for restart of your system on a secondary machine in case the machine running the database server becomes unusable.

A live backup runs continuously, terminating only if the server shuts down. If you suffer a system failure, the backed up transaction log can be used for a rapid restart of the system. However, depending on the load that the server is processing, the live backup may lag behind and may not contain all committed transactions.

🔗 For more information on making a live backup, see "Making a live backup" on page 340.

🔗 For information on restarting database using a live backup, see "Recovering from a live backup" on page 342.

Differences between live backups and transaction log mirrors

Both a live backup and a transaction log mirror appear to provide a secondary copy of the transaction log. However, there are several differences between using a live backup and using a transaction log mirror:


- ◆ **In general, a live backup is made to a different machine** Running a transaction log mirror on a separate machine is not recommended. It can lead to performance and data corruption problems, and stops the database server if the connection between the machines goes down.

By running the Backup utility on a separate machine, the database server does not do the writing of the backed up log file, and the data transfer is done by the Adaptive Server Anywhere client/server communications system. Therefore, performance impact is decreased and reliability is greater.

- ◆ **A live backup provides protection against a machine becoming unusable** Even if a transaction log mirror is kept on a separate device, it does not provide immediate recovery if the whole machine becomes unusable. You could consider an arrangement where two machines share access to a set of disks.
- ◆ **A live backup may lag behind the database server** A mirrored transaction log contains all the information required for complete recovery of committed transactions. Depending on the load that the server is processing, the live backup may lag behind and may not contain all the committed transactions.

Live backups and regular backups

The live backup of the transaction log is always the same length or shorter than the active transaction log. When a live backup is running, and another backup restarts the transaction log (**dbbackup -r** or **dbbackup -x**), the live backup automatically truncates the live backup log and restarts the live backup at the beginning of the new transaction log.

 For more information about how to make a live backup, see "Making a live backup" on page 340.

Controlling transaction log size

The size of the transaction log can determine what kind of backup is right for you, and can also affect recovery times.

You can control how fast the transaction log file grows by ensuring that all your tables have compact primary keys. If you carry out updates or deletes on tables that do not have a primary key or a unique index not allowing NULL, the entire contents of the affected rows are entered in the transaction log. If a primary key is defined, the database server needs to store only the primary key column values to uniquely identify a row. If the table contains many columns or wide columns, the transaction log pages fill up much faster if no primary key is defined. In addition to taking up disk space, this extra writing of data affects performance.

If a primary key does not exist, the engine looks for a UNIQUE NOT NULL index on the table (or a UNIQUE constraint). A UNIQUE index that allows NULL is not sufficient.

Backup and recovery internals

This section describes the internal mechanisms used during backup and during automatic recovery mechanism from system failures.

Backup internals

When you issue a backup instruction, the database may be in use by many people. If you later need to use your backup to restore your database, you need to know what information has been backed up, and what has not.

The database server carries out a backup as follows:

- 1 Issue a checkpoint. Further checkpoints are disallowed until the backup is complete. While the backup is taking place, any pages modified by other connections are saved before modification in the temporary file, instead of the database file, so that the backup image is made as of the checkpoint.
- 2 Make a backup of the database file, if the backup instruction is for a full backup.
- 3 Make a backup of the transaction log.

The backup includes all operations recorded in the transaction log before the final page of the log is read. This may include instructions issued after the backup instruction was issued.

The backup copy of the transaction log is generally smaller than the online transaction log. The database server allocates space to the online transaction logs in multiples of 64K, so the transaction log file size generally includes empty pages. However, only the non-empty pages are backed up.

- 4 On a database created with version 8.0 or later of Adaptive Server Anywhere, if the backup instruction requires the transaction log to be truncated or renamed, uncommitted transactions are carried forward to the new transaction log.

On a database created with version 7.x or earlier of Adaptive Server Anywhere, if the backup instruction requires the transaction log to be truncated or renamed, then the database server waits until there are no uncommitted transactions before truncating or renaming the log file. If the database is busy, this wait may be significant.

🔗 For more information about renaming and truncating the transaction log, see "Designing backup procedures" on page 308.

- 5 Mark the backup image of the database to indicate that recovery is needed. This causes any operations that happened since the start of the backup to be applied. It also causes operations that were incomplete at the checkpoint to be undone if they were not committed.

Restrictions during backup and recovery

The database server prevents the following operations from being executed while a backup is in progress:

- ◆ Another backup, with the exception of a live backup.
- ◆ A checkpoint, other than the one issued by the backup instruction itself.
- ◆ Any statement that causes a checkpoint. This includes data definition statements as well as the LOAD TABLE and TRUNCATE TABLE statements.


During recovery, including restoring backups, no action is permitted by other users of the database.

Checkpoints and the checkpoint log


The database file is composed of pages: fixed size portions of hard disk. The checkpoint log is located at the end of the database file. Pages are added to the checkpoint log as necessary during a session, and the entire checkpoint log is deleted at the end of the session.

Before any page is updated (made **dirty**), the database server carries out the following operations:


- ◆ It reads the page into memory, where it is held in the database cache.
- ◆ It makes a copy of the original page. These copied pages are the **checkpoint log**.



Changes made to the page are applied to the copy in the cache. For performance reasons they are not written immediately to the database file on disk.



When the cache is full, the changed page may get written out to disk. The copy in the checkpoint log remains unchanged.



A **checkpoint** is a point at which all dirty pages are written to disk. Following a checkpoint, the contents of the checkpoint log are deleted. The empty checkpoint log pages remain in the checkpoint log within a given session and can be reused for new checkpoint log data. As the checkpoint log increases in size, so does the database file.

At a checkpoint, all the data in the database is held on disk in the database file. The information in the database file matches that in the transaction log. The checkpoint represents a known consistent state of the database on disk. During recovery, the database is first recovered to the most recent checkpoint, and then changes since that checkpoint are applied.

The entire checkpoint log, including all empty checkpoint log pages, is deleted at the end of each session. Deleting the checkpoint log causes the database to shrink in size.

☞ For more information, see "How the database server decides when to checkpoint" on page 325.

Transactions and the rollback log

As changes are made to the contents of a database, a **rollback log** is kept for the purpose of canceling changes if a transaction is rolled back or if a transaction is uncommitted when a system failure occurs. There is a separate rollback log for each connection. When a transaction is committed or rolled back, the rollback log contents for that connection are deleted. The rollback logs are stored in the database, and rollback log pages are copied into the checkpoint log along with other pages that are changed.

The rollback log is also called the **undo log**.

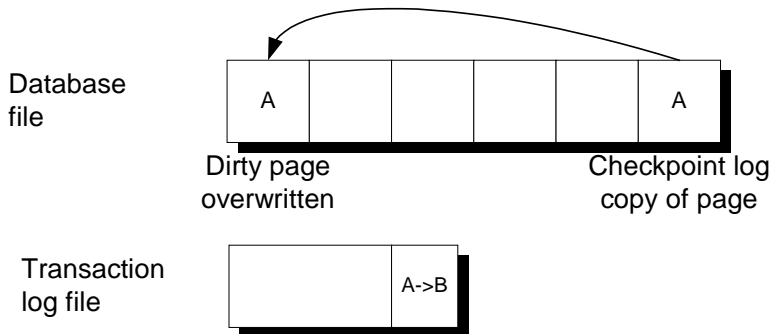
For more information about transaction processing, see "Using Transactions and Isolation Levels" on page 89 of the book *ASA SQL User's Guide*.

The automatic recovery process

When a database is shut down during normal operation, the database server carries out a checkpoint so that all the information in the database is held in the database file. This is a **clean** shutdown.

Each time you start a database, the database server checks whether the last shutdown was clean or the result of a system failure. If the database was not shut down cleanly, it automatically takes the following steps to recover from a system failure:

- 1 **Recover to the most recent checkpoint** All pages are restored to their state at the most recent checkpoint by copying the checkpoint log pages over the changes made since the checkpoint.



- 2 **Apply changes made since the checkpoint** Changes made between the checkpoint and the system failure, which are held in the transaction log, are applied.
- 3 **Rollback uncommitted transactions** Any uncommitted transactions are rolled back, using the rollback logs.

How the database server decides when to checkpoint

The priority of writing dirty pages to the disk increases as the time and the amount of work since the last checkpoint grows. The priority is determined by the following factors:

- ◆ **Checkpoint Urgency** The time that has elapsed since the last checkpoint, as a percentage of the checkpoint time setting of the database. The server `-gc` option controls the maximum desired time, in minutes, between checkpoints. You can also set the desired time using the `CHECKPOINT_TIME` option.

☞ For more information, see "`-gc` server option" on page 138.

- ◆ **Recovery Urgency** A heuristic to estimate the amount of time required to recover the database if it fails right now. The server `-gr` option controls the maximum desired time, in minutes, for recovery in the event of system failure. You can also set the desired time using the `RECOVERY_TIME` option.

☞ For more information, see "`-gr` server option" on page 143.

The checkpoint and recovery urgencies are important only if the server does not have enough idle time to write dirty pages.

Frequent checkpoints make recovery quicker, but also create work for the server writing out dirty pages.

There are two database options that allow you to control the frequency of checkpoints. `CHECKPOINT_TIME` controls the maximum desired time between checkpoints and `RECOVERY_TIME` controls the maximum desired time for recovery in the event of system failure.

The writing of dirty pages to disk is carried out by a task within the server called the **idle I/O task**. This task shares processing time with other database tasks.

There is a threshold for the number of dirty pages, below which writing of database pages does not take place.

When the database is busy, the urgency is low, and the cache only has a few dirty pages, the idle I/O task runs at a very low priority and no writing of dirty pages takes place.

Once the urgency exceeds 30%, the priority of the idle I/O task increases. At intervals, the priority is increased again. As the urgency becomes high, the engine shifts its primary focus to writing dirty pages until the number gets below the threshold again. However, the engine only writes out pages during the idle I/O task if the number of dirty pages is greater than the threshold.

If, because of other activity in the database, the number of dirty pages falls to zero, and if the urgency is 50% or more, then a checkpoint takes place automatically since it is a convenient time.

Both the checkpoint urgency and recovery urgency values increase in value until the checkpoint occurs, at which point they drop to zero. They do not decrease otherwise.

Validating the transaction log on database startup

When a database using a transaction log mirror starts up, the database server carries out a series of checks and automatic recovery operations to confirm that the transaction log and its mirror are not corrupted, and to correct some problems if corruption is detected.

On startup, the server checks that the transaction log and its mirror are identical by carrying out a full comparison of the two files; if they are identical, the database starts as usual. The comparison of log and mirror adds to database startup time.

If the database stopped because of a system failure, it is possible that some operations were written into the transaction log but not into the mirror. If the server finds that the transaction log and the mirror are identical up to the end of the shorter of the two files, the remainder of the longer file is copied into the shorter file. This produces an identical log and mirror. After this automatic recovery step, the server starts as usual.

If the check finds that the log and the mirror are different in the body of the shorter of the two, one of the two files is corrupt. In this case, the database does not start, and an error message is generated saying that the transaction log or its mirror is invalid.

Improving performance when validating databases

The `VALIDATE TABLE` statement can be slow when used on large databases running on servers with a cache size too small to contain the table and its largest index. It is often the case that all pages in the table are read at least once for each index. As well, if full compares are required for index lookups, the number of page reads can be proportional to the number of rows (not pages) in the table.

If you want to reduce the time taken to validate, you can use the `WITH EXPRESS CHECK` option with the `VALIDATE TABLE` statement, or the `-fx` option with the `dbvalid` utility. Depending on the size of your database, the size of your cache, and the type of validation you require, these two features can significantly reduce the time taken to perform validation.

Express validation causes each row of the table to be read and all columns evaluated, as is done with traditional validation using the `WITH DATA CHECK` clause. Each index is completely scanned once, and checks are done to ensure that the rows referenced in the index exist in the table. The Express check option also does checks on the validity of individual index pages. The number of rows in the table must match the number of entries in the index. The Express option saves time because it does not perform individual index lookups for each row.

Because the Express check feature does not perform individual lookups, it is possible (though unlikely) for some form of index corruption to go unnoticed by the Express validation feature. If index corruption should occur, data can be recovered by unloading and rebuilding the database since validation has confirmed that all of the data can be read.

Express validation is only supported for databases created with ASA 8.0 or later.


☞ For more information about the Express check option, see the "VALIDATE TABLE statement" on page 586 of the book *ASA SQL Reference Manual*, "Validating a database using the dbvalid command-line utility" on page 527, and the "sa_validate system procedure" on page 720 of the book *ASA SQL Reference Manual*.

Backup and recovery tasks

This section collects together instructions for tasks related to backup and recovery.


Implementing a backup and recovery plan

Regular backups and tested recovery commands are part of a comprehensive backup and recovery plan.

 For more information, see "Designing a backup and recovery plan" on page 315.


❖ To implement a backup and recovery plan:

- 1 Create and verify your backup and recovery commands.
- 2 Measure the time it takes to execute backup and recovery commands.
- 3 Document the backup commands and create written procedures outlining where your backups are kept and identify any naming convention used, as well as the kind of backups performed.
- 4 Set up your backup procedures on the production server.
- 5 Monitor backup procedures to avoid unexpected errors. Make sure any changes in the process are reflected in your documentation.

 For more information about carrying out backups, see "Making a full backup" on page 329, and "Making an incremental backup" on page 330.


Making a full backup

A full backup is a backup of the database file and the transaction log file.

 For more information about the difference between a full backup and an incremental backup, see "Types of backup" on page 308.

❖ To make a full backup (overview):

- 1 Ensure that you have DBA authority on the database.
- 2 Perform a validity check on your database to ensure that it is not corrupt. You can use the Validation utility or the *sa_validate* stored procedure.

 For more information, see "Validating a database" on page 331.

- 3 Make a backup of your database file and transaction log.

☞ For information on how to carry out the backup operation, see the following:

- ◆ "Making a backup, continuing to use the original transaction log" on page 333.
- ◆ "Making a backup, deleting the original transaction log" on page 334.
- ◆ "Making a backup, renaming the original transaction log" on page 335.

Notes

Validity checking requires exclusive access to entire tables on your database. For more information and alternative approaches, see "Ensuring your database is valid" on page 315.

If you validate your backup copy of the database, make sure that you do so in read-only mode. Start the database server with the `-r` option to use read-only mode.

Making an incremental backup

An incremental backup is a backup of the transaction log file only. Typically, you should make several incremental backups between each full backup.

☞ For more information about the difference between a full backup and an incremental backup, see "Types of backup" on page 308.

❖ To make an incremental backup (overview):

- 1 Ensure that you have DBA authority on the database.
- 2 Make a backup of your transaction log only, not your database file.

☞ For more information about how to carry out the backup operation, see the following:

- ◆ "Making a backup, continuing to use the original transaction log" on page 333.
- ◆ "Making a backup, deleting the original transaction log" on page 334.
- ◆ "Making a backup, renaming the original transaction log" on page 335.

Notes

The backup copies of the database file and transaction log file have the same names as the online versions of these files. For example, if you make a backup of the sample database, the backup copies are called *asademo.db* and *asademo.log*. When you repeat the backup statement, choose a new backup directory to avoid overwriting the backup copies.

☞ For more information about how to make a repeatable incremental backup command by renaming the backup copy of the transaction log, see "Renaming the backup copy of the transaction log during backup" on page 338.

Validating a database

Validating a database is a key part of the backup operation. For information, see "Ensuring your database is valid" on page 315.

For an overview of the backup operation, see "Making a full backup" on page 329.

❖ To check the validity of an entire database (Sybase Central):

- 1 Connect to the database as a user with DBA authority.
- 2 Right-click the database and choose Validate Database from the popup menu.
- 3 Follow the instructions in the wizard.

A message box indicates whether the database is valid or not.

❖ To check the validity of an entire database (SQL):

- 1 Connect to the database as a user with DBA authority.
- 2 Execute the *sa_validate* stored procedure:

```
call sa_validate
```

The procedure returns a single column, named *Messages*. If all tables are valid, the column contains No errors detected.

☞ For more information, see "sa_validate system procedure" on page 720 of the book *ASA SQL Reference Manual*.

❖ To check the validity of an entire database (Command line):

- 1 Connect to the database as a user with DBA authority.
- 2 Run the *dbvalid* utility:

```
dbvalid -c "connection_string"
```

☞ For more information, see "The Validation utility" on page 526.

Notes

If you are checking the validity of a backup copy, you should run the database in read-only mode so that it is not modified in any way. You can only do this when there were no transactions in progress during the backup.

☞ For information on running databases in read-only mode, see "-r server option" on page 149.

Validating a single table

You can check the validity of a single table either from Sybase Central or using a SQL statement. You must have DBA authority or be the owner of a table to check its validity.

❖ To check the validity of a table (Sybase Central):

- 1 Open the Tables folder.
- 2 Right-click the table and choose Validate from the popup menu.

A message box indicates whether the table is valid or not.

❖ To check the validity of a table (SQL):

- ◆ Execute the `VALIDATE TABLE` statement:

```
VALIDATE TABLE table_name
```

Notes

If errors are reported, you can drop all of the indexes and keys on a table and recreate them. Any foreign keys to the table will also need to be recreated. Another solution to errors reported by `VALIDATE TABLE` is to unload and reload your entire database. You should use the `-u` option of `dbunload` so that it does not try to use a possibly corrupt index to order the data.

Validating a transaction log

The process for validating a transaction log file depends on whether it is in use on a production database (online) or is not in use (offline, or a backup copy).

❖ To validate an online transaction log:


- ◆ Run your database with a mirrored transaction log. The database server automatically validates the transaction log each time the database is started.

❖ To validate an offline or backed up transaction log:

- ◆ Run the Log Translation utility (`dbtran`) against the log file. If the Log Translation utility can successfully read the log file, it is valid.

Making a backup, continuing to use the original transaction log

This task describes the simplest kind of backup, which leaves the transaction log untouched.

 For more information about when to use this type of backup, see "A backup scheme for when disk space is plentiful" on page 310.

❖ To make a backup, continuing to use the original transaction log (Sybase Central):

- 1 Start Sybase Central. Connect to the database as a user with DBA authority.
- 2 Right-click the database and choose Create Backup Images from the popup menu. The Create Backup Images wizard appears.
- 3 Click Next on the introductory page of the wizard.
- 4 Select the database that you want to back up.
- 5 On the next page, enter the name of a directory to hold the backup copies, and choose whether to carry out a complete backup (all database files) or an incremental backup (transaction log file only).
- 6 On the next page, select the option Continue To Use The Same Transaction Log.
- 7 Click Finish to start the backup.

The procedure describes a client-side backup. There are more options available for this kind of backup.

If you choose a server-side backup, and the server is running on a different machine from Sybase Central, you cannot use the Browse button to locate a directory in which to place the backups. The Browse button browses the client machine, while the backup directory is relative to the server.

❖ To make a backup, continuing to use the original transaction log (SQL):

- ◆ If you are using the BACKUP statement, use the following clauses only:

```
BACKUP DATABASE
  DIRECTORY directory_name
  [ TRANSACTION LOG ONLY ]
```

Include the TRANSACTION LOG ONLY clause if you are making an incremental backup.

❖ **To make a backup, continuing to use the original transaction log (Command line):**

- ◆ If you are using the *dbbackup* utility, use the following syntax:

```
dbbackup -c "connection_string" [ -t ] backup_directory
```

Include the `-t` option only if you are making an incremental backup.

Making a backup, deleting the original transaction log

If your database is not involved in replication, and if you have limited disk space on your online machine, you can delete the contents of the online transaction log (**truncate** the log) when you make a backup. In this case, you need to use every backup copy made since the last full backup during recovery from media failure on the database file.

☞ For more information about when to use this type of backup, see "A backup scheme for databases not involved in replication" on page 311.

❖ **To make a backup, deleting the transaction log (Sybase Central):**

- 1 Start Sybase Central. Connect to the database as a user with DBA authority.
- 2 Right-click the database and choose Create Backup Images from the popup menu. The Create Backup Images wizard appears.
- 3 Click Next on the introductory page of the wizard.
- 4 Select the database that you want to back up.
- 5 On the next page, enter the name of a directory to hold the backup copies, and choose whether to carry out a complete backup (all database files) or an incremental backup (transaction log file only).
- 6 On the next page, select the option Truncate The Transaction Log.
- 7 Click Finish to start the backup.

❖ **To make a backup, deleting the transaction log (SQL):**

- ◆ Use the BACKUP statement with the following clauses:

```
BACKUP DATABASE  
  DIRECTORY backup_directory  
  [ TRANSACTION LOG ONLY ]  
TRANSACTION LOG TRUNCATE
```

Include the TRANSACTION LOG ONLY clause only if you are making an incremental backup.

The backup copies of the transaction log and database file are placed in *backup_directory*. If you enter a path, it is relative to the working directory of the database server, not your client application.

❖ **To make a backup, deleting the transaction log (Command line):**

◆ From the command prompt, enter the following command:

```
dbbackup -c "connection_string" -x [ -t ] backup_directory
```

Include the `-t` option only if you are making an incremental backup.

The backup copies of the transaction log and database file are placed in *backup_directory*. If you enter a path, it is relative to the directory from which you run the command.

Notes

Before the online transaction log can be erased, all outstanding transactions must terminate. If there are outstanding transactions, your backup cannot complete. For information, see "Backup internals" on page 321.

You can determine which connection has an outstanding transaction using the *sa_conn_info* system procedure. If necessary, you can disconnect the user with a `DROP CONNECTION` statement.

☞ For more information, see "Determining which connection has an outstanding transaction" on page 337.

Making a backup, renaming the original transaction log

This set of backup options is typically used for databases involved in replication. In addition to making backup copies of the database file and transaction log, the transaction log at backup time is renamed to an offline log, and a new transaction log is started with the same name as the log in use at backup time.

☞ For more information about when to use this set of backup options, see "A backup scheme for databases involved in replication" on page 312.

❖ **To make a backup, renaming the transaction log (Sybase Central):**

- 1 Start Sybase Central. Connect to the database as a user with DBA authority.
- 2 Right-click the database and choose Create Backup Images from the popup menu. The Create Backup Images wizard appears.
- 3 Click Next on the introductory page of the wizard.
- 4 Select the database that you want to back up.

- 5 On the next page, enter the name of a directory to hold the backup copies, and choose whether to carry out a complete backup (all database files) or an incremental backup (transaction log file only).
- 6 On the next page, select the option Rename The Transaction Log.
- 7 Click Finish to start the backup.

❖ **To make a backup, renaming the transaction log (SQL):**

- ◆ Use the BACKUP statement with the following clauses:

```
BACKUP DATABASE  
    DIRECTORY backup_directory  
    [ TRANSACTION LOG ONLY ]  
    TRANSACTION LOG RENAME
```

Include the TRANSACTION LOG ONLY clause only if you are making an incremental backup.

The backup copies of the transaction log and database file are placed in *backup_directory*. If you enter a path, it is relative to the working directory of the database server, not your client application.

❖ **To make a backup, renaming the transaction log (Command line):**

- ◆ From the command prompt, enter the following command. You must enter the command on a single line:

```
dbbackup -c "connection_string" -r [ -t ] backup_directory
```

Include the -t option if you are making an incremental backup.

The backup copies of the transaction log and database file are placed in *backup_directory*. If you enter a path, it is relative to the directory from which you run the command.

Notes

Before the online transaction log can be renamed, all outstanding transactions must terminate. If there are outstanding transactions, your backup will not complete. For information, see "Backup internals" on page 321.

You can determine which connection has an outstanding transaction using the *sa_conn_info* system procedure. If necessary, you can disconnect the user with a DROP CONNECTION statement.

🔗 For more information, see "Determining which connection has an outstanding transaction" on page 337.

Determining which connection has an outstanding transaction

If you are carrying out a backup that renames or deletes the transaction log on a database created with version 8.0 or later of Adaptive Server Anywhere, incomplete transactions are carried forward to the new transaction log.


If you are carrying out a backup that renames or deletes the transaction log on a database created with version 7.x or earlier of Adaptive Server Anywhere, however, and if there are outstanding transactions, the backup must wait until those transactions are complete before it can complete.

You can use a system procedure to determine which user has outstanding transactions. If there are not too many connections, you can also use the Console utility to determine which connection has outstanding connections.

❖ To determine which connection has an outstanding transaction (SQL):

- 1 Connect to the database from Interactive SQL or another application that can call stored procedures.
- 2 Execute the *sa_conn_info* system procedure:


```
call sa_conn_info
```
- 3 Inspect the **UncmtOps** column to see which connection has uncommitted operations.

 For more information, see "sa_conn_info system procedure" on page 688 of the book *ASA SQL Reference Manual*.

❖ To determine which connection has an outstanding transaction (Console utility):

- 1 Connect to the database from the Console utility.
For example, the following command connects to the default database using user ID DBA and password SQL:

```
dbconsole -c "uid=DBA;pwd=SQL"
```

 For more information, see "The Console utility" on page 450.

- 2 Double-click each connection, and inspect the Uncommitted Ops entry to see which users have uncommitted operations. If necessary, you can disconnect the user to enable the backup to finish.

Renaming the backup copy of the transaction log during backup

By default, the backup copy of the transaction log file has the same name as the online file. For each backup operation, you must assign a different name or location for the backup copy, or you must move the backup copy before the next backup is done.

You can make a repeatable incremental backup command by renaming the backup copy of the transaction log.

❖ To rename the backup copy of the transaction log (SQL):

- ◆ Use the **MATCH** keyword in the **BACKUP** statement. For example, the following statement makes an incremental backup of the *asademo* database to the directory *c:\backup*. The backup copy of the transaction log is called *YYMMDDxx.log*, where *YYMMDD* is the date and *xx* is a counter, starting from *AA*.

```
BACKUP DATABASE
DIRECTORY 'c:\\backup'
TRANSACTION LOG ONLY
TRANSACTION LOG RENAME MATCH
```

❖ To rename the backup copy of the transaction log (Command line):

- ◆ Supply the **-n** option to *dbbackup*. For example, the following command makes an incremental backup of the sample database, renaming the backup copy of the transaction log.

```
dbbackup -c "uid=DBA;pwd=SQL;dbn=asademo" -r -t -n c:\backup
```

Notes

The backup copy of the transaction log is named *YYMMDDxx.log*, where *YY* is the year, *MM* is the month, *DD* is the day of the month, and *xx* runs from *AA* to *ZZ*, incrementing if there is more than one backup per day. There is no Year 2000 issue with the two-digit year in the *YYMMDDxx.log* filenames. The names are used for distinguishability only, not for ordering.

Backing up a database directly to tape

An archive backup makes a copy of the database file and transaction log file in a single archive destination. Only server-side, full backups can be made in this manner. When you make an archive backup in Sybase Central, you have the option of backing up the database directly to tape or to disk.

☞ For more information, see "Types of backup" on page 308.

❖ **To make an archive backup to tape (Sybase Central):**

- 1 Start Sybase Central. Connect to the database as a user with DBA authority.
- 2 Right-click the database and choose Backup Database from the popup menu. The Backup Database wizard appears.
- 3 Click Next on the introductory page of the wizard.
- 4 Select the database that you want to back up to tape.
- 5 On the second page of the wizard, select the On Tape To The Following Device option and type the tape drive in the text box below.
- 6 Click Finish to start the backup.

❖ **To make an archive backup to tape (SQL):**

- 1 Use the BACKUP statement, with the following clauses:

```
BACKUP DATABASE
TO archive_root
[ ATTENDED { ON | OFF } ]
[ WITH COMMENT comment string ]
```

If you set the ATTENDED option to OFF, the backup fails if it runs out of tape or disk space. If ATTENDED is set to ON, you are prompted to take an action, such as replacing the tape, when there is no more space on the backup archive device.

Notes

The BACKUP statement makes an entry in the text file *backup.syb*, in the same directory as the server executable.

☞ For more information about restoring from an archive backup, see "Restoring an archive backup" on page 343.

Example

The following statement makes a backup to the first tape drive on a Windows NT machine:

```
BACKUP DATABASE
TO '\\\\.\tape0'
ATTENDED OFF
WITH COMMENT 'May 6 backup'
```

The first tape drive on Windows NT is `\\.\tape0`. Because the backslash is an escape character in SQL strings, each backslash is preceded by another.

The following statement makes an archive file on disk named *c:\backup\archive.1*.

```
BACKUP DATABASE
TO 'c:\\backup\\archive'
```

Making a live backup

You carry out a live backup of the transaction log using the *dbbackup* utility with the `-l` option.

☞ For more information about live backups, see "Protecting against total machine failure" on page 318.

❖ To make a live backup:

- 1 Set up a secondary machine from which you can run the database if the online machine fails. For example, ensure that you have Adaptive Server Anywhere installed on the secondary machine.
- 2 Periodically, carry out a full backup to the secondary machine.
- 3 Run a live backup of the transaction log to the secondary machine.

```
dbbackup -l path\filename.log -c "connection_string"
```

You should normally run the *dbbackup* utility from the secondary machine.

If the primary machine becomes unusable, you can restart your database using the secondary machine. The database file and the transaction log hold the information needed to restart.

Recovering from media failure on the database file

The recovery process depends on whether you leave the transaction log untouched on incremental backup in your backup process. If your backup operation deletes or renames the transaction log, you may have to apply changes from several transaction logs. If your backup operation leaves the transaction log untouched, you need to use only the online transaction log in recovery.

☞ For more information about the backup types discussed here, see "Designing backup procedures" on page 308.

❖ To recover from media failure on the database file:

- 1 Make an extra backup copy of the current transaction log. The database file is gone, and the only record of changes since the last backup is in the transaction log.
- 2 Create a **recovery directory** to hold the files you use during recovery.
- 3 Copy the database file from the last full backup to the recovery directory.

- 4 Apply the transactions held in the backed up transaction logs to the recovery database.

For each log file, in chronological order,

- ◆ Copy the log file into the recovery directory.
- ◆ Start the database server with the apply transaction log (-a) option, to apply the transaction log:

```
dbeng8 db_name.db -a log_name.log
```

The database server shuts down automatically once the transactions have been applied.

- 5 Copy the online transaction log into the recovery directory. Apply the transactions from the online transaction log to the recovery database.

```
dbeng8 db_name.db -a db_name.log
```

- 6 Perform validity checks on the recovery database.
- 7 Make a post-recovery backup.
- 8 Move the database file to the production directory.
- 9 Allow user access to the production database.

Recovering from media failure on an unmirrored transaction log

If your database is a primary site in a Replication Server installation, or a consolidated database in a SQL Remote installation, you should use a mirrored transaction log, or hardware equivalent.

🔗 For more information, see "Protecting against media failure on the transaction log" on page 317.

❖ To recover from media failure on an unmirrored transaction log (partial recovery):

- 1 Make an extra backup copy of the database file immediately. With the transaction log gone, the only record of the changes between the last backup and the most recent checkpoint is in the database file.
- 2 Delete or rename the transaction log file.
- 3 Restart the database with the -f option.

```
dbeng8 asademo.db -f
```

Caution

This command should only be used when the database is not participating in a SQL Remote or Replication Server replication system. If your database is a consolidated database in a SQL Remote replication system, you may have to re-extract the remote databases.

Without the `-f` option, the server reports the lack of a transaction log as an error. With the option, the server restores the database to the most recent checkpoint and then rolls back any transactions that were not committed at the time of the checkpoint. A new transaction log is then created.

Recovering from media failure on a mirrored transaction log

❖ **To recover from media failure on a mirrored transaction log:**

- 1 Make an extra copy of the backup of your database file taken at the time the transaction log was started.
- 2 Identify which of the two files is corrupt. Run the Log Translation utility on the transaction log and on its mirror to see which one generates an error message. The Log Translation utility is accessible from Sybase Central or as the `dbtran` utility.

The following command line translates a transaction log named `asademo.log`, placing the translated output into `asademo.SQL`:

```
dbtran asademo.log
```

The Log Translation utility properly translates the intact file, and reports an error while translating the corrupt file.

- 3 Copy the correct file over the corrupt file so that you have two identical files again.
- 4 Restart the server.

Recovering from a live backup

A live backup is made to a separate machine from the primary machine that is running your production database. To restart a database from a live backup, you must have Adaptive Server Anywhere installed on the secondary machine.

☞ For more information about live backups, see "Protecting against total machine failure" on page 318.

❖ **To restart a database using a live backup:**

- 1 Start the database server on the secondary machine with the apply transaction log (-a) option to apply the transaction log and bring the database up to date:

```
dbeng8 asademo.db -a filename.log
```

The database server shuts down automatically once the transaction log is applied.

- 2 Start the database server in the normal way, allowing user access. Any new activity is appended to the current transaction log.

Restoring an archive backup

If you use an archive backup (typically to tape), you use the RESTORE statement to recover your data.

☞ For more information about making archive backups, see "Backing up a database directly to tape" on page 338.

❖ **To restore a database from an archive backup (Sybase Central):**

- 1 In Sybase Central, connect to a database with DBA authority.
- 2 Open the Utilities folder (located within the server folder).
- 3 In the right pane, double-click Restore Database.
- 4 Follow the instructions in the wizard.

❖ **To restore a database from an archive backup (Interactive SQL):**

- 1 Start a personal database server. Use a command such as the following, which starts a server named *restore*:

```
dbeng8 -n restore
```

- 2 Start Interactive SQL. On the Identification tab of the Connect dialog, enter a user ID of **DBA** and a password of **SQL**. Leave all other fields on this tab blank.
- 3 Click the Database tab and enter a database name of **utility_db**. Leave all other fields on this tab blank.
- 4 Click OK to connect.

- 5 Execute the RESTORE statement, specifying the archive root. At this time, you can choose to restore an archived database to its original location (default), or to a different machine with different device names using the RENAME clause.

🔗 For more information, see "RESTORE DATABASE statement" on page 511 of the book *ASA SQL Reference Manual*.

Example

The following statement restores a database from a tape archive to the database file `c:\newdb\newdb.db`.

```
RESTORE DATABASE 'c:\newdb\newdb.db'  
FROM '\\.\tape0'
```

The following statement restores a database from an archive backup in file `c:\backup\archive.1` to the database file `c:\newdb\newdb.db`. The transaction log name and location are specified in the database.

```
RESTORE DATABASE 'c:\newdb\newdb.db'  
FROM 'c:\backup\archive'
```

🔗 For more information, see "RESTORE DATABASE statement" on page 511 of the book *ASA SQL Reference Manual*.

Recovering uncommitted operations

When recovering from media failure on the database file, the transaction log is intact. Recovery reapplies all committed transactions to the database. In some circumstances, you may wish to find information about transactions that were incomplete at the time of the failure.

❖ To recover uncommitted operations from a transaction log (Sybase Central):

- 1 Do one of the following:
 - ◆ If you are connected to a database, open the server for that database and then open the Utilities folder. In the right pane, double-click Translate Log File.
 - ◆ If you are not connected to a database, click Tools ► Adaptive Server Anywhere 8 ► Translate Log File.
- 2 Follow the instructions in the wizard.
- 3 Edit the translated log (SQL command file) in a text editor and identify the instructions you need.

❖ **To recover uncommitted operations from a transaction log (Command line):**

- 1 Run *dbtran* to convert the transaction log into a SQL command file, using the -a option to include uncommitted transactions. For example, the following command uses *dbtran* to convert a transaction log:

```
dbtran -a sample.log changes.SQL
```

- 2 Edit the translated log (SQL command file) in a text editor and identify the instructions you need.

🔗 For more information on the Log Translation utility, see "The Log Translation utility" on page 488.

Note

The transaction log may or may not contain changes right up to the point where a failure occurred. It does contain any changes made before the end of the most recently committed transaction that made changes to the database.

Changing the location of a transaction log

The database must not be running when you change the location of a transaction log.

🔗 For more information about how to choose the location of a transaction log, see "Protecting against media failure on the database file" on page 317.

❖ **To change the location of a transaction log (Sybase Central):**

- 1 Do one of the following:
 - ◆ If you are already connected to the database associated with the log file, open the server for that database and then open the Utilities folder. In the right pane, double-click Change Log File Settings.
 - ◆ If you are not connected to the database, click Tools►Adaptive Server Anywhere 8►Change Log File Settings.
- 2 Follow the instructions in the wizard.

❖ **To start a transaction log mirror for an existing database (Command line):**

- 1 Ensure that the database is not running.
- 2 Enter the following command at the command prompt:

```
dblog -t new-log-file database-file
```

🔗 For more information about *dblog* options, see "Transaction log utility options" on page 509.

Creating a database with a transaction log mirror

You can choose to maintain a transaction log mirror when you create a database. This option is available either from the CREATE DATABASE statement, from Sybase Central, or from the *dbinit* utility.

☞ For more information about why you may wish to use a transaction log mirror, see "Protecting against media failure on the transaction log" on page 317.

❖ To create a database that uses a transaction log mirror (Sybase Central):

- 1 Do one of the following:
 - ◆ If you are connected to a database, open the server for that database and then open the Utilities folder. In the right pane, double-click Create Database.
 - ◆ If you are not connected to a database, click Tools►Adaptive Server Anywhere 8►Create Database.
- 2 Follow the instructions in the wizard.

❖ To create a database that uses a transaction log mirror (SQL):

- ◆ Use the CREATE DATABASE statement, with the TRANSACTION LOG clause.

☞ For more information, see "CREATE DATABASE statement" on page 273 of the book *ASA SQL Reference Manual*.

❖ To create a database that uses a transaction log mirror (Command line):

- ◆ Use the *dbinit* command with the *-m* option. For example, the following command (which should be entered on one line) initializes a database named *company.db*, with a transaction log kept on a different device and a mirror on a third device.

```
dbinit -t d:\log_dir\company.log -m  
e:\mirr_dir\company.mlg c:\db_dir\company.db
```

☞ For more information about initialization options, see "Initialization utility options" on page 467.

Starting a transaction log mirror for an existing database

Using the Transaction Log utility, you can choose to maintain a transaction log mirror for an existing database any time the database is not running. This option is available from either Sybase Central or the *dblog* utility.

☞ For more information about why you may wish to use a transaction log mirror, see "Protecting against media failure on the transaction log" on page 317.

❖ To start a transaction log mirror for an existing database (Sybase Central):

- 1 Do one of the following:
 - ◆ If you are already connected to the database associated with the log mirror, open the server for that database and then open the Utilities folder. In the right pane, double-click Change Log File Settings.
 - ◆ If you are not connected to the database, click Tools►Adaptive Server Anywhere 8►Change Log File Settings.
- 2 Follow the instructions in the wizard.

❖ To start a transaction log mirror for an existing database (Command line):

- 1 Ensure that the database is not running.
- 2 Enter the following command at the command prompt:

```
dblog -m mirror-file database-file
```

☞ For more information about *dblog* options, see "Transaction log utility options" on page 509.

You can also use the *dblog* utility and Sybase Central to stop a database from using a transaction log mirror.

PART THREE

Permissions, Security, and Replication

This section describes how to use user IDs and permissions to maintain a secure database. It also describes how to replicate data with the Replication Server.

C H A P T E R 1 2

Managing User IDs and Permissions

About this chapter Each user of a database must have a name they type when connecting to the database, called a user ID. This chapter describes how to manage user IDs.

Contents	<table><tr><th>Topic</th><th>Page</th></tr><tr><td>Database permissions overview</td><td>352</td></tr><tr><td>Setting user and group options</td><td>356</td></tr><tr><td>Managing individual user IDs and permissions</td><td>357</td></tr><tr><td>Managing connected users</td><td>368</td></tr><tr><td>Managing groups</td><td>369</td></tr><tr><td>Database object names and prefixes</td><td>376</td></tr><tr><td>Using views and procedures for extra security</td><td>378</td></tr><tr><td>Changing Ownership on Nested Objects</td><td>381</td></tr><tr><td>How user permissions are assessed</td><td>383</td></tr><tr><td>Managing the resources connections use</td><td>384</td></tr><tr><td>Users and permissions in the system tables</td><td>385</td></tr></table>	Topic	Page	Database permissions overview	352	Setting user and group options	356	Managing individual user IDs and permissions	357	Managing connected users	368	Managing groups	369	Database object names and prefixes	376	Using views and procedures for extra security	378	Changing Ownership on Nested Objects	381	How user permissions are assessed	383	Managing the resources connections use	384	Users and permissions in the system tables	385
Topic	Page																								
Database permissions overview	352																								
Setting user and group options	356																								
Managing individual user IDs and permissions	357																								
Managing connected users	368																								
Managing groups	369																								
Database object names and prefixes	376																								
Using views and procedures for extra security	378																								
Changing Ownership on Nested Objects	381																								
How user permissions are assessed	383																								
Managing the resources connections use	384																								
Users and permissions in the system tables	385																								

Database permissions overview

Proper management of user IDs and permissions lets users of a database carry out their jobs effectively, while maintaining the security and privacy of information within the database.

You use SQL statements for assigning user IDs to new users of a database, granting and revoking permissions for database users, and finding out the current permissions of users.

Database permissions are assigned to user IDs. Throughout this chapter, the term **user** is used as a synonym for user ID. Remember, however, that you grant and revoke permissions for each user ID.

Setting up individual user IDs

Even if there are no security concerns regarding a multi-user database, there are good reasons for setting up an individual user ID for each user. The administrative overhead is very low if a group with the appropriate permissions is set up. This chapter discusses groups of users.

You may want to use individual user IDs, since:

- ◆ The log translation utility can selectively extract the changes made by individual users from a transaction log. This is very useful when troubleshooting or piecing together what happened if data is incorrect.
- ◆ Sybase Central displays much more useful information with individual user IDs as you can tell which connections are which users.
- ◆ Row locking messages (with the BLOCKING option set to OFF) are more informative.

DBA authority overview

When you create a database, you also create a single usable user ID. This first user ID is **DBA**, and the password is initially **SQL**. The **DBA** user ID automatically has DBA authority within the database. This level of permission enables DBA users to carry out any activity in the database. They can create tables, change table structures, create new user IDs, revoke permissions from users, and so on.

Users with DBA authority

A user with DBA authority becomes the **database administrator**. In this chapter, references made to the database administrator, or **DBA**, include *any user or users with DBA authority*.

Although DBA authority may be granted or transferred to other user IDs, this chapter assumes that the **DBA** user ID is the database administrator, and that the abbreviation *DBA* means both the **DBA** user ID and any user ID with DBA authority.

Adding new users

The DBA has the authority to add new users to the database. As the DBA adds users, they are also granted permissions to carry out tasks on the database. Some users may need to simply look at the database information using SQL queries, others may need to add information to the database, and others may need to modify the structure of the database itself. Although some of the responsibilities of the DBA may be handed over to other user IDs, the DBA is responsible for the overall management of the database by virtue of the DBA authority.

The DBA has authority to create database objects and assign ownership of these objects to other user IDs.

RESOURCE authority overview

RESOURCE authority is the permission to create database objects, such as tables, views, stored procedures, and triggers. RESOURCE authority may be granted only by the DBA.

To create a trigger, a user needs both RESOURCE authority and ALTER permissions on the table to which the trigger applies.

Ownership permissions overview

The creator of a database object becomes the owner of that object. Ownership of a database object carries with it permissions to carry out actions on that object. These permissions are not assigned to users in the same way that other permissions in this chapter are assigned.

Owners

A user who creates a new object within the database is called the **owner** of that object, and automatically has permission to carry out any operation on that object. The owner of a table may modify the structure of that table, for instance, or may grant permissions to other database users to update the information within the table.

The DBA has permission to modify any component within the database, and so could delete a table created by another user. The DBA has all the permissions regarding database objects that the owners of each object have. As well, the DBA can also create database objects for other users. In this case the owner of an object is not the user ID that executed the CREATE statement. A use for this ability is discussed in "Groups without passwords" on page 373. Despite this possibility, this chapter refers interchangeably to the owner and creator of database objects as the same person.

Table and views permissions overview

There are several distinct permissions you may grant to user IDs concerning tables and views:

Permission	Description
ALTER	Permission to alter the structure of a table or create a trigger on a table
DELETE	Permission to delete rows from a table or view
INSERT	Permission to insert rows into a table or view
REFERENCES	Permission to create indexes on a table, and to create foreign keys that reference a table
SELECT	Permission to look at information in a table or view
UPDATE	Permission to update rows in a table or view. This may be granted on a set of columns in a table only
ALL	All the above permissions

Group permissions overview

Setting permissions individually for each user of a database can be a time-consuming and error-prone process. For most databases, permission management based on groups, rather than on individual user IDs, is a much more efficient approach.

You can assign permissions to a group in exactly the same way as to an individual user. You can then assign membership in appropriate groups to each new user of the database, and they gain a set of permissions by virtue of their group membership.

Example

For example, you may create groups for different departments in a company database (sales, marketing, and so on) and assign these groups permissions. Each salesperson becomes a member of the sales group, and automatically gains access to the appropriate areas of the database.

Each user ID can be a member of multiple groups, and they inherit all permissions from each of the groups.

Setting user and group options

In Sybase Central, configurable options for users and groups are located in the User Options and Group Options dialogs (the same dialog as for setting database options). In Interactive SQL, you can specify an option in a SET OPTION statement.

❖ **To set options for a user or group (Sybase Central):**

- 1 Connect to a database and open the Users & Groups folder.
- 2 Right-click the desired user or group and choose Options from the popup menu.
- 3 Edit the desired values.

❖ **To set the options for a user or group (SQL):**

- ◆ Specify the desired properties within a SET OPTION statement.

☞ See also

- ◆ "Setting properties for database objects" on page 34 of the book *ASA SQL User's Guide*
- ◆ "Database Options" on page 535

Managing individual user IDs and permissions

This section describes how to create new users and grant permissions to them. For most databases, the bulk of permission management should be carried out using **groups**, rather than by assigning permissions to individual users one at a time. However, as a group is simply a user ID with special properties, read and understand this section before moving on to the discussion of managing groups.

Creating new users

In both Sybase Central and Interactive SQL, you can create new users. In Sybase Central, you can manage users or groups in the Users & Groups folder. In Interactive SQL, you can add a new user using the GRANT CONNECT statement. For both tools, you need DBA authority to create new users.

All new users are automatically added to the PUBLIC group. Once you have created a new user, you can:

- ◆ add it to other groups
- ◆ set its permissions on tables, views, and procedures
- ◆ set it as the publisher or as a remote user of the database

Initial permissions for new users

By default, new users are not assigned any permissions beyond connecting to the database and viewing the system tables. To access tables in the database, they need to be assigned permissions.

The DBA can set the permissions granted automatically to new users by assigning permissions to the special PUBLIC user group, as discussed in "Special groups" on page 374.

❖ To create a new user (Sybase Central):

- 1 Open the Users & Groups folder.
- 2 In the right pane, double-click Add User.
- 3 Follow the instructions in the wizard.


❖ To create a new user (SQL):

- 1 Connect to a database with DBA authority.
- 2 Execute a GRANT CONNECT TO statement.

Example

Add a new user to the database with the user ID of M_Haneef and a password of welcome.

```
GRANT CONNECT TO M_Haneef  
IDENTIFIED BY welcome
```

 See also

- ◆ "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*

Changing a password

Changing a user's password

Using the GRANT statement, you can change your password or that of another user if you have DBA authority. For example, the following command changes the password for user ID M_Haneef to new_password:

```
GRANT CONNECT TO M_Haneef  
IDENTIFIED BY new_password
```

Changing the DBA password

The default password for the **DBA** user ID for all databases is **SQL**. You should change this password to prevent unauthorized access to your database. The following command changes the password for user ID **DBA** to **new_password**:

```
GRANT CONNECT TO DBA  
IDENTIFIED BY new_password
```

Granting DBA and RESOURCE authority

You can grant DBA and RESOURCE authority in the same manner.

❖ **To grant RESOURCE permissions to a user ID:**

- 1 Connect to the database as a user with DBA authority.
- 2 Type and execute the SQL statement:

```
GRANT RESOURCE TO userid
```

For DBA authority, the appropriate SQL statement is:

```
GRANT DBA TO userid
```

Notes

- ◆ Only the DBA may grant DBA or RESOURCE authority to database users.

- ◆ DBA authority is very powerful, since anyone with this authority has the ability to carry out any action on the database and as well as access to all the information in the database. It is wise to grant DBA authority to only a few people.
- ◆ Consider giving users with DBA authority two user IDs, one with DBA authority and one without, so that they connect as DBA only when necessary.
- ◆ RESOURCE authority allows the user to create new database objects, such as tables, views, indexes, procedures, or triggers.

Granting permissions on tables

You can assign a set of permissions on individual tables and grant users combinations of these permissions to define their access to a table.

You can use either Sybase Central or Interactive SQL to set permissions. In Interactive SQL, you can use the following SQL statements to grant permissions on tables.

- ◆ The ALTER permission allows a user to alter the structure of a table or to create triggers on a table. The REFERENCES permission allows a user to create indexes on a table, and to create foreign keys. These permissions grant the authority to modify the database schema, and so will not be assigned to most users. These permissions do not apply to views.
- ◆ The DELETE, INSERT, and UPDATE permissions grant the authority to modify the data in a table. Of these, the UPDATE permission may be restricted to a set of columns in the table or view.
- ◆ The SELECT permission grants authority to look at data in a table, but does not give permission to alter it.
- ◆ ALL permission grants all the above permissions.

❖ To grant permissions on tables or columns (Sybase Central):

- 1 Connect to the database.
- 2 Open the Tables folder for that database.
- 3 Right-click a table and choose Properties from the popup menu.
- 4 On the Permissions tab of the Table property sheet, configure the permissions for the table:
 - ◆ Click Grant to select users or groups to which to grant full permissions.

- ◆ Click in the fields beside the user or group to set specific permissions. Permissions are indicated by a check mark, and grant options are indicated by a check mark with two '+' signs.
- ◆ Select a user and click the button beside References, Select, or Update to set that type of permission on individual columns.
- ◆ Select a user or group in the list and click Revoke to revoke all permissions.


Tips

Legend for the columns on the Permissions page: A=Alter, D=Delete, I=Insert, R=Reference, S=Select, U=Update.

You can also assign permissions from the user/group property sheet. To assign permissions to many users and groups at once, use the table's property sheet. To assign permissions to many tables at once, use the user's property sheet.

❖ **To grant permissions on tables or columns (SQL):**

- 1 Connect to the database with DBA authority or as the owner of the table.
- 2 Execute a GRANT statement to assign the permission.

 For more information, see "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*.

Example 1

All table permissions are granted in a very similar fashion. You can grant permission to M_Haneef to delete rows from the table named *sample_table* as follows:

- 1 Connect to the database as a user with DBA authority, or as the owner of *sample_table*.
- 2 Type and execute the following SQL statement:

```
GRANT DELETE
ON sample_table
TO M_Haneef
```


Example 2

You can grant permission to M_Haneef to update the *column_1* and *column_2* columns only in the table named *sample_table* as follows:

- 1 Connect to the database as a user with DBA authority, or as the owner of *sample_table*.
- 2 Type and execute the following SQL statement:

```
GRANT UPDATE (column_1, column_2)
ON sample_table
TO M_Haneef
```

Table view permissions are limited in that they apply to all the data in a table (except for the UPDATE permission which may be restricted). You can fine-tune user permissions by creating procedures that carry out actions on tables, and then granting users the permission to execute the procedure.

 See also

- ◆ "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*

Granting permissions on views

Setting permissions on views is similar to setting them on tables; for more information about the SQL statements involved, see "Granting permissions on tables" on page 359.

A user may perform an operation through a view if one or more of the following are true:

- ◆ The appropriate permission(s) on the view for the operation has been granted to the user by a DBA.
- ◆ The user has the appropriate permission(s) on all the base table(s) for the operation.
- ◆ The user was granted appropriate permission(s) for the operation on the view by a non-DBA user. This user must be either the owner of the view or have WITH GRANT OPTION of the appropriate permission(s) on the view. The owner of the view must be either:
 - ◆ a DBA.
 - ◆ a non-DBA, but also the owner of all the base table(s) referred to by the view.
 - ◆ a non-DBA, and not the owner of some or all of the base table(s) referred to by the view. However, the view owner has SELECT permission WITH GRANT OPTION on the base table(s) not owned and any other required permission(s) WITH GRANT OPTION on the base table(s) not owned for the operation.

Instead of the owner having permission(s) WITH GRANT OPTION on the base table(s), permission(s) may have been granted to PUBLIC. This includes SELECT permission on system tables.

UPDATE permissions can be granted only on an entire view. Unlike tables, UPDATE permissions cannot be granted on individual columns within a view.

❖ **To grant permissions on views (Sybase Central):**

- 1 Connect to the database.
- 2 Open the Views folder for that database.
- 3 Right-click a view and choose Properties from the popup menu.
- 4 On the Permissions tab of the View property sheet, configure the permissions for the view:
 - ◆ Click Grant to select users or groups to which to grant full permissions.
 - ◆ Click in the fields beside the user or group to set specific permissions. Permissions are indicated by a check mark, and grant options are indicated by a check mark with two '+' signs.
 - ◆ Select a user or group in the list and click Revoke to revoke all permissions.

Tip

You can also assign permissions from the User or Group property sheet. To assign permissions to many users and groups at once, use the view's property sheet. To assign permissions to many views at once, use the User's or Group's property sheet.

Behavior change

There was a behavior change with Version 5 of the software concerning the permission requirements. Previously, permissions on the underlying tables were required to grant permissions on views.

☞ See also

- ◆ "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*

Granting users the right to grant permissions

You can assign each of the table and view permissions described in "Granting permissions on tables" on page 359 with the **WITH GRANT OPTION**. This option gives the right to pass on the permission to other users. In the context of groups, you can read about this feature in section "Permissions of groups" on page 372.

In Sybase Central, you can specify a grant option by displaying the property sheet of a user, group, or table, clicking the Permissions tab, and double-clicking in the fields provided so that a check mark with two '+' signs appears.

Example

You can grant permission to M_Haneef to delete rows from the table named *sample_table*, and the right to pass on this permission to other users, as follows:

- 1 Connect to the database as a user with DBA authority, or as the owner of *sample_table*:
- 2 Type and execute the SQL statement:

```
GRANT DELETE ON sample_table  
TO M_Haneef  
WITH GRANT OPTION
```

 For more information, see "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*

Granting permissions on procedures

The DBA or the owner of the procedure (the user ID that created the procedure) may grant permission to execute stored procedures. The **EXECUTE** permission is the only permission that may be granted on a procedure. This permission executes (or calls) the procedure.

The method for granting permissions to execute a procedure is similar to that for granting permissions on tables and views, discussed in "Granting permissions on tables" on page 359. However, the **WITH GRANT** option clause of the **GRANT** statement does not apply to the granting of permissions on procedures.

You can use either Sybase Central or Interactive SQL to set permissions.

❖ To grant permissions on procedures (Sybase Central):

- 1 Connect to the database.
- 2 Open the Procedures & Functions folder for that database.

- 3 Right-click a procedure and choose Properties from the popup menu.
- 4 On the Permissions tab of the Procedure property sheet, configure the permissions for the procedure:
 - ◆ Click Grant to select users or groups to which to grant full permissions.
 - ◆ Click beside users in the Execute column to toggle between granting or not granting permission.
 - ◆ Select a user or group in the list and click Revoke to revoke all permissions.

Tip

You can also assign permissions from the User or Group property sheet. To assign permissions to many users and groups at once, use the procedure's property sheet. To assign permissions to many views at once, use the User's or Group's property sheet.

❖ **To grant permissions on procedures (SQL):**

- 1 Connect to the database with DBA authority or as the owner of the procedure.
- 2 Execute a GRANT EXECUTE ON statement.

Example

You can grant M_Haneef permission to execute a procedure named *my_procedure*, as follows:

- 1 Connect to the database as a user with DBA authority or as owner of *my_procedure* procedure.
- 2 Execute the SQL statement:

```
GRANT EXECUTE
ON my_procedure
TO M_Haneef
```

**Execution
permissions of
procedures**

Procedures execute with the permissions of their owner. Any procedure that updates information on a table will execute successfully only if the owner of the procedure has UPDATE permissions on the table.


As long as the procedure owner has the proper permissions, the procedure executes successfully when called by any user assigned permission to execute it, whether or not they have permissions on the underlying table. You can use procedures to allow users to carry out well-defined activities on a table, without having any general permissions on the table.

 See also

- ◆ "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*

Execution permissions of triggers

The server executes triggers in response to a user action. Triggers do not require permissions to be executed. When a trigger executes, it does so with the permissions of the creator of the table with which they are associated.

 For more information on trigger permissions, see "Trigger execution permissions" on page 527 of the book *ASA SQL User's Guide*.

Granting and revoking remote permissions

In Sybase Central, you can manage the remote permissions of both users and groups. Remote permissions allow normal users and groups to become remote users in a SQL Remote replication setup in order to exchange replication messages with the publishing database.

Granting remote permissions

You cannot grant remote permissions to a user until you define at least one message type in the database.

To grant remote permissions to a group, you must explicitly grant remote permissions to each user in the group. You cannot grant remote permissions to a group itself.

Revoking remote permissions

Revoking remote permissions reverts a remote user to a normal user. Revoking these permissions also automatically unsubscribes that user from all publications.

❖ To grant remote permissions to users (Sybase Central):

- 1 Connect to a database.
- 2 Open the Users & Groups folder.
- 3 Right-click the desired user and choose Change to Remote User from the popup menu.
- 4 In the resulting dialog, enter the desired values.

Once you have granted remote permissions to a user, you can subscribe it to publications:

❖ **To revoke remote permissions from remote users:**

- 1 Open either the Users & Groups folder or the Remote Users folder (located within the SQL Remote folder).
- 2 Right-click the desired remote user and choose Revoke Remote from the popup menu.

☞ For more information, see "SQL Remote Concepts" on page 7 of the book *SQL Remote User's Guide*.

Revoking user permissions

Any user's permissions are a combination of those that have been granted and those that have been revoked. By revoking and granting permissions, you can manage the pattern of user permissions on a database.

The REVOKE statement is the exact converse of the GRANT statement. To disallow M_Haneef from executing *my_procedure*, the command is:

```
REVOKE EXECUTE
ON my_procedure
FROM M_Haneef
```

The DBA or the owner of the procedure must issue this command.

Permission to delete rows from *sample_table* may be revoked by issuing the command:

```
REVOKE DELETE
ON sample_table
FROM M_Haneef
```

☞ For more information on using Sybase Central to grant or revoke permission, see the following sections:

- ◆ "Granting permissions on tables" on page 359
- ◆ "Granting permissions on views" on page 361
- ◆ "Granting permissions on procedures" on page 363

Deleting users from the database

You can delete a user from the database using both Sybase Central and Interactive SQL. The user being removed cannot be connected to the database during this procedure.

Deleting a user also deletes all database objects (such as tables) that they own.

Only the DBA can delete a user.

❖ **To delete a user from the database (Sybase Central):**

- 1 Open the Users & Groups folder.
- 2 Right-click the desired user and choose Delete from the popup menu.

Tips

You cannot delete users when you select them within a group container.


❖ **To delete a user from the database (SQL):**

- 1 Connect to a database.
- 2 Execute a REVOKE CONNECT FROM statement.

Example

Remove the user M_Haneef from the database.

```
REVOKE CONNECT FROM M_Haneef
```

 See also

- ◆ "REVOKE statement" on page 516 of the book *ASA SQL Reference Manual*
- ◆ "Revoking user permissions" on page 366
- ◆ "Deleting groups from the database" on page 375

Managing connected users

If you are working with Sybase Central, you can keep track of all users connected to the database. You can view properties of these connected users, and you can disconnect them if you want.

❖ **To display a list of all users connected to a database:**

- 1 In Sybase Central, connect to a database.
- 2 Open the Connected Users folder for that database.

This folder displays all other users currently connected to a given database (including yourself), regardless of the client that they used to connect (Sybase Central, Interactive SQL, a custom client application, etc.).

❖ **To inspect the properties of a user's connection to a database:**

- 1 In Sybase Central, connect to a database.
- 2 Open the Connected Users folder for that database.
- 3 Right-click the desired user and choose Properties from the popup menu.
- 4 Inspect the desired properties.

❖ **To disconnect users from a database:**

- 1 In Sybase Central, connect to a database.
- 2 Open the Connected Users folder for that database.
- 3 Right-click the desired user and choose Disconnect from the popup menu.

Managing groups

DBA, RESOURCE,
and GROUP
permissions


Once you understand how to manage permissions for individual users (as described in the previous section), working with groups is straightforward. A user ID identifies a group, just like it does a single user. However, a group user ID has the permission to have members.


When you grant permissions to a group or revoke permissions from a group for tables, views, and procedures, all members of the group inherit those changes. The DBA, RESOURCE, and GROUP permissions are not inherited: you must assign them individually to each individual user IDs requiring them.


A group is simply a user ID with special permissions. You grant permissions to a group and revoke permissions from a group in exactly the same manner as any other user, using the commands described in "Managing individual user IDs and permissions" on page 357.

You can construct a hierarchy of groups, where each group inherits permissions from its parent group. That means that a group can also be a member of a group. As well, each user ID may belong to more than one group, so the user-to-group relationship is many-to-many.

The ability to create a group without a password enables you to prevent anybody from signing on using the group user ID.

 For more information about this security feature, see "Groups without passwords" on page 373.

 For more information on altering database object properties, see "Setting properties for database objects" on page 34 of the book *ASA SQL User's Guide*.

 For more information about granting remote permissions for groups, see "Granting and revoking remote permissions" on page 365.

Creating groups

You can create a new group in both Sybase Central and Interactive SQL. You need DBA authority to create a new group.

The GROUP permission, which gives the user ID the ability to have members, is not inherited by members of a group. Otherwise, every user ID would automatically be a group as a consequence of its membership in the special PUBLIC group.

❖ **To create a new group (Sybase Central):**

- 1 Open the Users & Groups folder.
- 2 In the right pane, double-click Add Group.
- 3 Follow the instructions in the wizard.

❖ **To create a new group (SQL):**

- 1 Connect to a database.
- 2 Execute a GRANT GROUP TO statement. If the user ID you cite in this statement has not already been created, you need to create it first.

Example

Create the user ID personnel.

```
GRANT CONNECT  
TO personnel  
IDENTIFIED BY group_password
```

Make the user ID personnel a group.

```
GRANT GROUP TO personnel
```

 See also

- ◆ "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*
- ◆ "Creating new users" on page 357

Granting group membership to existing users or groups

You can add existing users to groups or add groups to other groups in both Sybase Central and Interactive SQL. In Sybase Central, you can control group membership on the property sheets of users or groups. In Interactive SQL, you can make a user a member of a group with the GRANT statement.

When you assign a user membership in a group, they inherit all the permissions on tables, views, and procedures associated with that group.

Only the DBA can grant membership in a group.

❖ **To add a user or group to another group (Sybase Central):**

- 1 Open the Users & Groups folder.
- 2 Right-click the user/group that you want to add to another group and choose Properties from the popup menu.

- 3 Click the Membership tab of the property sheet.
- 4 Click Join to open the Join Groups dialog.
- 5 Select the desired group and click OK. The user or group associated with the property sheet is then added to this desired group.

❖ **To add a user or group to another group (SQL):**

- 1 Connect to a database.
- 2 Execute a GRANT MEMBERSHIP IN GROUP statement, specifying the desired group and the users involved.

Example

Grant the user M_Haneef membership in the personnel group:

```
GRANT MEMBERSHIP
IN GROUP personnel
TO M_Haneef
```

 See also

- ◆ "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*
- ◆ "Creating new users" on page 357

Revoking group membership

You can remove users or groups from a group in both Sybase Central and Interactive SQL.

Removing a user or group from a group does *not* delete them from the database (or from other groups). To do this, you must delete the user/group itself.

Only the DBA can revoke membership in a group.

❖ **To remove a user or group from another group (Sybase Central):**

- 1 Open the Users & Groups folder.
- 2 Right-click the desired user/group and choose Properties from the popup menu.
- 3 Click the Membership tab of the property sheet.
- 4 Select the desired group to leave and click Leave. The user or group associated with the property sheet is then removed from this desired group.

Tip

You can perform this action by opening the Users & Groups folder, right-clicking the user or group that is currently a member of another group, and choosing Leave.

❖ To remove a user or group from another group (SQL):

- 1 Connect to a database.
- 2 Execute a REVOKE MEMBERSHIP IN GROUP statement, specifying the desired group and the users involved.

Example

Remove the user M_Haneef from the personnel group:

```
REVOKE MEMBERSHIP  
IN GROUP personnel  
FROM M_Haneef
```

 See also

- ◆ "REVOKE statement" on page 516 of the book *ASA SQL Reference Manual*
- ◆ "Creating new users" on page 357
- ◆ "Deleting users from the database" on page 366
- ◆ "Deleting groups from the database" on page 375

Permissions of groups

You may grant permissions to groups in exactly the same way as to any other user ID. Permissions on tables, views, and procedures are inherited by members of the group, including other groups and their members. Some complexities to group permissions exists, that database administrators need to keep in mind.

Notes

Members of a group do not inherit the DBA, RESOURCE, and GROUP permissions. Even if the personnel user ID has RESOURCE permissions, the members of personnel do not have RESOURCE permissions.

Ownership of database objects is associated with a single user ID and is not inherited by group members. If the user ID *personnel* creates a table, then the *personnel* user ID is the owner of that table and has the authority to make any changes to the table, as well as to grant privileges concerning the table to other users. Other user IDs who are members of *personnel* are not the owners of this table, and do not have these rights. Only granted permissions are inherited. For example, if the DBA or the *personnel* user ID explicitly grants *SELECT* authority to the *personnel* user ID, all group members do have select access to the table.

Referring to tables owned by groups

Groups are used for finding tables and procedures in the database. For example, the query

```
SELECT * FROM SYSGROUPS
```

always finds the table *SYSGROUPS*, because all users belong to the *PUBLIC* group, and *PUBLIC* belongs to the *SYS* group which owns the *SYSGROUPS* table. (The *SYSGROUPS* table contains a list of *group_name*, *member_name* pairs representing the group memberships in your database.)

If a table named *employees* is owned by the user ID *personnel*, and if **M_Haneef** is a member of the *personnel* group, then **M_Haneef** can refer to the *employees* table simply as *employees* in SQL statements. Users who are not members of the *personnel* group need to use the qualified name *personnel.employees*.

Creating a group to own the tables

A good practice to follow that allows everyone to access the tables without qualifying names, is to create a group whose only purpose is to own the tables. Do not grant any permissions to this group, but make all users members of the group. You can then create permission groups and grant users membership in these permission groups as warranted.

☞ For an example, see the section "Database object names and prefixes" on page 376.

Groups without passwords

Users connected to a group's user ID have certain permissions. This user ID can grant and revoke membership in the group. Also, this user would have ownership permissions over any tables in the database created in the name of the group's user ID.

It is possible to set up a database so that only the DBA handles groups and their database objects, rather than permitting other user IDs to make changes to group membership. You can do this by disallowing connection as the group's user ID when creating the group. To do this, type the GRANT CONNECT statement without a password. Thus:

```
GRANT CONNECT  
TO personnel
```

creates a user ID personnel. This user ID can be granted group permissions, and other user IDs can be granted membership in the group, inheriting any permissions that have been given to personnel. However, nobody can connect to the database using the personnel user ID, because it has no valid password.

The user ID personnel can be an owner of database objects, even though no user can connect to the database using this user ID. The CREATE TABLE statement, CREATE PROCEDURE statement, and CREATE VIEW statement all allow the owner of the object to be specified as a user other than that executing the statement. Only the DBA can carry out this assignment of ownership.

Special groups

When you create a database, the SYS and PUBLIC groups are also automatically created. Neither of these groups has passwords, so it is not possible to connect to the database as either SYS or as PUBLIC. However, the two groups serve important functions in the database.

The SYS group

The SYS group owns the system tables and views for the database, which contain the full description of database structure, including all database objects and all user IDs.

☞ For more information about the system tables and views, together with a description of access to the tables, see the chapters "System Tables" on page 595 of the book *ASA SQL Reference Manual*, and also "System Views" on page 673 of the book *ASA SQL Reference Manual*.

The PUBLIC group

The PUBLIC group has CONNECT permissions to the database and SELECT permission on the system tables. As well, the PUBLIC group is a member of the SYS group, and has read access for some of the system tables and views, so any user of the database can find out information about the database schema. If you wish to restrict this access, you can REVOKE PUBLIC's membership in the SYS group.

Any new user ID is automatically a member of the PUBLIC group and inherits any permissions specifically granted to that group by the DBA. You can also REVOKE membership in PUBLIC for users if you wish.

Deleting groups from the database

You can delete a group from the database using both Sybase Central and Interactive SQL.

Deleting users or groups from the database is different from *removing* them from other groups. Deleting a group from the database does *not* delete its members from the database, although they lose membership in the deleted group.

Only the DBA can delete a group.

❖ To delete a group from the database (Sybase Central):

- 1 Open the Users & Groups folder.
- 2 Right-click the desired group and choose Delete from the popup menu.

❖ To delete a group from the database (SQL):

- 1 Connect to a database.
- 2 Execute a REVOKE CONNECT FROM statement.

Example

Remove the group personnel from the database.

```
REVOKE CONNECT FROM personnel
```

 See also

- ◆ "REVOKE statement" on page 516 of the book *ASA SQL Reference Manual*
- ◆ "Revoking user permissions" on page 366
- ◆ "Deleting users from the database" on page 366

Database object names and prefixes

The name of every database object is an identifier. The rules for valid identifiers appear in "Identifiers" on page 7 of the book *ASA SQL Reference Manual*.

In queries and sample SQL statements throughout this guide, database objects from the sample database are generally referred to using their simple name. For example:

```
SELECT *
FROM employee
```

Tables, procedures, and views all have an owner. The owner of the tables in the sample database is the user ID **DBA**. In some circumstances, you must prefix the object name with the owner user ID, as in the following statement.

```
SELECT *
FROM "DBA".employee
```

The *employee* table reference is said to be **qualified**. In other circumstances it is sufficient to give the object name. This section describes when you need to use the owner prefix to identify tables, views and procedures, and when you do not.

When referring to a database object, you require a prefix unless:

- ◆ You are the owner of the database object.
- ◆ The database object is owned by a group ID of which you are a member.

Example

Consider the following example of a corporate database. The user ID company created all the tables, and since the this user ID belongs to the database administrator, it therefore has DBA authority.

```
GRANT CONNECT TO company
IDENTIFIED BY secret;
GRANT DBA TO company;
```

The company user ID created the tables in the database.

```
CONNECT USER company IDENTIFIED BY secret;
CREATE TABLE company.Customers ( ... );
CREATE TABLE company.Products ( ... );
CREATE TABLE company.Orders ( ... );
CREATE TABLE company.Invoices ( ... );
CREATE TABLE company.Employees ( ... );
CREATE TABLE company.Salaries ( ... );
```

Not everybody in the company should have access to all information. Consider two user IDs in the sales department, Joe and Sally, who should have access to the *Customers*, *Products* and *Orders* tables. To do this, you create a Sales group.

```
GRANT CONNECT TO Sally IDENTIFIED BY xxxxxx;  
GRANT CONNECT TO Joe IDENTIFIED BY xxxxxx;  
GRANT CONNECT TO Sales IDENTIFIED BY xxxxxx;  
GRANT GROUP TO Sales;  
GRANT ALL ON Customers TO Sales;  
GRANT ALL ON Orders TO Sales;  
GRANT SELECT ON Products TO Sales;  
GRANT MEMBERSHIP IN GROUP Sales TO Sally;  
GRANT MEMBERSHIP IN GROUP Sales TO Joe;
```

Now Joe and Sally have permission to use these tables, but they still have to qualify their table references because the table owner is company, and Sally and Joe are not members of the company group:

```
SELECT *  
FROM company.customers
```

To rectify the situation, make the Sales group a member of the company group.

```
GRANT GROUP TO company;  
GRANT MEMBERSHIP IN GROUP company TO Sales;
```

Now Joe and Sally, being members of the Sales group, are indirectly members of the company group, and can reference their tables without qualifiers. The following command now works:

```
SELECT *  
FROM Customers
```

Note

Joe and Sally do not have any extra permissions because of their membership in the company group. The company group has not been explicitly granted any table permissions. (The company user ID has implicit permission to look at tables like *Salaries* because it created the tables and has DBA authority.) Thus, Joe and Sally still get an error executing either of these commands:

```
SELECT *  
FROM Salaries;  
  
SELECT *  
FROM company.Salaries
```

In either case, Joe and Sally do not have permission to look at the *Salaries* table.

Using views and procedures for extra security

For databases that require a high level of security, defining permissions directly on tables has limitations. Any permission granted to a user on a table applies to the whole table. There are many cases when users' permissions need to be shaped more precisely than on a table-by-table basis. For example:

- ◆ It is not desirable to give access to personal or sensitive information stored in an employee table to users who need access to other parts of the table.
- ◆ You may wish to give sales representatives update permissions on a table containing descriptions of their sales calls, but limit such permissions to their own calls.

In these cases, you can use views and stored procedures to tailor permissions to suit the needs of your organization. This section describes some of the uses of views and procedures for permission management.

🌀 For more information about how to create views, see "Working with views" on page 51 of the book *ASA SQL User's Guide*.

🌀 For more information about view permissions, see "Granting permissions on views" on page 361.

Using views for tailored security

Views are computed tables that contain a selection of rows and columns from base tables. Views are useful for security when it is appropriate to give a user access to just one portion of a table. The portion can be defined in terms of rows or in terms of columns. For example, you may wish to disallow a group of users from seeing the salary column of an employee table, or you may wish to limit a user to see only the rows of a table they have created.

Example

The Sales manager needs access to information in the database concerning employees in the department. However, there is no reason for the manager to have access to information about employees in other departments.

This example describes how to create a user ID for the sales manager, create views that provides the information she needs, and grants the appropriate permissions to the sales manager user ID.

- 1 Create the new user ID using the GRANT statement. While logged in as a user ID with DBA authority, enter the following:

```
CONNECT "DBA"  
IDENTIFIED by SQL ;
```

```
GRANT CONNECT  
TO SalesManager  
IDENTIFIED BY sales
```

- 2 Define a view which only looks at sales employees as follows:

```
CREATE VIEW emp_sales AS  
  SELECT emp_id, emp_fname, emp_lname  
  FROM "DBA".employee  
  WHERE dept_id = 200
```

The table should therefore be identified as *DBA.employee*, with the owner of the table explicitly identified, for the SalesManager user ID to be able to use the view. Otherwise, when SalesManager uses the view, the SELECT statement refers to a table that user ID does not recognize.

- 3 Give SalesManager permission to look at the view:

```
GRANT SELECT  
ON emp_sales  
TO SalesManager
```

You use exactly the same command to grant permission on views and on tables.

Example 2

The next example creates a view which allows the Sales Manager to look at a summary of sales orders. This view requires information from more than one table for its definition:

- 1 Create the view.

```
CREATE VIEW order_summary AS  
  SELECT order_date, region, sales_rep, company_name  
  FROM "DBA".sales_order  
  KEY JOIN "DBA".customer
```

- 2 Grant permission for the Sales Manager to examine this view.

```
GRANT SELECT  
ON order_summary  
TO SalesManager
```

- 3 To check that the process has worked properly, connect to the SalesManager user ID and look at the views you created:

```
CONNECT SalesManager  
IDENTIFIED BY sales ;  
  
SELECT *  
FROM "DBA".emp_sales ;  
  
SELECT *  
FROM "DBA".order_summary ;
```

No permissions have been granted to the Sales Manager to look at the underlying tables. The following commands produce permission errors.

```
SELECT * FROM "DBA".employee ;
SELECT * FROM "DBA".sales_order
```

Other permissions
on views

The previous example shows how to use views to tailor SELECT permissions. You can grant INSERT, DELETE, and UPDATE permissions on views in the same way.

☞ For more information about allowing data modification on views, see "Using views" on page 53 of the book *ASA SQL User's Guide*.

Using procedures for tailored security

While views restrict access on the basis of data, procedures restrict the actions a user may take. As described in "Granting permissions on procedures" on page 363, a user may have EXECUTE permission on a procedure without having any permissions on the table or tables on which the procedure acts.

Strict security

For strict security, you can disallow all access to the underlying tables, and grant permissions to users or groups of users to execute certain stored procedures. This approach strictly defines the manner in which data in the database can be modified.

Changing Ownership on Nested Objects

Views and procedures can access underlying objects that are owned by different users. For example, if `usera`, `userb`, `userc`, and `userd` were four different users, `userd.viewd` could be based on `userc.viewc`, which could be based on `userb.viewb`, which could be based on `usera.tablea`. Similarly for procedures, `userd.procd` could call `userc.procc`, which could call `userb.procb`, which could insert into `usera.tablea`.

The following Discretionary Access Control (DAC) rules apply to nested views and tables:

- ◆ To create a view, the user must have `SELECT` permission on all of the base objects (for example tables and views) in the view.
- ◆ To access a view, the view owner must have been granted the appropriate permission on the underlying tables or views with the `GRANT OPTION` and the user must have been granted the appropriate permission on the view.
- ◆ Updating with a `WHERE` clause requires both `SELECT` and `UPDATE` permission.
- ◆ If a user owns the tables in a view definition, the user can access the tables through a view, even if the user is not the owner of the view and has not been granted access on the view.

The following DAC rules apply to nested procedures:

- ◆ A user does not require any permissions on the underlying objects (for example tables, views or procedures) to create a procedure.
- ◆ For a procedure to execute, the owner of the procedure needs the appropriate permissions on the objects that the procedure references.
- ◆ Even if a user owns all the tables referenced by a procedure, the user will not be able to execute the procedure to access the tables unless the user has been granted `EXECUTE` permission on the procedure.

Following are some examples that describe this behavior.

Example 1: User1 creates table1, and user2 creates view2 on table1

- ◆ User1 can always access `table1`, since `user1` is the owner.
- ◆ User1 can always access `table1` through `view2`, since `user1` is the owner of the underlying table. This is true even if `user2` does not grant permission on `view2` to `user1`.

- ◆ User2 can access table1 directly or through view2 if user1 grants permission on table1 to user2.
- ◆ User3 can access table1 if user1 grants permission on table1 to user3
- ◆ User3 can access table1 through view2 if user1 grants permission on table1 to user2 with grant option AND user2 grants permission on view2 to user3.

Example 2: User2 creates procedure2 that accesses table1

- ◆ User1 can access table1 through procedure2 if user2 grants EXECUTE permission on procedure2 to user1. Note that this is different from the case of view2, where user1 did not need permission on view2.

Example 3: User1 creates table1, user2 creates table2, and user3 creates view3 joining table1 and table2

- ◆ User3 can access table1 and table2 through view3 if user1 grants permission on table1 to user3 AND user2 grants permission on table2 to user3.
- ◆ If user3 has permission on table1 but not on table2, then user3 cannot use view3, even to access the subset of columns belonging to table1.
- ◆ User1 or user2 can use view3 if (a) user1 grants permission with grant option on table1 to user3, (b) user2 grants permission with grant option on table2 to user3, AND (c) user3 grants permission on view3 to that user.

How user permissions are assessed

Groups do introduce complexities in the permissions of individual users. Suppose user M_Haneef has SELECT and UPDATE permissions on a specific table individually, but is also a member of two groups. Suppose one of these groups has no access to the table at all, and one has only SELECT access. What are the permissions in effect for this user?

Adaptive Server Anywhere decides whether a user ID has permission to carry out a specific action in the following manner:

- 1 If the user ID has DBA authority, the user ID can carry out any action in the database.
- 2 Otherwise, permission depends on the permissions assigned to the individual user. If the user ID has been granted permission to carry out the action, then the action proceeds.
- 3 If no individual settings have been made for that user, permission depends on the permissions of each of the groups to which the member belongs. If any of these groups has permission to carry out the action, the user ID has permission by virtue of membership in that group, and the action proceeds.

This approach minimizes problems associated with the order in which permissions are set.

Managing the resources connections use

Building a set of users and groups allows you to manage permissions on a database. Another aspect of database security and management is to limit the resources an individual user can use.


For example, you may wish to prevent a single connection from taking too much of the available memory or CPU resources, so you can avoid a having a connection slow down other users of the database.

Adaptive Server Anywhere provides a set of database options that the DBA can use to control resources. These options are called **resource governors**.

Setting options

You can set database options using the SET OPTION statement, with the following syntax:

```
SET [ TEMPORARY ] OPTION  
... [ userid. | PUBLIC. ] option-name = [ option-value ]
```

 For more information about options, see "Database Options" on page 535. For information on the SET OPTION statement, see "SET OPTION statement" on page 539 of the book *ASA SQL Reference Manual*.

Resources that can be managed

You can use the following options to manage resources:

- ◆ **JAVA_HEAP_SIZE** Sets the maximum size (in bytes) of the part of the memory allocated to Java applications on a per connection basis.
- ◆ **MAX_CURSOR_COUNT** Limits the number of cursors for a connection.
- ◆ **MAX_STATEMENT_COUNT** Limits the number of prepared statements for a connection.
- ◆ **BACKGROUND_PRIORITY** Limits the impact requests on the current connection have on the performance of other connections.

Database option settings are not inherited through the group structure.

Users and permissions in the system tables

The database system tables and system views stores information about the current users of a database and about their permissions.

☞ For more information about each of these tables, see "System Tables" on page 595 of the book *ASA SQL Reference Manual*.

The special user ID SYS owns the system tables. You cannot connect to the SYS user ID.

The DBA has SELECT access to all system tables, just as to any other tables in the database. The access of other users to some of the tables is limited. For example, only the DBA has access to the SYS.SYSUSERPERM table, which contains all information about the permissions of users of the database, as well as the encrypted passwords of each user ID. However, SYS.SYSUSERPERMS is a view containing all information in SYS.SYSUSERPERM except for the password, and by default all users have SELECT access to this view. You can fully modify all permissions and group memberships set up in a new database for SYS, PUBLIC, and DBA.

The following table summarizes the system tables containing information about user IDs, groups, and permissions. The user ID SYS owns all tables and views, and so their qualified names are SYS.SYSUSERPERM and so on.

Appropriate SELECT queries on these tables generates all the user ID and permission information stored in the database.

Table	Default	Contents
<i>SYSUSERPERM</i>	DBA only	Database-level permissions and password for each user ID
<i>SYSGROUP</i>	PUBLIC	One row for each member of each group
<i>SYSTABLEPERM</i>	PUBLIC	All permissions on table given by the GRANT commands
<i>SYSCOLPERM</i>	PUBLIC	All columns with UPDATE permission given by the GRANT command
<i>SYSDDUMMY</i>	PUBLIC	Dummy table, can be used to find the current user ID
<i>SYSROCPERM</i>	PUBLIC	Each row holds one user granted permission to use one procedure

The following table summarizes the system views containing information about user IDs, groups, and permissions.


Views	Default	Contents
<i>SYSUSERAUTH</i>	DBA only	All information in SYSUSERPERM except for user numbers
<i>SYSUSERPERMS</i>	PUBLIC	All information in SYSUSERPERM except for passwords
<i>SYSUSERLIST</i>	PUBLIC	All information in SYSUSERAUTH except for passwords
<i>SYSGROUPS</i>	PUBLIC	Information from SYSGROUP in a more readable format
<i>SYSTABAUTH</i>	PUBLIC	Information from SYSTABLEPERM in a more readable format
<i>SYSCOLAUTH</i>	PUBLIC	Information from SYSCOLPERM in a more readable format
<i>SYSPROCAUTH</i>	PUBLIC	Information from SYSPROCPerm in a more readable format

In addition to these, there are tables and views that contain information about each object in the database.

Keeping Your Data Secure

About this chapter This chapter describes Adaptive Server Anywhere features that help make your database secure. In particular, this chapter describes auditing, database encryption, and C2 certification. It presents overviews of other security features, providing pointers to where you can find more detailed information.

Database administrators are responsible for data security. In this chapter, unless otherwise noted, you require DBA authority to carry out the tasks described.

 User IDs and permissions are major security-related topics. For information on these topics, see "Managing User IDs and Permissions" on page 351.







Contents

Topic	Page
Security features overview	388
Controlling database access	390
Auditing database activity	393
Running the database server in a secure fashion	397
Encrypting a database	399
Security tips	403

Security features overview


Since databases may contain proprietary, confidential, or private information, ensuring that the database and the data in it are designed for security is very important.

Adaptive Server Anywhere has several features to assist in building a secure environment for your data:


- ◆ **User identification and authentication** These features control who has access to a database.
 For information on these subjects, see "Creating new users" on page 357.
- ◆ **Discretionary access control features** These features control the actions a user can carry out while connected to a database.
 For more information, see "Database permissions overview" on page 352.
- ◆ **Auditing** This feature helps you maintain a record of actions on the database.
 For more information, see "Auditing database activity" on page 393.
- ◆ **Database server options** These features let you control who can carry out operations (for example, loading databases). These options are set when you start the database server.
 For more information, see "Controlling permissions from the command line" on page 12.
- ◆ **Views and stored procedures** These features allow you to specify the data a user can access and the operations a user can execute.
 For more information, see "Using views and procedures for extra security" on page 378.
- ◆ **Database encryption** Database encryption features allow you to choose the level of database encryption. You can choose to secure your database either with simple encryption, or with strong encryption. Simple encryption is equivalent to obfuscation. Strong encryption renders the database completely inaccessible without the key.
 For more information, see "Initialization utility options" on page 467.

- ◆ **Communication encryption** You can encrypt client/server communications with simple or strong encryption for greater security as they pass over the network. Strong encryption is only supported over the TCP/IP port on Solaris, Linux, NetWare, and all supported Windows operating systems except Windows CE.

Communication encryption is a separately licensable component and must be ordered before you can install it. To order this component, see the card in your SQL Anywhere Studio package or see <http://www.sybase.com/detail?id=1015780>.

 For more information, see "Encrypting client/server communications" on page 398.

- ◆ **C2 certification** C2 is a set of security guidelines established by the U.S. government to maintain consistency within their organization. If you are running Adaptive Server Anywhere 7.0, and if you have the appropriate hardware, you can set up your machine to run in a C2 certified manner. The C2-certified documentation is available at <http://my.sybase.com/detail?id=1010458>.

 For information on running the current version of Adaptive Server Anywhere in a manner equivalent to the C2-certified environment, see "Installation" on page 3 of the book *ASA C2 Security Supplement*.

Controlling database access

By assigning user IDs and passwords, the database administrator controls who has access to a database. By granting permissions to each user ID, the database administrator controls what tasks each user can carry out when connected. This section describes the features available for controlling database access.

Permission scheme is based on user IDs

When you log onto the database, you have access to all database objects that meet *any* of the following criteria:

- ◆ objects you created.
- ◆ objects to which you received explicit permission.
- ◆ objects to which a group you belong to received explicit permission.

The user cannot access any database object that does not meet these criteria. In short, users can access only the objects they own or objects to which they explicitly received access permissions.

☞ For more information, see the following:

- ◆ "Managing User IDs and Permissions" on page 351
- ◆ "CONNECT statement [ESQL] [Interactive SQL]" on page 268 of the book *ASA SQL Reference Manual*
- ◆ "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*
- ◆ "REVOKE statement" on page 516 of the book *ASA SQL Reference Manual*

Using integrated logins

Integrated logins allow users to use a single login name and password to log onto both the Windows NT/2000/XP operating systems and onto a database. An external login name is associated with a database user ID. When you attempt an integrated login, you log onto the operating system by giving both a login name and password. The operating system then tells the server who you are, and the server logs you in as the associated database user ID. No additional login name or password are required.




There are some security implications of integrated logins to consider. For example, leaving the user profile *Guest* enabled with a blank password can permit unrestricted access to a database that is hosted by that server. Literally any user can log in to the server using any login ID and any password because they are logged in by default to the Guest user profile.

☞ For more information, see the following:

- ◆ "Security concerns: unrestricted database access" on page 87

- ◆ "Using integrated logins" on page 83
- ◆ "LOGIN_MODE option" on page 577

Increasing password security

Implement minimum password lengths	<p>Passwords are an important part of any database security system. To be secure, passwords must be difficult to guess, and they must not be easily accessible on users' hard drives or other locations.</p> <p>By default, passwords can be any length. For greater security, you can enforce a minimum length requirement on all new passwords. You do this by setting the MIN_PASSWORD_LENGTH database option to a value greater than zero. The following statement enforces passwords to be at least 8 bytes long.</p> <pre>SET OPTION PUBLIC.MIN_PASSWORD_LENGTH = 8</pre> <p> For more information, see "MIN_PASSWORD_LENGTH option" on page 583.</p>
Do not include password in ODBC data sources	<p>Passwords are the key to accessing databases. They should not be easily available to unauthorized people in a security-conscious environment.</p> <p>When you create an ODBC data source, or a Sybase Central connection profile, you can optionally include a password. Avoid including password for greater security.</p> <p> For information on creating ODBC data sources, see "Creating an ODBC data source" on page 53.</p>
Encrypt command files containing password	<p>When you create a configuration file, you can optionally include password information. To protect your password, consider hiding the contents of configuration files with simple encryption, using the File Hiding [dbfhide] utility.</p> <p> For information on the File Hiding [dbfhide] utility, see "The File Hiding utility" on page 446.</p>

Controlling the tasks users can perform

Users can access only those objects to which they have been granted access.

You grant permission on an object to another user with the GRANT statement. You can also grant a user permission to pass on the permissions on an object to other users.

The GRANT statement also gives more general permissions to users. Granting CONNECT permissions to a user allows them to connect to the database and change their passwords. Granting RESOURCE authority allows the user to create tables, views, procedures, and so on. Granting DBA authority to a user gives that user the ability to see and do anything in the database. The DBA would also use the GRANT statement to create and administer groups.

The REVOKE statement is the opposite of the GRANT statement—any permission that GRANT has explicitly given, REVOKE can take away. Revoking CONNECT from a user removes the user from the database, including all objects owned by that user.

Negative permissions

Adaptive Server Anywhere does not support **negative permissions**. This means that you cannot revoke a permission that was not explicitly granted.

For example, suppose user *bob* is a member of a group called *sales*. If a user grants DELETE permission on a table, T, to sales, then bob can delete rows from T. If you want to prevent bob from deleting from T, you cannot simply execute a REVOKE DELETE on T from bob, since the DELETE ON T permission was never granted directly to bob. In this case, you would have to revoke bob's membership in the sales group.

☞ For more information, see:

- ◆ "GRANT statement" on page 443 of the book *ASA SQL Reference Manual*
- ◆ "REVOKE statement" on page 516 of the book *ASA SQL Reference Manual*

Designing database objects for security

Views and stored procedures provide alternative ways of tuning the data users can access and the tasks they can perform.

☞ For more information on these features, see:

- ◆ "Benefits of procedures and triggers" on page 510 of the book *ASA SQL User's Guide*
- ◆ "Using views and procedures for extra security" on page 378


Auditing database activity

Auditing is a way of keeping track of the activity performed on a database. The record of activities stays in the transaction log. By turning on auditing, the DBA increases the amount of data saved in the transaction log to include the following:

- ◆ All login attempts (successful and failed), including the terminal ID.
- ◆ Accurate timestamps of all events (to a resolution of milliseconds)
- ◆ All permissions checks (successful and failed), including the object on which the permission was checked (if applicable)
- ◆ All actions that require DBA authority.

The transaction log

Each database has an associated transaction log file. The transaction log is used for database recovery. It is a record of transactions executed against a database.

 For information about the transaction log, see "The transaction log" on page 305.

The transaction log stores all executed data definition statements, and the user ID that executed them. It also stores all updates, deletes, and inserts and which user executed those statements. However, this is insufficient for some auditing purposes. By default, the transaction log does not contain the time of the event, just the order in which events occurred. It also contains neither failed events, nor select statements.

Turning on auditing

The database administrator can turn on **auditing** to add security-related information to the transaction log.

Auditing is off by default. To enable auditing on a database, the DBA must set the value of the public option `AUDITING` to `ON`. Auditing then remains enabled until explicitly disabled, by setting the value of the `AUDITING` option to `OFF`. You must have DBA permissions to set this option.

❖ To turn on auditing:

- 1 Ensure that your database is upgraded to at least version 6.0.2.
- 2 If you had to upgrade your database, create a new transaction log.
- 3 Execute the following statement:

```
SET OPTION PUBLIC.AUDITING = 'ON'
```

🔗 For more information, see "AUDITING option" on page 553.

Retrieving audit information

You can use the Log Translation [dbtran] utility to retrieve audit information. You can access this utility from Sybase Central or from the command prompt. It operates on a transaction log to produce a SQL script containing all of the transactions, along with some information on what user executed each command. By using the `-g` option, *dbtran* includes more comments containing the auditing information.

To ensure a complete and readable audit record, the `-g` option automatically sets the following options:

- ◆ **-d** Display output in chronological order.
- ◆ **-t** Include trigger-generated operations in the output.
- ◆ **-a** Include rolled back transactions in the output.

You can run the Log Translation Utility against a running database server or against a database log file.

❖ To retrieve auditing information from a running database server:

- 1 Make sure your user ID has DBA authority.
- 2 With the database server running, execute the following statement at a system command prompt:

```
dbtran -g -c "uid=DBA;pwd=SQL;..." -n asademo.SQL
```

🔗 For information about connection strings, see "Connection parameters" on page 70.

❖ To retrieve auditing information from a transaction log file:

- 1 Close the database server to ensure the log file is available.
- 2 At a system command prompt, execute the following statement to place the information from the file *asademo.log* and into the file *asademo.SQL*.

```
dbtran -g asademo.log
```

The `-g` option includes auditing information in the output file.

🔗 For more information, see "The Log Translation utility" on page 488.

Adding audit comments

You can add comments to the audit trail using the *sa_audit_string* system stored procedure. It takes a single argument, which is a string of up to 200 bytes. You must have DBA permissions to call this procedure.

For example:

```
call sa_audit_string( 'Started audit testing here.' )
```

This comment is stored in the transaction log as an audit statement.

An auditing example

This example shows how the auditing feature records attempts to access unauthorized information.

- 1 As database administrator, turn on auditing.

You can do this from Sybase Central as follows:

- ◆ Connect to the ASA 7.0 Sample data source. This connects you as the **DBA** user.
- ◆ Right-click the **asademo** database icon and choose Options from the popup menu.
- ◆ Select Auditing from the list of options, and enter the value ON in the Public Setting box. Click Set Permanent Now to set the option and Close to exit.

Alternatively, you can use Interactive SQL. Connect to the sample database from Interactive SQL as user ID **DBA** with password **SQL** and execute the following statement:

```
SET OPTION PUBLIC.AUDITING = 'ON'
```

- 2 Add a user to the sample database, named *BadUser*, with password **BadUser**. You can do this from Sybase Central. Alternatively, you can use Interactive SQL and enter the following statement:

```
GRANT CONNECT TO BadUser  
IDENTIFIED BY 'BadUser'
```

- 3 Using Interactive SQL, connect to the sample database as *BadUser* and attempt to access confidential information in the **employee** table with the following query:

```
SELECT emp_lname, salary  
FROM DBA.employee
```

You receive an error message: do not have permission to select from employee.

- 4 From a command prompt, change directory to your Adaptive Server Anywhere installation directory, which holds the sample database, and execute the following command:

```
dbtran -g -c "dsn=ASA 7.0 Sample" -n asademo.SQL
```

This command produces a file named *asademo.SQL*, containing the transaction log information and a set of comments holding audit information. The lines indicating the unauthorized *BadUser* attempt to access the employee table are included in the file as follows:

```
--AUDIT-1001-0000287812 -- 1999/02/11 13:59:58.765
Checking Select permission on employee - Failed
--AUDIT-1001-0000287847 -- 1999/02/11 13:59:58.765
Checking Select permission on employee(salary) -
Failed
```

- 5 Restore the sample database to its original state so other examples you try in this documentation give the expected results.

Connect as the DBA user, and carry out the following operations:

- ◆ Revoke Connect privileges from the user ID **BadUser**.
- ◆ Set the PUBLIC.AUDITING option to 'OFF'.

Auditing actions outside the database server

Some database utilities act on the database file directly. In a secure environment, only trusted users should have access to the database files.

To provide auditing of actions, under Windows NT only, any use of *dbtran*, *dbwrite*, and *dblog* generates a text file in the same directory as the database file, with the extension *.alg*. For example, for *asademo.db*, the file is called *asademo.alg*. Records containing the tool name, Windows user name, and date/time are appended to this file. Records are only added to the *.alg* file if the AUDITING option is set to ON.

Running the database server in a secure fashion

There are several security features you can set either when starting the database server or during server operation, including:

- ◆ **Starting and stopping databases** By default, any user can start an extra database on a running server. The `-gd` option allows you to limit access to this option to users with a certain level of permission in the database to which they are already connected. The permissible values include **DBA**, **all**, or **none**.

☞ For more information, see "`-gd` server option" on page 139.

- ◆ **Creating and deleting databases** By default, any user can use the `CREATE DATABASE` statement to create a database file. The `-gu` option allows you to limit access to this option to users with a certain level of permission in the database to which they are connected. The permissible values include **DBA**, **all**, **none**, or **utility_db**.

☞ For information, see "`-gu` server option" on page 143.

- ◆ **Stopping the server** The `dbstop` utility stops a database server. It is useful in batch files, or in other cases where interactive stopping of the server (by clicking Shutdown on the server window) is impractical. By default, any user can run `dbstop` to shut down a server. The `-gk` option allows you to limit access to this option to users with a certain level of permission in the database. The permissible values include **DBA**, **all**, or **none**.

☞ For more information, see "`-gk` server option" on page 140.

- ◆ **Loading and unloading data** The `LOAD TABLE`, `UNLOAD TABLE`, and `UNLOAD` statements all access the file system on the database server machine. If you are running the personal database server, you already have access to the file system and this is not a security issue. If you are running the network database server, unwarranted file system access may be a security issue. The `-gl` option allows you to control the database permissions required to carry out loading and unloading of data. The permissible values are **DBA**, **all**, or **none**.

☞ For more information, see "`-gl` server option" on page 141.

- ◆ **Encrypting client/server communications over the network** For greater security, you can force client/server network communications to be encrypted as they pass over the network.

☞ For more information, see "Encrypting client/server communications" on page 398.

Encrypting client/server communications

Client/server communication encryption is a separately licensable component and must be ordered before you can install it. To order this component, see the card in your SQL Anywhere Studio package or see <http://www.sybase.com/detail?id=1015780>.

You can set client/server encryption when you start the database server or in the client connection parameters. You can encrypt all native Adaptive Server Anywhere packets (Embedded SQL, ODBC, and OLEDB) that are transmitted to and from all clients. TDS packets (Java connections, including Sybase Central and Interactive SQL, as well as Open Client connections) are not encrypted.

When you use strong encryption by specifying `-ec ECC_TLS` or `RSA_TLS` in the server command, all connections to the server must perform a Certicom handshake. This handshake cannot be faked and Certicom encryption ensures that invalid packets, which may be intended to harm the server, are discarded.

❖ To force encryption of client/server communications from the server:

- ◆ Start the database server using the `-ec` option. For example:

```
dbsrv8 -ec simple,ECC_TLS -x tcpip "c:\Program  
Files\Sybase\SQL Anywhere 8\asademo.db"
```

🔗 For more information, see "-ec server option" on page 135.

❖ To force encryption of client/server communications from a particular client:

- ◆ Add the **Encryption (ENC)** connection parameter to your connection string.

```
"UID=DBA;PWD=SQL;ENG=myeng;LINKS=tcpip;  
Encryption=ECC_TLS  
(trusted_certificates=sample.crt)"
```

You can also set this parameter on the Advanced tab of the Connect dialog and on the Network tab of the ODBC data source dialog.

🔗 For more information, see "Encryption connection parameter" on page 177.

Encrypting a database

As a database administrator, you can use database encryption to make it more difficult for someone to decipher the data in your database. You can choose to secure your database either with simple or with strong encryption.

Simple encryption

Simple encryption is equivalent to obfuscation and makes it more difficult for someone using a disk utility to look at the file to decipher the data in your database. Simple encryption does not require a key to encrypt the database. Simple encryption technology is supported in previous versions of SQL Anywhere.

Strong encryption

Strong database file encryption technology makes the database inoperable without the key (password). As well, it scrambles the information contained in your database and transaction log files so they cannot be deciphered simply by looking at the files using a disk utility. The data is completely inaccessible without the key.

Two algorithms have been chosen to implement strong encryption: AES, a block encryption algorithm chosen as the new Advanced Encryption Standard (AES) for block ciphers by the National Institute of Standards and Technology (NIST); and MDSR, a new algorithm developed by Casio.

A database can be strongly encrypted using the `ENCRYPTION` and `KEY` options with the `CREATE DATABASE` statement. Similarly, the database administrator can initialize a database using the `dbinit` utility in combination with various options to enable strong encryption. You can also use the Sybase Central Create Database wizard to create a strongly encrypted database. Using the `dbinit` utility with the `-ea` option enables strong encryption and sets the algorithm to either AES or MDSR. Using the `dbinit` utility in combination with the `-ek` or `-ep` option enables strong encryption and indicates whether the key is to be specified in a prompt box or at the command prompt.

❖ To create a strongly encrypted database (SQL):

- 1 Connect to an existing database from Interactive SQL.
- 2 Execute a `CREATE DATABASE` statement that includes the `ENCRYPTION` and `KEY` options. For example, the following statement creates a database file named *myencrypteddb.db* in the C:\ directory.

```
CREATE DATABASE 'c:\myencrypteddb'
TRANSACTION LOG ON
ENCRYPTED ON
KEY '0kZ2o52AK#'
ALGORITHM 'MDSR'
```

❖ To create a strongly encrypted database (command prompt):

- 1 At a command prompt, use the *dbinit* utility to create a database. You must include the following options:
 - ◆ `-ea` to specify the encryption algorithm.
 - ◆ `-ek` or `-ep` to specify the encryption key and whether you want to enter it at the command prompt or in a dialog box.

The following command (entered all on one line) creates a strongly encrypted database and specifies that the encryption key is entered as part of the command.

```
dbinit -ea MDSR -ek "0kZ2o56AK#" "myencrypteddb.db"
```

- 2 Start the database from the command prompt.

```
dbeng8 myencrypteddb.db -ek "0kZ2o56AK#"
```

🔗 For more information about the encryption key, see "Encryption Key connection parameter" on page 179.

As with most passwords, it is best to choose a key value that cannot be easily guessed. We recommend that you choose a value for your key that includes between 8 and 30 digits, a combination of upper and lower case characters, and numbers, letters, and special characters.

Caution


Be sure to store a copy of your key in a safe location. You require the key each time you want to start or modify the database. A lost key will result in a completely inaccessible database, from which there is no recovery.

Controlling strong encryption

In Adaptive Server Anywhere, the database administrator has control over four aspects of strong encryption, including: strong encryption status, the encryption key, protection of the encryption key, and the encryption algorithm.

Strong encryption status

Although you can't simply turn strong encryption on or off in an existing database, you can choose from two options when it comes to implementing strong encryption. You can either create a database from scratch with strong encryption, or you can rebuild an existing database and change the encryption status at that time. Rebuilding the database unloads all of the data and schema of an existing database, creates a new database (at which point you can change a variety of settings including strong encryption status), and reloads the data into the new database. You need to know the key to unload a strongly encrypted database.

 For more information on these features, see


- ◆ "Reloading a Database" on page 444 of the book *ASA SQL User's Guide*
- ◆ "CREATE DATABASE statement" on page 273 of the book *ASA SQL Reference Manual*

The encryption key

As with most passwords, it is best to choose a key value that cannot be easily guessed. The key can be of arbitrary length, but generally the longer the key, the better because a shorter key is easier to guess than a longer one. As well, including a combination of numbers, letters, and special characters decreases the chances of someone guessing the key. You must supply this key each time you want to start the database. Lost or forgotten keys result in completely inaccessible databases.

Protection of the encryption key

You can choose whether the encryption key is entered at the command prompt (the default) or into a prompt box. Choosing to enter the key in a prompt box provides an extra measure of security because the key is never visible in plain sight. Clients are required to specify the key each time they start the database. In cases where the database administrator starts the database, clients never need to have access to the key.

 For more information, see "-ep server option" on page 137.

The encryption algorithm

When you strongly encrypt a database, you can choose to use either the AES algorithm (the default if an algorithm isn't specified explicitly) or MDSR.

AES has recently been through a period of international evaluation and has now been chosen as the new Advanced Encryption Standard block cipher algorithm. It has many properties that lend itself well to encryption of Adaptive Server Anywhere databases in terms of performance and size.

☞ For more information about database encryption algorithms, see:

- ◆ "Initialization utility options" on page 467
- ◆ "CREATE DATABASE statement" on page 273 of the book *ASA SQL Reference Manual*

Performance issues

Performance of Adaptive Server Anywhere is somewhat slower when the database is encrypted. The performance impact depends on how often pages are read from or written to disk, and can be minimized by ensuring that the server is using an adequate cache size.

You can increase the starting size of the cache with the `-c` option when you start the server. For operating systems that support dynamic resizing of the cache, the cache size that is used may be restricted by the amount of memory that is available; to increase the cache size, increase the available memory.

☞ For more information, see:

- ◆ "Using the cache to improve performance" on page 152 of the book *ASA SQL User's Guide*
- ◆ "`-c` server option" on page 127

Security tips

As database administrator, there are many actions you can take to improve the security of your data. For example, you can:

- ◆ **Change the default user ID and password** The default user ID and password for a newly created database is **DBA** and **SQL**. You should change this password before deploying the database.
- ◆ **Require long passwords** You can set the `MIN_PASSWORD_LENGTH` public option to disallow short (and therefore easily guessed) passwords.

🔗 For information, see "MIN_PASSWORD_LENGTH option" on page 583.

- ◆ **Restrict DBA authority** You should restrict DBA authority only to users who absolutely require it since it is very powerful. Users with DBA authority can see and do anything in the database.

You may consider giving users with DBA authority two user IDs: one with DBA authority and one without, so they can connect as DBA only when necessary.

- ◆ **Drop external system functions** The following external functions present possible security risks: `xp_cmdshell`, `xp_startmail`, `xp_startsmtp`, `xp_sendmail`, `xp_stopmail`, and `xp_stopsmtp`.

The `xp_cmdshell` procedure allows users to execute operating system commands or programs.

The e-mail commands allow users to have the server send e-mail composed by the user. Malicious users could use either the e-mail or command shell procedures to perform operating-system tasks with authorities other than those they have been given by the operating system. In a security-conscious environment, you should drop these functions.


🔗 For information on dropping procedures, see "DROP statement" on page 397 of the book *ASA SQL Reference Manual*.

- ◆ **Protect your database files** You should protect the database file, log files, dbspace files, and write files from unauthorized access. Do not store them within a shared directory or volume.
- ◆ **Protect your database software** You should similarly protect Adaptive Server Anywhere software. Only give users access to the applications, DLLs, and other resources they require.

- ◆ **Run the database server as a service or a daemon** To prevent unauthorized users from shutting down or gaining access to the database or log files, run the database server as a Windows service. On UNIX, running the server as a daemon serves a similar purpose.

 For more information, see "Running the server outside the current session" on page 21.

- ◆ **Set ASTMP to a unique directory** To make the engine secure on UNIX platforms, set ASTMP to a unique directory, and make the directory read, write, and execute protected against all other users. Doing so forces all connections to use TCP/IP, which is more secure than the shared memory connection.
- ◆ **Strongly encrypt your database** Strongly encrypting your database makes it completely inaccessible without the key. You cannot open the database, or view the database or transaction log files using any other means.

 For more information, see "-ep server option" on page 137 and "-ek database option" on page 159.

Replicating Data with Replication Server

About this chapter


This chapter describes how you can use Replication Server to replicate data between an Adaptive Server Anywhere database and other databases. Other databases in the replication system may be Adaptive Server Anywhere databases or other kinds of database.

Contents

Topic	Page
Introduction to replication	406
Tutorial: Replicate data using Replication Server	409
Configuring databases for Replication Server	419
Using the LTM	422

Before you begin

Replication Server administrators who are setting up Adaptive Server Anywhere to take part in their Replication Server installation will find this chapter especially useful. You should have knowledge of Replication Server documentation, and familiarity with the Replication Server product. This chapter does not describe Replication Server itself.

 For information about Replication Server, including design, commands, and administration, see your Replication Server documentation.

Introduction to replication

Data replication is the sharing of data among physically distinct databases. Changes made to shared data at any one database are copied precisely to the other databases in the replication system.

Data replication brings some key benefits to database users.

Data availability	Replication makes data available locally, rather than through potentially expensive, less reliable and slower connections to a single, central database. Since data is accessible locally, you always have access to data, even in the event of a long-distance network connection failure.
Response time	Replication improves response times for data requests for two reasons. First, retrieval rates are faster since requests are processed on a local server without accessing some wide area network. Second, competition for processor time decreases, since local processing offloads work from a central database server.

Sybase replication technologies

Sybase provides the following replication technologies for Adaptive Server Anywhere:

- ◆ **MobiLink** MobiLink is designed for two-way replication involving a consolidated database and large numbers of remote databases, typically including many mobile databases.
- ◆ **SQL Remote** Like MobiLink, SQL Remote is designed for two-way replication involving a consolidated database and large numbers of remote databases, typically including many mobile databases. Administration and resource requirements at the remote sites are minimal.
- ◆ **Replication Server** Replication Server is designed for replication among relatively small numbers of data servers, with a typical time lag between primary data and replicate data of a few seconds, and generally with an administrator at each site.

Each replication technology has its own documentation. This chapter describes how to use Adaptive Server Anywhere with Replication Server.

☞ For information about SQL Remote, see the book *SQL Remote User's Guide*. For information about MobiLink, see the book *MobiLink Synchronization User's Guide*.

Replicate sites and primary sites


In a Replication Server installation, the data to be shared among databases is arranged in **replication subscriptions**.

For each replication definition, there is a **primary site**, where changes to the data in the replication occur. The sites that receive the data in the replication are called **replicate sites**.

Replicate site components

You can use Adaptive Server Anywhere as a replicate site with no additional components.

The following diagram illustrates the components required for Adaptive Server Anywhere to participate in a Replication Server installation as a replicate site.



- ◆ Replication Server receives data changes from primary site servers.
- ◆ Replication Server connects to Adaptive Server Anywhere to apply the changes.
- ◆ Adaptive Server Anywhere makes the changes to the database.


Asynchronous procedure calls

The Replication Server can use asynchronous procedure calls (APC) at replicate sites to alter data at a primary site database. If you are using APCs, the above diagram does not apply. Instead, the requirements are the same as for a primary site.

Primary site components

To use an Adaptive Server Anywhere database as a primary site, you need to use the Log Transfer Manager (LTM), or Replication Agent. The LTM supports Replication Server version 10.0 and greater.

The following diagram illustrates the components required for Adaptive Server Anywhere to participate in a Replication Server installation as a primary site. The arrows in the diagram represent data flow.



- ◆ The Adaptive Server Anywhere database server manages the database.
- ◆ The Adaptive Server Anywhere Log Transfer Manager connects to the database. It scans the transaction log to pick up changes to the data, and sends them to Replication Server.
- ◆ Replication Server sends the changes to replicate site databases.

Tutorial: Replicate data using Replication Server

This section provides a step-by-step tutorial describing how to replicate data from a primary database to a replicate database. Both databases in the tutorial are Adaptive Server Anywhere databases.

Replication Server assumed

This section assumes you have a running Replication Server. For more information about how to install or configure Replication Server, see the Replication Server documentation.

What is in the tutorial

This tutorial describes how to replicate only tables. For information about replicating procedures, see "Preparing procedures and functions for replication" on page 424.

The tutorial uses a simple example of a (very) primitive office news system: a single table with an ID column holding an integer, a column holding the user ID of the author of the news item, and a column holding the text of the news item. The **id** column and the **author** column make up the primary key.

Before you work through the tutorial, create a directory (for example, *c:\tutorial*) to hold the files you create in the tutorial.

Lesson 1: Set up the Adaptive Server Anywhere databases

This section describes how to create and set up the Adaptive Server Anywhere databases for replication.

You can create a database using Sybase Central or the *dbinit* utility. For this tutorial, we use the *dbinit* utility.

❖ Create the primary site database:

- ◆ Enter the following command from the tutorial directory you created (for example *c:\tutorial*).

```
dbinit primedb
```

This creates a database file *primedb.db* in the current directory.

❖ Create the replicate site database:

- ◆ Enter the following command from the tutorial directory you created (for example *c:\tutorial*).

```
dbinit repdb
```

This creates a database file *repdb.db* in the current directory.

What's next?

Next, you have to start database servers running on these databases.

Lesson 2: Start the database servers

You need to run the primary site database server, with the primary database loaded.

❖ Start the primary site database server:

- 1 Change to the tutorial directory.
- 2 Enter the following command to start a network database server running the **primedb** database. You should be using the TCP/IP network communication protocol on the default communications port (2638):

```
dbsrv8 -x tcpip primedb.db
```

❖ Start the replicate site database server:

- 1 Change to the tutorial directory.
- 2 Enter the following command to start a network database server running the **repdb** database, but on a different port:

```
dbsrv8 -x tcpip(port=2639) -n REPSV repdb.db
```

What's next?


Next, you have to make entries for each of the Adaptive Server Anywhere servers in an interfaces file, so Replication Server can communicate with these database servers.

Lesson 3: Set up the Open Servers in your system

You need to add a set of Open Servers to the list of Open Servers in your system.

Adding Open Servers

Open Servers are defined in your interfaces file (*SQL.ini*) using the *DSEdit* utility. For NetWare and UNIX users, the interfaces file is named *interfaces*, and the utility is named *sybinit*.

 For full instructions on how to add definitions to your interfaces file, see "Configuring Open Servers" on page 110.

Required Open Servers

For each Open Server definition you must provide a **name** and an **address**. Do not alter the other attributes of the definition. You need to add an Open Server entry for each of the following:

- ◆ **The primary database** Create an entry named PRIMEDB with address as follows:
 - ◆ **Protocol** NLWNSCK
 - ◆ **Network address** localhost,2638
- ◆ **The replicate database** Create an entry named REPDB with address as follows:
 - ◆ **Protocol** NLWNSCK
 - ◆ **Network address** localhost,2639
- ◆ **The LTM at the primary database** This is necessary so you can shut down the LTM properly. Create an entry named PRIMELTM with address as follows:
 - ◆ **Protocol** NLWNSCK
 - ◆ **Network address** localhost,2640
- ◆ **Your Replication Server** This tutorial assumes you already have the Replication Server Open Server defined.

What's next?

Next, confirm that the Open Servers are configured properly.

Lesson 4: Confirm the Open Servers are configured properly

You can confirm that each Open Server is available by selecting ServerObject ► Ping Server from the DSEdit utility.

Alternatively, you can confirm that each Open Server is configured properly by connecting to the database using an Open Client application such as the *isql* utility.

To start *isql* running on the primary site database, type

```
isql -U DBA -P SQL -S PRIMEDB
```

The Open Client *isql* utility is not the same as the Adaptive Server Anywhere Interactive SQL utility.

Lesson 5: Add Replication Server information to the primary database

You need to add Replication Server tables and procedures to the primary site database for the database to participate in a Replication Server installation. You also need to create two user IDs for use by Replication Server. The SQL command file *rssetup.sql* comes with Adaptive Server Anywhere, and carries out these tasks.

The *rssetup.sql* command file must be run on the Adaptive Server Anywhere server from the Interactive SQL utility.

❖ Run the *rssetup* script:

- 1 From Interactive SQL, connect to the Adaptive Server Anywhere database as user ID **DBA** using password **SQL**.
- 2 Run the *rssetup* script using the following command:

```
read "path\rssetup.sql"
```

where *path* is your Adaptive Server Anywhere installation directory.

You can alternatively use File►Run Script, and browse to the file.

Actions carried out
by *rssetup.sql*

The *rssetup.sql* command file carries out the following functions:

- ◆ Creates a user named **dbmaint**, with password **dbmaint** and with DBA permissions. This is the maintenance user name and password required by Replication Server to connect to the primary site database.
- ◆ Creates a user named **sa**, with password **sysadmin** and with DBA permissions. This is the user ID used by Replication Server when materializing data.
- ◆ Adds **sa** and **dbmaint** to a group named **rs_systabgroup**.

Passwords and
user IDs

While the hard-wired user IDs (**dbmaint** and **sa**) and passwords are useful for test and tutorial purposes, you should change the password and perhaps also the user IDs when running databases that require security. Users with DBA permissions have full authority in an Adaptive Server Anywhere database.

The user ID **sa** and its password must match that of the system administrator account on the Replication Server. Adaptive Server Anywhere does not currently accept A NULL password.

Permissions

The *rssetup.sql* script carries out a number of operations, including some permissions management. The permissions changes made by *rssetup.sql* are outlined here. *You do not have to make these changes yourself.*

For replication, ensure that the **dbmaint** and **sa** users can access this table without explicitly specifying the owner. To do this, the table owner user ID must have group membership permissions, and the **dbmaint** and **sa** users must be members of the table owner group. To grant group permissions, you must have DBA authority.

For example, if user DBA owns the table, you should grant group permissions to the DBA user ID:

```
GRANT GROUP
TO DBA
```

You should then grant the **dbmaint** and **sa** users membership in the DBA group. To grant group membership, you must either have DBA authority or be the group ID.

```
GRANT MEMBERSHIP
IN GROUP "DBA"
TO dbmaint ;

GRANT MEMBERSHIP
IN GROUP "DBA"
TO sa ;
```

Lesson 6: Create the table for the primary database

In this section, you create a single table in the primary site database, using *isql*. First, make sure you are connected to the primary site database:

```
isql -U DBA -P SQL -S PRIMEDB
```

Next, create a table in the database:

```
CREATE TABLE news (
  ID int,
  AUTHOR char( 40 ) DEFAULT CURRENT USER,
  TEXT char( 255 ),
  PRIMARY KEY ( ID, AUTHOR )
)
go
```

Identifier case sensitivity

In Adaptive Server Anywhere, all identifiers are case insensitive. In Adaptive Server Enterprise, identifiers are case sensitive by default. Even in Adaptive Server Anywhere, ensure the case of your identifiers matches in all parts of the SQL statement to ensure compatibility with Adaptive Server Enterprise.

In Adaptive Server Anywhere, the database determines case sensitivity. For example, passwords are case insensitive in case insensitive databases, and case sensitive in case sensitive databases, while user IDs, being identifiers, are case insensitive in all Adaptive Server Anywhere databases.

For **news** to act as part of a replication primary site, you must set the REPLICATE flag to ON for the table using an ALTER TABLE statement:

```
ALTER TABLE news
REPLICATE ON
go
```

This is equivalent to running the **sp_setreplicate** or **sp_setreptable** procedure on the table in Adaptive Server Enterprise. You cannot set REPLICATE ON in a CREATE TABLE statement.

Lesson 7: Add Replication Server information to the replicate database

You should run the *rssetup.sql* command file on the replicate database in exactly the same manner as it ran on the primary database. Also ensure that the **dbmaint** and **sa** users can access this table without explicitly specifying the table owner.

☞ These tasks are the same as those carried out on the primary database. For a complete explanation, see "Lesson 5: Add Replication Server information to the primary database" on page 412.

Lesson 8: Create the tables for the replicate database

The replicate site database needs to have tables to hold the data it receives. Now is a good time to create these tables. As long as the database elements are in place, no extra statements are necessary for them to act as a replicate site in a Replication Server installation. In particular, you do not need to set the REPLICATE flag to ON, which is necessary only at the primary site.

Replication Server allows replication between tables and columns with different names. As a simple example, however, create a table in the replicate database identical in definition to that in the primary database (except for the **REPLICATE** flag, which is not set to **ON** in the replicate database). The table creation statement for this is:

```
CREATE TABLE news (
    ID int,
    AUTHOR char( 40 ) DEFAULT CURRENT USER,
    TEXT char( 255 ),
    PRIMARY KEY ( ID, AUTHOR )
)
go
```

For the tutorial, the **CREATE TABLE** statement must be *exactly* the same as that at the primary site.

You must ensure that the users **dbmaint** and **sa** can access this table without specifying the owner name. Also, these user IDs must have **SELECT** and **UPDATE** permissions on the table.

Lesson 9: Set up Replication Server

You need to carry out the following tasks on the Replication Server:

- ◆ Create a connection for the primary site data server.
- ◆ Create a connection for the replicate site data server.
- ◆ Create a replication definition.
- ◆ Create a subscription to the replication.

This section describes each of these tasks. It also describes starting the Adaptive Server Anywhere LTM.

Create a connection for the primary site

Using *isql*, connect to Replication Server and create a connection to the primary site Adaptive Server Anywhere database.

The following command creates a connection to the **primedb** database on the **PRIMEDB** Open Server.

```
create connection to PRIMEDB.primedb
set error class rs_sqlserver_error_class
set function string class rs_sqlserver_function_class
set username dbmaint
set password dbmaint
with log transfer on
go
```

If you have changed the **dbmaint** user ID and password in the *rssetup.sql* command file, make sure you replace the **dbmaint** username and password in this command.

Replication Server does not actually use the **primedb** database name; instead, the database name is read from the command line of the **PRIMEDB** Open Server. You must, however, include a database name in the **CREATE CONNECTION** statement to conform to the syntax.

☞ For a full description of the create connection statement, see the chapter "Replication Server Commands" in the *Replication Server Reference Manual*.

Create a connection for the replicate site.

Using *isql*, connect to Replication Server and create a connection to the replicate site Adaptive Server Anywhere database.

The following command creates a connection to the **repdb** database on the **REPDB** Open Server.

```
create connection to REPDB.repdb
set error class rs_sqlserver_error_class
set function string class rs_sqlserver_function_class
set username dbmaint
set password dbmaint
go
```

This statement differs from the primary site server statement in that there is no **with log transfer on** clause in this statement.

If you have changed the **dbmaint** user ID and password in the *rssetup.sql* command file, make sure you replace the **dbmaint** username and password in this command.

Create a replication definition

Using *isql*, connect to Replication Server and create a replication definition. The following statement creates a replication definition for the news table on the **primedb** database:

```
create replication definition news
with primary at PRIMEDB.primedb
( id int, author char(40), text char(255) )
primary key ( id, author )
go
```

For a full description of the **CREATE REPLICATION DEFINITION** statement, see your *Replication Server Reference Manual*.

If you set the **qualify_table_owner** option to on in the LTM configuration file, you must specify the table owner in the statement, for all replicating tables.

Configure and start the Adaptive Server Anywhere LTM

For replication to take place, the Adaptive Server Anywhere LTM must be running against the primary site server. Before you start the Adaptive Server Anywhere LTM, make sure it is properly configured by editing an LTM configuration file.

Below is a sample configuration file for the **primedb** database. If you are following the examples, you should make a copy of this file as *primeltn.cfg*:

```
#
# Configuration file for 'PRIMELTM'
#
SQL_server=PRIMEDB
SQL_database=primedb
SQL_user=sa
SQL_pw=sysadmin
RS_source_ds=PRIMEDB
RS_source_db=primedb
RS=your_rep_server_name_here
RS_user=sa
RS_pw=sysadmin
LTM_admin_user=DBA
LTM_admin_pw=SQL
LTM_charset=cp850
scan_retry=2
APC_user=sa
APC_pw=sysadmin
SQL_log_files=C:\TUTORIAL
```

If you have changed the user ID and password in the *rssetup.sql* command file from **sa** and **sysadmin**, you should use the new user ID and password in this configuration.

To start the Adaptive Server Anywhere LTM running on the primary site server, enter the following command:

```
dbltm -S PRIMELTM -C primeltn.cfg
```

The connection information is in *primeltn.cfg*. In this command, PRIMELTM is the server name of the LTM.

You can find usage information about the Adaptive Server Anywhere LTM by typing the following statement:

```
dbltm -?
```

You can run the Adaptive Server Anywhere LTM as a Windows service. For information on running programs as services, see "Running the server outside the current session" on page 21.

Create a subscription for your replication

Using *isql*, connect to Replication Server and create a subscription for the replication.

The following statement creates a subscription for the news replication defined in "Create a replication definition" on page 416 with replicate site as the **repdb** database.

```
create subscription NEWS_SUBSCRIPTION
for news
with replicate at REPDB.repdb
go
```

You have now completed your installation. Try replicating data to confirm that the setup is working properly.

Lesson 10: Enter data at the primary site for replication

You can now replicate data from the primary database to the replicate database. As an example, connect to the primary database using the *isql* utility, and enter a row in the **news** table.

```
insert news (id, text)
values (1, 'Test news item.' )
commit
go
```

The Adaptive Server Anywhere LTM sends only committed changes to the Replication Server. The data change is replicated next time the LTM polls the transaction log.

Tutorial complete

You have now completed the tutorial.


Configuring databases for Replication Server

Each Adaptive Server Anywhere database that participates in a Replication Server installation needs to be configured before it can do so. Configuring the database involves the following tasks:

- ◆ Selecting a secure user ID for the maintenance user and the name used by Replication Server when materializing data.
- ◆ Setting up the database for Replication Server.
- ◆ Configuring the language and character set, where necessary.

Configuring the LTM

Each primary site Adaptive Server Anywhere database requires an LTM to send data to Replication Server. Each primary or replicate site Adaptive Server Anywhere database requires an Open Server definition so that Replication Server can connect to the database.

 For information on configuring the LTM, see "Configuring the LTM" on page 425.

Setting up the database for Replication Server

Once you have created your Adaptive Server Anywhere database, and created the necessary tables and so on within the database, you can make database ready for use with Replication Server. You do this using a setup script supplied with the Adaptive Server Anywhere Replication Agent product. The script is named *rssetup.sql*.

When you need to run the setup script

You need to run the setup script at any Adaptive Server Anywhere database that is taking part in a Replication Server installation, whether as a primary or a replicate site.

What the setup script does

The setup script creates user IDs required by Replication Server when connecting to the database. It also creates a set of stored procedures and tables used by Replication Server. The tables begin with the characters **rs_**, and the procedures begin with the characters **sp_**. Procedures include some that are important for character set and language configuration.

Prepare to run the setup script

Replication Server uses a special data server **maintenance user** login name for each local database containing replicated tables. This allows Replication Server to maintain and update the replicated tables in the database.

The maintenance user

The setup script creates a maintenance user with name **dbmaint** and password **dbmaint**. The maintenance user has DBA permissions in the Adaptive Server Anywhere database, which allows it full control over the database. For security reasons, you should change the maintenance user ID and password.

❖ To change the maintenance user ID and password:

- 1 Open the *rssetup.sql* setup script in a text editor. The script is held in the *scripts* subdirectory of your Adaptive Server Anywhere installation directory.
- 2 Change all occurrences of the *dbmaint* user ID to the new maintenance user ID of your choice.
- 3 Change the **dbmaint** password to the new maintenance user password of your choice. The password occurs in the following place at the top of the setup script file:

```
GRANT CONNECT TO dbmaint  
IDENTIFIED BY dbmaint
```

The materialization user ID

When Replication Server connects to a database to materialize the initial copy of the data in the replication, it does so using the Replication Server system administrator account.

The Adaptive Server Anywhere database must have a user ID and password that match the Replication Server system administrator user ID and password. Adaptive Server Anywhere does not accept a NULL password.

The setup script assumes a user ID of *sa* and a password of **sysadmin** for the Replication Server administrator. You should change this to match the actual name and password.

❖ To change the system administrator user ID and password:

- 1 Open the *rssetup.sql* setup script in a text editor.
- 2 Change all occurrences of the **sa** user ID to match the Replication Server system administrator user ID.
- 3 Change the **sa** password to match the Replication Server system administrator password. The password has the initial setting of **sysadmin**.

Run the setup script

Once you have modified the setup script to match the user IDs and passwords appropriately, you can run the setup script to create the maintenance and system administrator users in the Adaptive Server Anywhere database.

❖ To run the setup script:

- 1 Start the Adaptive Server Anywhere database on an Adaptive Server Anywhere database engine or server.
- 2 Start the Interactive SQL utility, and connect to the database as a user with DBA authority. When you create an Adaptive Server Anywhere database, it has a user with user ID **DBA** and password **SQL**, which has DBA authority.
- 3 Run the script by entering the following command in the SQL Statements pane:

```
read path\rssetup.sql
```

where *path* is the path to the setup script.

Character set and language issues

Upon creation, each Adaptive Server Anywhere database is assigned a specific collation (character set and sort order). Replication Server uses a different set of identifiers for character sets and sort orders.

Set the character set and language parameters in the LTM configuration file. If you are unsure of the characters set label to specify, you can do the following to determine the character set of the server:

❖ To determine the character set:

- ◆ Execute the following command:

```
exec sp_serverinfo csname
```

The language is one of the language labels listed in "Language label values" on page 263.

Using the LTM

Since the Adaptive Server Anywhere LTM relies on information in the Adaptive Server Anywhere transaction log, take care not to delete or damage the log without storing backups (for example, using a transaction log mirror).

✍ For more information about transaction log management, see the section "Transaction log and backup management" on page 429.

You cannot substitute an Adaptive Server Anywhere LTM for an Adaptive Server Enterprise LTM since the transaction logs have different formats.

The Adaptive Server Anywhere LTM supports replication of inserts, updates, and deletes, as well as replication of Transact-SQL-dialect stored procedure calls.

The Adaptive Server Enterprise LTM sends data changes to the Replication Server before they are committed. The Replication Server holds the changes until a COMMIT statement arrives. By contrast, the Adaptive Server Anywhere LTM sends only committed changes to Replication Server. For long transactions, this may lead to some added delay in replication, since all changes have to go through the Replication Server before distribution.

Configuring tables for replication

Adaptive Server Anywhere does not support the **sp_setreplicate** system procedure. Instead, a table is identified as a primary data source using the ALTER TABLE statement with a single clause:

```
ALTER TABLE table-name  
SET REPLICATE ON
```

The effects of
setting
REPLICATE ON
for a table

Setting REPLICATE ON places extra information into the transaction log. Whenever an UPDATE, INSERT, or DELETE action occurs on the table. The Adaptive Server Anywhere Replication Agent uses this extra information to submit the full pre-image of the row, where required, to Replication Server for replication.

Even if only some of the data in the table needs to be replicated, all changes to the table are submitted to Replication Server. It is Replication Server's responsibility to distinguish the data to be replicated from that which is not.

When you update, insert, or delete a row, the pre-image of the row is the contents of the row before the action, and the post-image is the contents of the row after the action. For INSERTS, only the post-image is submitted (the pre-image is empty). For DELETES, the post-image is empty and only the pre-image is submitted. For UPDATES, both the pre-image and the updated values are submitted.

The following data types are supported for replication:

Data type	Description (Open Client/Open Server type)
Exact integer data types	int, smallint, tinyint
Exact decimal data types	decimal, numeric
Approximate numeric data types	float (8-byte), real
Money data types	money, smallmoney
Character data types	char(n), varchar(n), text
Date and time data types	datetime, smalldatetime
Binary data types	binary(n), varbinary(n), image
Bit data types	bit

Notes

Adaptive Server Anywhere supports data of zero length that is not NULL. However, non-null long varchar and long binary data of zero length is replicated to a replicate site as NULL.

If a primary table has columns with unsupported data types, you can replicate the data if you create a replication definition using a compatible supported data type. For example, to replicate a DOUBLE column, you could define the column as FLOAT in the replication definition.


Side effects of setting REPLICATE ON for a table

There can be a replication performance hit for heavily updated tables. You could consider using replicated procedures if you experience performance problems that may be related to replication traffic, since replicated procedures send only the call to the procedure instead of each individual action.

Since setting REPLICATE ON sends extra information to the transaction log, this log grows faster than for a non-replicating database.

Minimal column replication definitions

The Adaptive Server Anywhere LTM supports the Replication Server replicate minimal columns feature. This feature is enabled at Replication Server.

 For more information on replicate minimal columns, see your Replication Server documentation.

Preparing procedures and functions for replication

You can use stored procedures to modify the data in tables. Updates, inserts, and deletes execute from within the procedure.

Replication Server can replicate procedures as long as they satisfy certain conditions. The first statement in a procedure must carry out an update for the procedure to be replicated. See your Replication Server documentation for a full description of how Replication Server replicates procedures.

Adaptive Server Anywhere supports two dialects for stored procedures: the Watcom-SQL dialect, based on the draft ISO/ANSI standard, and the Transact-SQL dialect. You can use either dialect in writing stored procedures for replication.

Function APC format

The Adaptive Server Anywhere LTM supports the Replication Server **function APC** format. To make use of these functions, set the configuration parameter **rep_func** to **on** (the default is **off**).

The LTM interprets all replicated APCs as either table APCs or function APCs. A single Adaptive Server Anywhere database cannot combine function APCs with other table APCs.

☞ For more information about replicate functions, see your Replication Server documentation.

SQL Statements for controlling procedure replication

A procedure can be configured to act as a replication source using the **ALTER PROCEDURE** statement.

The following statement makes the procedure **MyProc** act as a replication source.

```
ALTER PROCEDURE MyProc
REPLICATE ON
```

The following statement prevents the procedure **MyProc** from acting as a replication source.

```
ALTER PROCEDURE MyProc
REPLICATE OFF
```

These statements have the same effect as running **sp_setrepl** or **sp_setrepproc 'table'** on the procedure in Adaptive Server Enterprise. The **sp_setrepproc 'function'** syntax is not supported.

The effects of
setting
REPLICATE ON
for a procedure

When a procedure is used as a replication data source, calling the procedure sends extra information to the transaction log.

Asynchronous procedures

A procedure called at a replicate site database to update data at a primary site database is an **asynchronous procedure**. The procedure carries out no action at the replicate site, but rather, the call to the procedure is replicated to the primary site, where a procedure of the same name executes. This is called an **asynchronous procedure call** (APC). The changes made by the APC are then replicated from the primary to the replicate database in the usual manner.

 For information about APCs, see your Replication Server documentation.

The APC_user and
APC support

Support for APCs in Adaptive Server Anywhere is different from that in Adaptive Server Enterprise. In Adaptive Server Enterprise, each APC executes using the user ID and password of the user who called the procedure at the replicate site. In Adaptive Server Anywhere, however, the transaction log does not store the password, and so it is not available at the primary site. To work around this difference, the LTM configuration file holds a single user ID with associated password, and this user ID (the **APC_user**) executes the procedure at the primary site. The APC_user must, therefore, have appropriate permissions at the primary site for each APC that may be called.

Configuring the LTM

You control LTM behavior by modifying the LTM **configuration file**, which is a plain text file created and edited using a text editor. The LTM configuration file contains information the LTM needs, such as the Adaptive Server Anywhere server it transfers a log from, the Replication Server it transfers the log to. You need a valid configuration file to run the LTM.

Creating a
configuration file

You must create a configuration file, using a text editor, before you can run the LTM. The `-c` LTM command specifies the name of the configuration file to use, and has a default of *dbltm.cfg*.

Configuration file
format

The LTM configuration file shares the same format as the Replication Server configuration file described in your *Replication Server Administration Guide*. In summary:

- ◆ The configuration file contains one entry per line.

- ◆ An entry consists of a parameter, followed by the = character, followed by the value:

Entry=value

- ◆ Lines beginning with a # character are comments ignored by the LTM.
- ◆ The configuration file cannot contain leading blanks.
- ◆ Entries are case sensitive.

☞ For the full list of available configuration file parameters, see "The LTM configuration file" on page 483.

Example configuration file

- ◆ The following is a sample Adaptive Server Anywhere LTM configuration file.

```
# This is a comment line
# Names are case sensitive.
SQL_user=sa
SQL_pw=sysadmin
SQL_server=PRIMESV
SQL_database=primedb
RS_source_ds=PRIMESV
RS_source_db=primedb
RS=MY_REPSEVER
RS_user=sa
RS_pw=sysadmin
LTM_admin_user=DBA
LTM_admin_pw=SQL
LTM_charset=cp850
scan_retry=2
SQL_log_files=e:\logs\backup
APC_user=sa
APC_pw=sysadmin
```

Replicating transactions in batches

Effects of buffering transactions

The LTM allows buffering of replication commands to Replication Server. Buffering the replication commands and sending them in batches results in fewer messages being sent, and can significantly increase overall throughput, especially on high volume installations.

How batch mode works

By default, the LTM buffers transactions. The buffer flushes (the transactions sent to Replication Server when the buffer:

- ◆ **Reaches maximum number of commands** The **batch_ltl_sz** parameter sets the maximum number of LTL (log transfer language) commands stored in the buffer before it flushes. The default setting is 200.

- ◆ **Reaches maximum memory used** The **batch_ltl_mem** parameter sets the maximum memory that the buffer can occupy before flushes. The default setting is 256K.
- ◆ **Completes transaction log processing** If there are no more entries in the transaction log to process (that is, the LTM is up to date with all committed transactions), then the buffer flushes.

Turning off
buffering

You can turn off buffering of transactions by setting the **batch_ltl_cmds** parameter to **off**:

```
batch_ltl_cmds=off
```

Language and character set issues

Language and character set issues are an important consideration in many replication sites. Each database and server in the system uses a specific collation (character set and sorting order) for storing and ordering strings. Adaptive Server Anywhere character set support is carried out in a different manner to character set support in Adaptive Server Enterprise and other Open Client/Open Server based applications.

This section describes how to configure the Adaptive Server Anywhere LTM such that data in an Adaptive Server Anywhere database can be shared with Replication Server and hence with other databases.

The LTM automatically uses the default Open Client/Open Server language, sort order, and character set. You can override these defaults by adding entries to the LTM configuration file.

Open Client/Open Server collations

Adaptive Server Enterprise, Replication Server, and other Open Client/Open Server applications share a common means of managing character sets.

🔗 For information on Open Client/Open Server character set support, see the chapter "Configuring Character Sets, Sort Orders, and Languages" in the *Adaptive Server Enterprise System Administration Guide*. For more information about character set issues in Replication Server, see the chapter "International Replication Design Considerations" in the *Replication Server Design Guide*.

This section provides a brief overview of Open Client/Open Server character set support.

Internationalization files

Files that support data processing in a particular language are called **internationalization files**. Several types of internationalization files come with Adaptive Server Enterprise and other Open Client/Open Server applications.

There is a directory named charsets under your Sybase directory. Charsets has a set of subdirectories, including one for each character set available to you. Each character set contains a set of files, as described in the following table

File	Description
Charset.loc	Character set definition files that define the lexical properties of each character such as alphanumeric, punctuation, operand, upper or lower case.
*.srt	Defines the sort order for alphanumeric and special characters.
*.xlt	Terminal-specific character translation files for use with utilities.

Character set settings in the LTM configuration file

Three settings in the LTM configuration file refer to character set issues:

- ◆ **LTM_charset** The character set for the LTM to use. You can specify any Sybase-supported character set.
- ◆ **LTM_language** The **language** used by the LTM to print its messages to the error log and to its clients. You can choose any language to which the LTM has been localized, as long as it is compatible with the LTM character set.

The Adaptive Server Anywhere LTM has been localized to several languages.

Notes

Character set In an Open Client/Open Server environment, an LTM should use the same character set as the data server and Replication Server attached to it.

Adaptive Server Anywhere character sets are specified differently than Open Client/Open Server character sets. Consequently, the requirement is that the Adaptive Server Anywhere character set be compatible with the LTM character set.

Language The *locales.dat* file in the *locales* subdirectory of the Sybase release directory contains valid map settings. However, the LTM output messages in the user interface are currently available in those languages to which the LTM has been localized.

Sort order All sort orders in your replication system should be the same. You can find the default entry for your platform in the *locales.dat* file in the *locales* subdirectory of the Sybase release directory.

Example

- ◆ The following settings are valid for a Japanese installation:

```
LTM_charset=SJIS  
LTM_language=Japanese
```

Transaction log and backup management

One of the differences between the Adaptive Server Enterprise LTM and the Adaptive Server Anywhere LTM is that while the Adaptive Server Enterprise LTM depends on a temporary recovery database for access to old transactions, the Adaptive Server Anywhere LTM depends on access to old transaction logs. No temporary recovery database exists for the Adaptive Server Anywhere LTM.

Replication depends on access to operations in the transaction log, and for Adaptive Server Anywhere primary site databases, sometimes access to old transaction logs. This section describes how to set up backup procedures at an Adaptive Server Anywhere primary site to ensure proper access to old transaction logs.

Consequences of
lost transaction
logs

Good backup practices at Adaptive Server Anywhere primary database sites are crucial. A lost transaction log could mean rematerializing replicate site databases. At primary database sites, a transaction log mirror is recommended. For information on transaction log mirrors and other backup procedure information, see the Adaptive Server Anywhere *User's Guide*.

The LTM configuration file contains a directory entry, which points to the directory where backed up transaction logs are kept. This section describes how to set up a backup procedure to ensure that such a directory stays in proper shape.

Backup utility
options

With the Backup utility, you have the option of renaming the transaction log on backup and restart. For the *dbbackup* utility, this is the *-r* option. It is recommended that you use this option when backing up the consolidated database and remote database transaction logs.

For example, consider a database named *primedb.db*, in directory *c:\prime*, with a transaction log in directory *d:\primelog\primedb.log*. Backing up this transaction log to a directory *e:\primebak* using the rename and restart option carries out the following tasks:

- 1 Backs up the transaction log, creating a backup file
e:\primebak\primedb.log.

- 2 Renames the existing transaction log to *d:\primelog\YYMMDDxx.log*, where **xx** are sequential characters ranging from **AA** to **ZZ**.
- 3 Starts a new transaction log, as *d:\primelog\primedb.log*.


After several backups, the directory *d:\primelog* will contain a set of sequential transaction logs. The log directory should not contain any transaction logs other than the sequence of logs generated by this backup procedure.

Using the DELETE_OLD_LOGS option

The DELETE_OLD_LOGS Adaptive Server Anywhere database option's default is OFF. If you set the default to ON, then the LTM automatically deletes the old transaction logs when Replication Server no longer needs access to the transactions. This option can help to manage disk space in replication setups.

You set the DELETE_OLD_LOGS option for the PUBLIC group:

```
SET OPTION PUBLIC.DELETE_OLD_LOGS = 'ON'
```

 For more information, see "DELETE_OLD_LOGS option" on page 566.

The Unload utility and replication

If a database participates in replication, care must be taken when unloading and reloading in order to avoid needing to re-materialize the database. Replication is based on the transaction log, and unloading and reloading a database deletes the old transaction log. For this reason, good backup practices are especially important when participating in replication.

Replicating an entire database

Adaptive Server Anywhere provides a short cut for replicating an entire database, so you don't have to set each table in the database as a replicated table.

You can set a PUBLIC database option called REPLICATE_ALL using the SET OPTION statement. You can designate a whole database for replication using the following command:

```
SET OPTION PUBLIC.Replicate_all='ON'
```

You require DBA authority to change this and other PUBLIC option settings. You must restart the database for the new setting to take effect. The REPLICATE_ALL option has no effect on procedures.

 For more information, see "REPLICATE_ALL option" on page 595.

Stopping the LTM

You can shut down the LTM from the user interface in Windows NT/2000/XP, or in other circumstances by issuing a command.

❖ **To stop the LTM in Windows NT/2000/XP, when the LTM is not running as a service:**

- ◆ Click SHUTDOWN on the user interface.

❖ **To stop the LTM by issuing a command:**

- 1 Connect to the LTM from *isql* using the LTM_admin_user login name and password in the LTM configuration file. The user ID and password are case sensitive.
- 2 Stop the LTM using the SHUTDOWN statement.

Example

- ◆ The following statements connect *isql* to the LTM PRIMELTM, and shut it down:

```
isql -SPRIMELTM -UDBA -PSQL
1> shutdown
2> go
```


PART FOUR

Utilities, Options, and Properties

This section describes database administration utilities, options, properties, and limitations.

CHAPTER 15

Database Administration Utilities

About this chapter

Adaptive Server Anywhere includes a set of utility programs for backing up databases and performing other database administration tasks. This chapter provides reference information for each of the database administration utilities.


Contents

Topic	Page
Administration utilities overview	437
The Backup utility	438
The Collation utility	442
The File Hiding utility	446
The Compression utility	448
The Console utility	450
The Data Source utility	451
The Debugger utility	456
The Erase utility	458
The Histogram utility	461
The Information utility	463
The Initialization utility	465
The Interactive SQL utility	472
The Language utility	477
The License utility	479
The Log Transfer Manager	481
The Log Translation utility	488
The Ping utility	494
The Rebuild utility	497
The Server Location utility	498
The Service Creation utility	499

The Spawn utility	503
The Stop utility	505
The Transaction Log utility	507
The Uncompression utility	511
The Unload utility	513
The Upgrade utility	521
The Validation utility	526
The Write File utility	530

Administration utilities overview

This chapter presents reference information on the programs and database administration utilities that are part of Adaptive Server Anywhere. Each of the utilities can be accessed from one or more of Sybase Central, Interactive SQL, or at the command prompt.

 For comprehensive documentation on Sybase Central, see the Sybase Central online Help. For an introduction to the Sybase Central database administration tool, see "Tutorial: Managing Databases with Sybase Central" on page 49 of the book *Introducing SQL Anywhere Studio*.

The administration utilities use a set of system environment variables. These variables are described in "Environment variables" on page 208.

Database file administration statements

A set of SQL statements are available that carry out some of the tasks that the administration utilities carry out. These statements are listed in "SQL Statements" on page 199 of the book *ASA SQL Reference Manual*.

The Backup utility

You can use the Backup utility to back up running databases, database files, transaction logs, and write files.

You can access the Backup utility in the following ways:

- ◆ From Sybase Central, using the Create Backup Images wizard.
- ◆ At the command prompt, using the *dbbackup* command. This is useful for incorporating backup procedures into batch or command files.

The Backup utility makes a backup copy of all the files for a single database. A simple database consists of two files: the main database file and the transaction log. More complicated databases can store tables in multiple files, with each file as a separate dbspace. All backup filenames are the same as the database filenames.

Running the Backup utility on a running database is equivalent to copying the database files when the database is not running. Thus, you can back up the database while other applications or users are using it.

Exit codes are 0 (success) or non-zero (failure).

🔗 For a description of suggested backup procedures, see "Backup and Data Recovery" on page 299.

🔗 For more information about the backup utility, see "BACKUP statement" on page 245 of the book *ASA SQL Reference Manual*.

Backing up a database using the Create Backup Images wizard

❖ To back up a database file or a running database:

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Create Backup Images in the right pane.

The Create Backup Images wizard appears.

- 3 Follow the instructions in the wizard.

🔗 For full information on backing up a database from Sybase Central, see "Backup and Data Recovery" on page 299. For more information about options, see "Backup utility options" on page 440

Tip

You can also access the Create Backup Images Database wizard from within Sybase Central using any of the following methods:

- ◆ choosing Tools ► Adaptive Server Anywhere 8 ► Create Backup Images.
- ◆ Selecting a database in the left pane, and choosing Create Backup Images from the File menu.
- ◆ Right-clicking a database, and choosing Create Backup Images from the popup menu.

Backing up a database using the dbbackup command-line utility

Syntax

dbbackup [*options*] *directory*

Option	Description
-c " <i>keyword=value; ...</i> "	Supply database connection parameters
-d	Only back up the main database file
-l <i>file</i>	Live backup of the transaction log to a file
-n	Change the naming convention for the backup transaction log
-o <i>filename</i>	Log output messages to a file
-q	Quiet mode—do not print messages
-r	Rename and start a new transaction log
-t	Only back up the transaction log
-w	Only back up the write file
-x	Delete and restart the transaction log
-xo	Delete and restart the transaction log without making a backup
-y	Replace files without confirmation

Description

If none of the options **-d**, **-t**, or **-w** are used, all database files are backed up.

☞ For more information about the Backup utility options, see "Backup utility options" on page 440.

Backup utility options

Directory The directory to which the backup files are copied. If the directory does not exist, it is created. However, the parent directory must exist.

Connection parameters (-c) For a description of the connection parameters, see "Connection parameters" on page 70. If the connection parameters are not specified, connection parameters from the SQLCONNECT environment variable are used, if set. The **user ID** must have DBA authority or REMOTE DBA authority.

For example, the following command backs up the **asademo** database running on the server **sample_server**, connecting as user ID **DBA** with password **SQL**, into the *asabackup* directory:

```
dbbackup -c "eng=sample_server;dbn=asademo;uid=DBA;pwd=SQL" asabackup
```

Backup main database only (-d) Back up the main database files only, without backing up the transaction log file or a write file, if one exists.

Live backup (-l lower-case L) This option is provided to enable a secondary system to be brought up rapidly in the event of a server crash. A live backup does not terminate, but continues running while the server runs. It runs until the primary server crashes. At that point, it is shut down, but the backed up log file is intact and can be used to bring a secondary system up quickly.

Change backup transaction log naming convention (-n) This option is used in conjunction with **-r**. It changes the naming convention of the backup transaction log file to *yymmddxx.log*, where *xx* are sequential letters ranging from AA to ZZ and *yymmdd* represents the current year, month and day.

The backup copy of the transaction log file is stored in the directory specified in the command, and with the *yymmddxx.log* naming convention. This allows backups of multiple versions of the transaction log file to be kept in the same backup directory.

The two-digit year notation does not cause any year 2000 problems. The names are used solely for identification, not for ordering.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Rename and start new transaction log (-r) This option forces a checkpoint and the following three steps to occur:

- 1 The current working transaction log file is copied and saved to the directory specified in the command.
- 2 The current transaction log remains in its current directory, but is renamed using the format *yymmddxx.log*, where *xx* are sequential characters starting at AA and running through to ZZ, and *yymmdd* represents the current year, month and day. This file is then no longer the current transaction log.
- 3 A new transaction log file is generated that contains no transactions. It is given the name of the file that was previously considered the current transaction log, and is used by the database server as the current transaction log.

Back up the transaction log file only (-t) This can be used as an incremental backup since the transaction log can be applied to the most recently backed up copy of the database file(s).

Back up the database write file only (-w) For a description of database write files, see "The Write File utility" on page 530.

Delete and restart the transaction log (-x) Back up the existing transaction log, then delete the original log and start a new transaction log. This option causes the backup to wait for a point when all transactions from all connections are committed.

Delete and restart the transaction log without a backup (-xo) Delete the current transaction log and start a new one. This operation does not carry out a backup; its purpose is to free up disk space in non-replication environments.

Operate without confirming actions (-y) Choosing this option creates the backup directory or the replacement of a previous backup file in the directory without confirmation.

The Collation utility


With the Collation utility, you can extract a collation (sorting sequence) into a file suitable for creating a database using a custom collation.

The file that is produced can be modified and used with Sybase Central or the `-z` option of *dbinit* to create a new database with a custom collation.

Exit codes are 0 (success) or non-zero (failure).

You must change the label on the following line in the collation file. If you do not, the custom collation will conflict with the original collation on which it is based.

```
Collation label (name)
```

 For more information about custom collating sequences, see "Creating a database with a custom collation" on page 293.

You can access the Collation utility in the following ways:

- ◆ From Sybase Central, using the Create Custom Collation wizard.
- ◆ At the command prompt, using the *dbcollat* command.

Extracting a collation using the Create Custom Collation wizard

❖ To use the Create Custom Collation wizard:

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Create Custom Collation in the right pane.
The Create Custom Collation wizard appears.
- 3 Follow the instructions in the wizard.

Tip

You can also access the Create Custom Collation wizard from within Sybase Central using any of the following methods:

- ◆ choosing Tools ► Adaptive Server Anywhere 8 ► Create Custom Collation.
- ◆ Selecting a database in the left pane, and choosing Create Custom Collation from the File menu.
- ◆ Right-clicking a database, and choosing Create Custom Collation from the popup menu.

Extracting a collation using the dbcollat command-line utility

Syntax

dbcollat [*options*] *output-file*

Options	Description
-c <i>"keyword=value; ..."</i>	Supply database connection parameters
-d <i>collation-file</i>	Convert the definition file to INSERT statements with collation mapping placed in <i>output-file</i>
-e	Include empty mappings
-o <i>filename</i>	Log output messages to a file
-q	Quiet mode—do not print messages
-x	Use hex for extended characters (7F-FF)
-y	Replace the file without confirmation
-z <i>col-seq</i>	Specify a collating sequence label

For more information about the Collation utility options, see "Collation utility options" on page 443.

Collation utility options

Connection parameters (-c) For a description of the connection parameters, see "Connection parameters" on page 70. If the connection parameters are not specified, connection parameters from the SQLCONNECT environment variable are used, if set.

For example, the following command extracts a collation file from the **asademo** database that is running on the **sample_server** server, and connects as user ID **DBA** with password **SQL**:

```
dbcollat -c "eng=sample_server;dbn=asademo;uid=DBA;pwd=SQL" c:\sample\col
```

Convert the collation definition file to INSERT statements that describe the collation (-d)

When a database is created, the collation is inserted into the SYS.SYSCOLLATION system table. A mapping from the collation to character sets and Sybase TDS collations is also inserted into the SYS.SYSCOLLATIONMAPPINGS system table. This collation is selected from the set of provided collations in the *collseqs.sql* file or from the custom collations in the *custom.sql* file in the *scripts* subdirectory of your Adaptive Server Anywhere installation directory.

For more information about the SYSCOLLATIONMAPPINGS system table, see "SYSCOLLATIONMAPPINGS system table" on page 606 of the book *ASA SQL Reference Manual*.

Custom collations are added to the *custom.sql* script. The -d option converts the collation definition file that you provide into INSERT statements that can be copied into *custom.sql*.

For example, you can use the -d option with the *dbcollat* command as follows:

```
dbcollat -d coll-defn-file custom-file
```

The *coll-defn-file* is read and parsed as a collation definition. Output is written to *custom-file*. The *custom-file* contents must be added to *custom.sql*.

For more information about creating a custom collation using the -d option, see "Creating a custom collation" on page 292.

Include empty mappings (-e) Normally, collations don't specify the actual value that a character is to sort to. Instead, each line of the collation sorts to one position higher than the previous line. However, older collations have gaps between some sort positions. Normally, the Collation utility skips the gaps and writes the next line with an explicit sort-position. This option causes the Collation utility to write empty mappings (consisting of just a colon (:)) for each line in the gap.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Use hexadecimal for extended characters [7F to FF] (-x) Extended single-byte characters (whose value is greater than hex 7F) may or may not appear correctly on your screen, depending on whether the code page on your computer is the same as the code page of the collation that you are extracting. This option causes the Collation utility to write any character to hex 7F or above as a two-digit hexadecimal number, in the form `\xdd`. For example:

```
\x80, \xFE
```

Without the `-x` option, only characters from `\x00` to `\x1F`, `\x7F` and `\xFF` are written in hexadecimal form.

Operate without confirming actions (-y) Choosing this option automatically replaces an existing collation file without prompting for confirmation.

Specify a collating sequence label (-z) Specify the label of the collation to be extracted. The names of the recommended collation sequences can be found by executing the following command:

```
dbinit -l
```

If the `-z` option is specified with one of the available collation labels, then *dbcollat* does not connect to a database. Otherwise, it connects to a database and extracts the collation of that database. If the collation label does not match the collation label of the database, an error is returned.

The File Hiding utility

With the File Hiding [dbfhide] utility you can add simple encryption to configuration files to hide the contents of the file.

You can access the File Hiding utility:

- ◆ At the command prompt, using the *dbfhide* command. This is useful for incorporating into batch or configuration files.

Hiding the contents of files using the dbfhide command-line utility

Syntax

dbfhide *original-configuration-file encrypted-configuration-file*

Description

Configuration files (also called command files) are used by some utilities to contain command-line options. These options may contain a password. You can use the File Hiding utility to add simple encryption to configuration files, and thereby obfuscate the contents of the file. The original file will not be modified.

The utilities that use configuration files include:

- ◆ dbsrv8
- ◆ dbmlsrv8
- ◆ dbmlsync
- ◆ dbltm
- ◆ dbremote
- ◆ ssremote
- ◆ ssqueue

Example

Create a configuration file that starts the personal database server and the sample database. It should set a cache of 10 Mb, and name this instance of the personal server *Elora*. The configuration file would be written as follows:

```
# Configuration file for server Elora
-n Elora
-c 10M
path\asademo.db
```

(Note that lines beginning with # are treated as comments.)

Name the file *sample.txt*. If you wanted to start the database using this configuration file, your command line would be:


```
dbeng8 @sample.txt
```


Now, add simple encryption to the configuration.

```
dbfhide sample.txt encrypted_sample.txt
```

Use the encrypted_sample.txt file to start a database.

```
dbsrv8 @encrypted_sample.txt
```

 For more information about Encryption, see "Keeping Your Data Secure" on page 387. For more information about using configuration files, see "Using configuration files" on page 10.

The Compression utility

With the Compression utility you can compress a database file. The Compression utility reads the given database file and creates a compressed database file. Compressed databases are usually 40 to 60 per cent of their original size. The database server cannot update compressed database files: they must be used as read-only files, in conjunction with write files.

The Compression utility does not compress files other than the main database file. The database to be compressed must not be running.

Exit codes are 0 (success) or non-zero (failure).

You can access the Compression utility in the following ways:

- ◆ From Sybase Central, using the Compress Database wizard.
- ◆ At the command prompt, using the *dbshrink* command. This is useful for incorporating into batch or command files.

Compressing a database using the Compress Database wizard

❖ To compress a database file:

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Compress Database in the right pane. The Compress Database wizard appears.
- 3 Follow the instructions in the wizard.

Tip

You can also access the Backup wizard from within Sybase Central by right-clicking a database, and choosing Compress Database from the popup menu.

Compressing a database using the dbshrink command-line utility

Syntax

dbshrink [*options*] *database-file* [*compressed-database-file*]

Option	Description
-ek <i>key</i>	Specify encryption key
-ep	Prompt for encryption key
-o <i>filename</i>	Log output messages to a file
-q	Quiet mode—do not print messages
-y	Replace an existing output file without confirmation

Description

The *dbshrink* utility reads *database-file* and creates a compressed database file. The compressed filename defaults to the same name as the original database file, but with an extension of *.cdb*. The output filename (with extension) must not be same as the input filename (with extension).

For more information about the Compression utility options, see "Compression utility options" on page 449.

Compression utility options

Specify encryption key (-ek) This option allows you to specify the encryption key for strongly encrypted databases directly in the command. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way. For strongly encrypted databases, you must specify either **-ek** or **-ep**, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Prompt for encryption key (-ep) This option allows you to specify that you want to be prompted for the encryption key. This option causes a dialog box to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text. For strongly encrypted databases, you must specify either **-ek** or **-ep**, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Operate without confirming actions (-y) Choosing this option automatically replaces existing compressed database file(s) without requesting confirmation.

The Console utility

The Console utility provides administration and monitoring facilities for database server connections. The Console utility is available on all operating systems.

Monitoring connections using the dbconsole command-line utility

Syntax	dbconsole [options]	
	Option	Description
	-c "keyword=value; ..."	Database connection parameters

Description

Start the server console utility.

The server console utility allows you to monitor the server from a client machine. You can use it to track who is logged on to a database server elsewhere on your network. You can also display both the server and the client statistics on your local client screen, disconnect users, and configure the database server. The Console utility can display information for multiple connections.

Exit codes are 0 (success) or non-zero (failure).

- ❖ **To disconnect a user from a database:**
- 1 Connect to the database from the Console utility.
 - 2 Double-click the connection to display the Connection Detail window.
 - 3 Click Disconnect to disconnect the specified connection from the database.

Console utility options

-c "keyword=value; ..." Specify connection parameters.

🔗 For a description of connection parameters, see "Connection parameters" on page 70.

The Data Source utility

The Data Source utility is a cross-platform alternative to the ODBC Administrator for creating, deleting, describing, and listing Adaptive Server Anywhere ODBC data sources.

Managing ODBC data sources using the dbdsn command-line utility

Syntax

```
dbdsn [ modifier-options ]
      { -l[ u | s ] [ qq ]
        | -d[ u | s ] dsn
        | -g[ u | s ] dsn
        | -w[ u | s ] dsn [details-options;...]
        | -cl[ qq ] }
```

Parameters

Major option	Description
-l[u s] [qq]	List either all Adaptive Server Anywhere users or all Adaptive Server Anywhere system data sources. You can also list both Adaptive Server Anywhere user and Adaptive Server Anywhere system data sources. User data sources is the default. Using -qq with this option lists the DSNs without any banner or titles.
-d[u s] <i>dsn</i>	Delete the named Adaptive Server Anywhere user or Adaptive Server Anywhere system data source. User data sources is the default.
-g[u s] <i>dsn</i>	List (get) details about the named Adaptive Server Anywhere user or Adaptive Server Anywhere system data source. User data sources is the default.
-w[u s] <i>dsn</i> [<i>details-options</i>]	Create (write) a user or system data source definition. User data sources is the default.
-cl[qq]	List available connection parameters. Using -qq with this option lists the available connection parameters without any banner or titles.

Modifier option	Description
-b	Brief. Print connection string for the data source
-q	Quiet. Do not print banner
-v	Verbose. Print connection parameters in tabular form
-y	Delete or overwrite data source without confirmation
Details-option	Description
-c "keyword=value;..."	Supply database connection parameters
-ec	Encrypt all network packets
-o filename	Write client message to filename
-p size	Set maximum network packet size
-r	Disable multiple record fetching
-tl seconds	Client liveness timeout period
-x list	List network drivers to run
-z	Display debugging information
server-name	Connect to named database server

See also "Working with ODBC data sources" on page 53
"Using ODBC data sources on UNIX" on page 63

Description The Data Source utility is a cross-platform alternative to the ODBC Administrator for creating, deleting, describing, and listing Adaptive Server Anywhere ODBC data sources. On Windows operating systems, the data sources are held in the registry. On UNIX operating systems, the data sources are held in the *.odbc.ini* file. The utility is useful for batch operations.

The modifier options can occur before or after the major option specification. The order makes a difference only when a connection parameter value is specified more than once. In such a case, the last value specified is used.

Exit codes are 0 (success) or non-zero (failure).

Examples Write a definition of the data source **newdsn**. Do not prompt for confirmation if the data source already exists.

```
dbdsn -y -x tcPIP -w newdsn -c "uid=DBA;pwd=SQL" -v
```

or, with a different option order,

```
dbdsn -w newdsn -c "uid=DBA;pwd=SQL" -x tcPIP -y
```

List all known user data sources, one data source name per line:

```
dbdsn -l
```

List all known system data sources, one data source name per line:

```
dbdsn -ls
```

List all data sources along with their associated connection string:

```
dbdsn -l -b
```

Report the connection string for user data source **MyDSN**:

```
dbdsn -g MyDSN
```

Report the connection string for system data source **MyDSN**:

```
dbdsn -gs MyDSN
```

Delete the data source **BadDSN**, but first list the connection parameters for **BadDSN** and prompt for confirmation:

```
dbdsn -d BadDSN -v
```

Delete the data source **BadDSN** without prompting for confirmation.

```
dbdsn -d BadDSN -y
```

Create a data source named **NewDSN** for the database server **MyServer**:

```
dbdsn -w NewDSN -c "uid=DBA;pwd=SQL;eng=bar"
```

If a **NewDSN** already exists, the previous definition is overwritten.

The following example connects to the **sample** database server. The server name **sample** overrides the previous specified value of **MyServer**:

```
dbdsn -w NewDSN -c "uid=DBA;pwd=SQL;eng=MyServer" sample
```

List all connection parameter names and their aliases:

```
dbdsn -cl
```

List all user DSNs (without banner and titles):

```
dbdsn -l -qq
```

List all connection parameter names (without banner and titles):

```
dbdsn -cl -qq
```

Data Source utility options

When you use the Data Source utility, you can choose from major options, modifier options, and details options.

Major options

List defined data sources (-l) Lists the available Adaptive Server Anywhere ODBC data sources. You can modify the list format using the `-b` or `-v` options. You can modify the option using the `u` (user) or `s` (system) specifiers. The default specifier is `u`.

Delete the named data source (-d) Deletes the named Adaptive Server Anywhere data source. You can modify the option using the `u` (user) or `s` (system) specifiers. `u` is the default specifier. If you supply `-y`, any existing data source is overwritten without confirmation.

List (get) details of the named data source (-g) List the definition of the named Adaptive Server Anywhere data source. You can modify the format of the output using the `-b` or `-v` options. You can modify the option using the `u` (user) or `s` (system) specifiers. `u` is the default specifier.

Create (write) a data source definition (-w) Creates a new data source, or overwrites one if one of the same name exists. You can modify the option using the `u` (user) or `s` (system) specifiers. `u` is the default specifier. If you supply `-y`, any existing data source is overwritten without confirmation.

List available connection parameters (-cl) This convenience option lists the connection parameters supported by the *dbdsn* utility.

Modifier options

Print connection string for the data source (-b) Format the output of the list as a single line connection string.

Do not print banner (-q) Suppress the informational banner.

Do not print banner or titles (-qq) Suppress both the informational banner and titles. This option can only be used with the `-l` and the `-cl` options.

Print connection parameters in tabular form (-v) Format the output of the list over several lines, as a table.

Delete or overwrite data source without confirmation (-y) Choosing this option automatically deletes or overwrites each data source without prompting you for confirmation.

Details options

Connection parameters (-c) Specify connection parameters as a connection string.

 For more information, see "Connection parameters" on page 70.

The remainder of the details options are supplied for users familiar with the *dbcli6.exe* utility shipped with Adaptive Server Anywhere version 6.

Encrypt network packets (-ec) Encrypt packets sent between the client application and the server.

☞ For more information, see "Encryption connection parameter" on page 177.

Log output messages to file (-o) Write output messages to the named file. By default, messages are written to the console.

☞ For more information, see "Logfile connection parameter" on page 183.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Set maximum network packet size (-p) The maximum packet size for network communications, in bytes. The value must be greater than 300, and less than 16000. The default setting is 1460.

☞ For more information, see "CommBufferSize connection parameter" on page 168.

Disable multiple-record fetching (-r) By default, when the database server gets a simple fetch request, the application asks for extra rows. You can disable this behavior using this option.

☞ For more information, see "DisableMultiRowFetch connection parameter" on page 175.

Set client liveness timeout (-tl) Terminates connections when they are no longer intact. The value is in seconds.

The default is server setting, which in turn has a default value of 120 seconds.

☞ For more information, see "LivenessTimeout connection parameter" on page 182.

Set communications links (-x) A comma separated list of network drivers to run.

☞ For more information, see "CommLinks connection parameter" on page 169.

Display debugging information (-z) Provide diagnostic information on communications links on startup.

Server name Connect to the named server. Only the first 40 characters are used.

☞ For more information, see "The Database Server" on page 119.

The Debugger utility

With the Adaptive Server Anywhere Object Debugger, you can debug SQL stored procedures (including triggers and event handlers) and Java classes executing in an Adaptive Server Anywhere database.

Exit codes are 0 (success) or non-zero (failure).

You can access the Debugger utility in the following ways:

- ◆ From Sybase Central, using the Database Object Debugger.
- ◆ At the command prompt, using the *dbprdbg* command.

Opening the Database Object Debugger in Sybase Central

❖ **To open the Database Debugger in Sybase Central:**

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Open Database Debugger in the right pane.

The Database Debugger appears.

Tip

You can also access the Database Debugger utility by choosing Start►Programs►SQL Anywhere 8►Adaptive Server Anywhere►Database Object Debugger, or from within Sybase Central using any of the following methods:

- ◆ choosing Tools►Adaptive Server Anywhere 8►Open Database Object Debugger.
- ◆ Selecting a database in the left pane, and choosing Open Database Object Debugger from the File menu.
- ◆ Right-clicking a database, and choosing Open Database Object from the popup menu.

Debugging database objects using the dbprdbg command-line utility


Syntax

dbprdbg [*options*]

Option	Description
-c "keyword=value;..."	Database connection parameters
-user username	User to debug

Description

With the Adaptive Server Anywhere Debugger, you can debug SQL stored procedures (including triggers and event handlers) and Java classes that are executing in an Adaptive Server Anywhere database.

 For information on connection parameters, see "Connection parameters" on page 70.

The debugger takes an Adaptive Server Anywhere connection string as its command-line argument. When started at the command prompt, the debugger is configured to debug all connections to the named database.

The debugger shares a connection dialog and connection logic with Interactive SQL and the Sybase Central Adaptive Server Anywhere plug-in.

See also

"Debugging Logic in the Database" on page 571 of the book *ASA SQL User's Guide*

Debugger utility options

-c Specify connection parameters. See "Connection parameters" on page 70 for a description of the connection parameters. If the debugger cannot connect, you are presented with a dialog where you can enter the connection parameters. Connection parameters are semi-colon delimited.

-user If specified, debug connections for the given user only. If not specified, debug all connections. The default is to debug all connections.

The Erase utility

With the Erase utility, you can erase a database file or write file and its associated transaction log, or you can erase a transaction log file or transaction log mirror file. All database files, write files, and transaction log files are marked read-only to prevent accidental damage to the database and accidental deletion of the database files.

Deleting a database file that references other dbspaces does not automatically delete the dbspace files. If you want to delete the dbspace files on your own, change the files from read-only to writable, and then delete the files individually. As an alternative, you can use the DROP DATABASE statement to erase a database and its associated dbspace files.

If you erase a database file or write file, the associated transaction log and transaction log mirror are also deleted. If you erase a transaction log for a database that also maintains a transaction log mirror, the mirror is not deleted.

The database to be erased must not be running when this utility is used.

Exit codes are 0 (success) or non-zero (failure).

You can access the Erase utility in the following ways:

- ◆ From Sybase Central, using the Erase Database wizard.
- ◆ At the command prompt, using the *dberase* command. This is useful for incorporating into batch or command files.

✎ For more information, see the "DROP DATABASE statement" on page 399 of the book *ASA SQL Reference Manual*.

Erasing a database using the Erase Database wizard

❖ To erase a database file:

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Erase Database in the right pane. The Erase Database wizard appears.
- 3 Follow the instructions in the wizard.

✎ For full information on erasing a database from Sybase Central, see "Erasing a database" on page 32 of the book *ASA SQL User's Guide*.

Tip

You can also access the Erase Database wizard from within Sybase Central using any of the following methods:

- ◆ choosing Tools►Adaptive Server Anywhere 8►Erase Database.
- ◆ Selecting a server in the left pane, and choosing Erase Database from the File menu.
- ◆ Right-clicking a server, and choosing Erase Database from the popup menu.

Erasing a database using the dberase command-line utility

Syntax

dberase [*options*] *database-file*

Option	Description
-ek <i>key</i>	Specify encryption key
-ep	Prompt for encryption key
-o <i>filename</i>	Log output messages to a file
-q	Operate quietly—do not print messages
-y	Erase files without confirmation

Description

The *database-file* may be a database file, write file, or transaction log file. The full filename must be specified, including extension. If a database file or write file is specified, the associated transaction log file (and mirror, if one is maintained) is also erased.

🔗 For more information about the Erase utility options, see "Erase utility options" on page 459.

Erase utility options

Specify encryption key (-ek) This option allows you to specify the encryption key for strongly encrypted databases directly in the command. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way. For strongly encrypted databases, you must specify either **-ek** or **-ep**, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Prompt for encryption key (-ep) This option allows you to specify that you want to be prompted for the encryption key. This option causes a dialog box to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text. For strongly encrypted databases, you must specify either `-ek` or `-ep`, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Operate without confirming actions (-y) Choosing this option automatically deletes each file without prompting you for confirmation.

The Histogram utility

The Histogram utility converts a histogram into a Microsoft Excel chart. The utility only works with Excel 97 and higher, and only on Windows.

Managing histograms using the dbhist command-line utility

Syntax `dbhist [options] -t table-name [excel-output-name]`

Parameters

Optional options	Description
<code>-c options</code>	Connection string. If not specified, the defaults are uid=dba, pwd=sql
<code>-n colname</code>	Column to associate with the histogram
<code>-u owner</code>	The table owner
<code>-t table-name</code>	The name of the table to generate histograms for
<code>excel-output-name:</code>	The name of the generated Excel file. If no name is specified, Excel prompts you to enter one with a Save As dialog.

Description

Histograms are stored in the **SYSCOLSTAT** system table and can also be retrieved with the *sa_get_histogram* stored procedure.

To determine the selectivity of a predicate over a string column, you should use the *ESTIMATE* or *ESTIMATE_SOURCE* functions. For string columns, both *sa_get_histogram* and the Histogram utility retrieve nothing from **SYSCOLSTAT**.

Exit codes are 0 (success) or non-zero (failure).

For more information about the *sa_get_histogram* stored procedure, see "sa_get_histogram system procedure" on page 694 of the book *ASA SQL Reference Manual*.

Example

Assuming that a histogram has been created for the column, the following statement will generate an Excel chart for the column *prod_id* in the table *sales_order_items* for database *asademo.db*, and save it as *histgram.xls*.

```
dbhist -c "uid=DBA;pwd=SQL;dbf=asademo.db" -n prod_id -t
sales_order_items histgram.xls
```

The following statement generates charts for every column with a histogram in the table *sales_order*, assuming that *asademo* is already loaded. This statement also attempts to connect using uid=dba, pwd=sql. No output file name is specified, so Excel prompts you to enter one.

```
dbhist -t sales_order
```

Histogram utility options

Specify a connection string (-c) Specify connection parameters.

🔗 For a description of the connection parameters, see "Connection parameters" on page 70.

Specify a column (-n) Specify the name of the column to associate the histogram with. If you do not specify a column, all columns that have histograms in the table are returned.

Specify an owner (-u) Specify the owner of the table.

Specify a table name (-t) Specify the name of the table to generate histograms for.

The Information utility

The Information utility can display information about a database. The utility indicates when the database was created, the name of any transaction log file or log mirror that is maintained, the page size, the version of installed Java classes, and other information. Optionally, it can also provide table usage statistics and details.

Use the *dbinfo* command to access the Information utility.

Obtaining database information using the dbinfo command-line utility

Syntax

dbinfo [*options*]

Option	Description
-c " <i>keyword=value; ...</i> "	Database connection parameters
-o <i>filename</i>	Log output messages to a file
-q	Operate quietly
-u	Output page usage statistics

Description

The *dbinfo* utility displays information about a database. It reports the time when the database was created, the name of any transaction log file or log mirror, the page size, the version of installed Java classes, the collation name and label, and other information. Optionally, it can also provide table usage statistics and details. Note that if your database does not have a SYSCOLLATION table, the collation name is not returned.

Exit codes are 0 (success) or non-zero (failure).

For more information about the options, see "Information utility options" on page 463.

Information utility options

Connection parameters (-c) Specify connection parameters.

For a description of the connection parameters, see "Connection parameters" on page 70.

Any valid user ID can run the Information utility, but to obtain page usage statistics you need DBA authority.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Page usage statistics (-u) Display information about the usage and size of all tables, including system and user-defined tables.

You can only request page usage statistics if no other users are connected to the database.

The Initialization utility

With the Initialization utility, you can initialize (create) a database. A number of database attributes are specified at initialization and cannot be changed later except by unloading, reinitializing, and rebuilding the entire database. These database attributes include:


- ◆ Case sensitivity or insensitivity
- ◆ Storage as an encrypted file
- ◆ Treatment of trailing blanks in comparisons
- ◆ Page size
- ◆ Collation sequence

In addition, the choice of whether to use a transaction log and a transaction log mirror is made at initialization. This choice can be changed later using the Transaction Log utility.

Exit codes are 0 (success) or non-zero (failure).

You can access the Initialization utility in the following ways:

- ◆ From Sybase Central, using the Create Database wizard.
- ◆ At the command prompt, using the *dbinit* command. This is useful for incorporating into batch or command files.

 For more information about creating a database, see "CREATE DATABASE statement" on page 273 of the book *ASA SQL Reference Manual*.

Creating a database using the Create Database wizard


❖ To create a database:

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Create Database in the right pane. The Create Database wizard appears.
- 3 Follow the instructions in the wizard.

Tip

You can also access the Create Database wizard from within Sybase Central using any of the following methods:

- ◆ choosing Tools►Adaptive Server Anywhere 8►Create Database.
- ◆ Selecting a server in the left pane, and choosing Create Database from the File menu.
- ◆ Right-clicking a server, and choosing Create Database from the popup menu.

 For full information on creating a database in Sybase Central, see "Creating a database" on page 29 of the book *ASA SQL User's Guide*.

Creating a database using the dbinit command-line utility

Syntax

dbinit [*options*] *new-database-file*

Option	Description
-b	Blank padding of strings for comparisons and fetching
-c	Case sensitivity for all string comparisons
-e	Encrypt the database using simple encryption
-ea <i>algorithm</i>	Specify which strong encryption algorithm to encrypt your database with: you can choose AES or MDSR
-ek <i>key</i>	Specify encryption key for strong encryption
-ep	Prompt for encryption key for strong encryption
-i	Do not install Sybase jConnect support
-j	Do not install runtime Java classes
-ja	Install default runtime Java classes
-jdk <i>version</i>	Install entries for the named version of the Java Development Kit
-k	Omit Watcom SQL compatibility views SYS.SYSCOLUMNS and SYS.SYSINDEXES
-l	List the recommended collating sequences
-m <i>file-name</i>	Use a transaction log mirror (default is no mirror)
-n	No transaction log

Option	Description
<code>-o filename</code>	Log output messages to a file
<code>-p page-size</code>	Set page size
<code>-q</code>	Quiet mode—do not print messages
<code>-t log-name</code>	Transaction log filename (default is the database name with <i>.log</i> extension)
<code>-z col-seq</code>	Collation sequence used for comparisons

Description

For example, the database *test.db* can be created with 1024 byte pages as follows:

```
dbinit -p 1024 test.db
```

For more information about the Initialization utility options, see "Initialization utility options" on page 467.

Initialization utility options

Blank padding (–b) Ignore trailing blanks for comparison purposes and pad strings that are fetched into character arrays. For example, the two strings

```
'Smith'
'Smith '
```

would be treated as equal in a database created with blank-padding.

This option is provided for compatibility with the ISO/ANSI SQL standard, which is to ignore trailing blanks in comparisons. The default is that blanks are significant for comparisons.

For related information, see "ANSINULL option" on page 553.

Case sensitivity for all string comparisons (–c) For databases created with this option, all values are considered to be case sensitive in comparisons and string operations. Identifiers in the database are case insensitive, even in case sensitive databases.

This option is provided for compatibility with the ISO/ANSI SQL standard. The default is that all comparisons are case insensitive.

User ID and password

All databases are created with at least one user ID, **DBA**, with password **SQL**. If you create a case-sensitive database, all passwords are case sensitive. User IDs, like other identifiers, are case insensitive even in case sensitive databases.

Encrypt the database using simple encryption (-e) Simple encryption makes it more difficult for someone to decipher the data in your database using a disk utility to look at the file. File compaction utilities cannot compress encrypted database files as much as unencrypted ones.

This simple encryption is equivalent to obfuscation and is intended only to keep data hidden in the event of casual direct access of the database file. For greater security, you can use strong encryption by supplying the **-ea** and **-ek** options.

Specify encryption algorithm (-ea) This option allows you to choose which strong encryption algorithm to encrypt your database with. You can choose either AES or MDSR. Algorithm names are case-insensitive. If you specify the **-ea** option, you must also specify either **-ep** or **-ek**.

Specify encryption key (-ek) This option allows you to create a strongly encrypted database by specifying an encryption key directly in the command. If you specify the **-ek** option without specifying an algorithm (**-ea** option), the algorithm used is AES.

Prompt for encryption key (-ep) This option allows you to specify that you want to create a strongly encrypted database by inputting the encryption key in a dialog box. This provides an extra measure of security by never allowing the encryption key to be seen in clear text. If you specify the **-ep** option without specifying an algorithm (**-ea** option), the algorithm used is AES.

You must input the encryption key twice to confirm that it was entered correctly. If the keys don't match, the initialization fails.

Do not install Sybase jConnect support (-i) If you wish to use the Sybase jConnect JDBC driver to access system catalog information, you need to install jConnect support. Use this option if you wish to exclude the jConnect system objects. You can still use JDBC, as long as you do not access system information. If you want, you can add Sybase jConnect support at a later time using Sybase Central.

For more information, see "Installing jConnect system objects into a database" on page 137 of the book *ASA Programming Guide*.

Do not install runtime Java classes (-j) If you do not intend to use Java classes, you can specify either no Java option, or the `-j` option to avoid including these classes in the database.

You can add Sybase runtime Java classes later using the Upgrade Database wizard, or the ALTER DATABASE statement.

☞ For more information, see "Java-enabling a database" on page 89 of the book *ASA Programming Guide*.

Install runtime Java classes (-ja) If you wish to use Java in your database, you must install the runtime Java classes. Specifying `-ja` installs version 1.3 of the JDK into the database.

For example, the following command creates a database that supports JDK 1.3 applications in the database.

```
dbinit -ja java2.db
```

The runtime classes add several megabytes to the size of a database, so if you do not intend to use Java classes, you can omit the `-ja` option to avoid installing them.

You can add the runtime Java classes at a later time using the Upgrade Database wizard, or the ALTER DATABASE statement.

☞ For more information, see "Java-enabling a database" on page 89 of the book *ASA Programming Guide*.

Install support for the named version of the JDK (-jdk) If you want to install a version of Java other than 1.3 into your database, use the `-jdk` option followed by the version number. Currently, the only other version of Java supported is 1.1.8

For example, the following command creates a database that supports JDK 1.1.8 applications in the database.

```
dbinit -jdk 1.1.8 java2.db
```

The runtime classes add several megabytes to the size of a database, so if you do not intend to use Java classes, you can omit the `-jdk` option to avoid installing them.

The `-jdk` option implies the `-ja` option.

You can add the runtime Java classes at a later time using the Upgrade Database wizard, or the ALTER DATABASE statement.

☞ For more information, see "Java-enabling a database" on page 89 of the book *ASA Programming Guide*.

Omit Watcom SQL compatibility views (-k) By default, database creation generates the views SYS.SYSCOLUMNS and SYS.SYSINDEXES for compatibility with system tables that were available in Watcom SQL (versions 4 and earlier of this software). These views will conflict with the Sybase Adaptive Server Enterprise compatibility views *dbo.syscolumns* and *dbo.sysindexes* unless the -c case sensitivity option is used.

List the available collating sequences (-l) *dbinit* lists the recommended collating sequences and then stops. No database is created. A list of available collating sequences is automatically presented in the Sybase Central Create Database wizard.

Use a transaction log mirror (-m) A transaction log mirror is an identical copy of a transaction log, usually maintained on a separate device, for greater protection of your data. By default, Adaptive Server Anywhere does not use a mirrored transaction log.

Do not use a transaction log (-n) Creating a database without a transaction log saves disk space. The transaction log is required for data replication and provides extra security for database information in case of media failure or system failure. Databases that do not use transaction logs typically run slower than databases that use transaction logs.

Log output messages to file (-o) Write output messages to the named file.

Page size (-p) The page size for a database can be (in bytes) 1024, 2048, 4096, 8192, 16384, or 32768, with 2048 being the default. Any other value for the size will be changed to the next larger size.

Large databases usually benefit from a larger page size. For example, the number of I/O operations required to scan a table is generally lower, as a whole page is read in at a time. Also, the number of levels in an index may be reduced as more entries can be stored on each page.

However, there are additional memory requirements for large page sizes. Also, the maximum number of rows stored on a page is 255, so that tables with small rows will not fill each page. For example, with 8 kb pages the average row size must be at least 32 bytes to fully use each page. For most applications, 16 kb or 32 kb page sizes are not recommended. You should not use page sizes of 16 kb or 32 kb in production systems unless you can be sure that a large database server cache is always available, and only after you have investigated the tradeoffs of memory and disk space with its performance characteristics.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Set the transaction log filename (-t) The transaction log is a file where the database server logs all changes, made by all users, no matter what application system is being used. The transaction log plays a key role in backup and recovery (see "The transaction log" on page 305), and in data replication. If the filename has no path, it is placed in the same directory as the database file. If you run *dbinit* without specifying *-t* or *-n*, a transaction log is created with the same filename as the database file, but with extension *.log*.

Collating sequence (-z) The collation sequence is used for all string comparisons in the database.

If you want to create a custom collation, use the Collation utility to create a file containing the collation. Once you have modified the collation and inserted it into the appropriate scripts, you use the Initialization utility to create the database and specify the new collation.

You must change the collation label in the custom collation file. Otherwise, the Initialization utility prevents the insertion of the new collation, since it conflicts with an existing collation.

🔗 For more information on custom collating sequences, see "International Languages and Character Sets" on page 249. For information on the Collation utility, see "The Collation utility" on page 442.

In order to change the collation that an existing database uses, it is necessary to unload the database, create a new database using the appropriate collation, and then reload the database. It may be necessary to translate the data as well.

If *-z* is not specified, the default collation is used. Normal ASCII (binary) ordering is used for the lower 128 characters, subject to the case sensitivity setting (*-c*). For the upper 128 characters (also called the extended characters), any character that is an accented form of a letter in the lower 128 are sorted to the same position as the unaccented form. The determination of whether or not an extended character is an accented letter is based upon code page 850 (multilingual code page).

🔗 For a list of the available collating sequence labels, see "Understanding collations" on page 269.

The Interactive SQL utility

Interactive SQL provides an interactive environment for database browsing and for sending SQL statements to the database server.

You can access Interactive SQL:

- ◆ from Sybase Central, using the Open Interactive SQL option.
- ◆ at the command prompt, using the *dbisql* command.

Starting Interactive SQL from Sybase Central

❖ **To open Interactive SQL from Sybase Central:**

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Open Interactive SQL in the right pane.

The Interactive SQL window appears.

Tip

You can also access Interactive SQL by choosing Start ► Programs ► SQL Anywhere 8 ► Adaptive Server Anywhere 8 ► Interactive SQL, or from within Sybase Central by choosing any of the following options:

- ◆ choosing Tools ► Adaptive Server Anywhere 8 ► Open Interactive SQL.
- ◆ Selecting a database in the left pane, and choosing Open Interactive SQL from the File menu.
- ◆ Right-clicking a database, and choosing Open Interactive SQL from the popup menu.
- ◆ Right-clicking a stored procedure, and choosing Execute From Interactive SQL from the popup menu. Interactive SQL opens with the procedure in the SQL Statements pane.

Opening Interactive SQL using the dbisql command-line utility

Syntax

dbisql [*options*] [*dbisql-command* | *command-file*]

Option	Description
-c <i>"keyword=value; ..."</i>	Supply database connection parameters
-codepage <i>codepage</i>	Specify a codepage to use when reading or writing files
-d delimiter	Use the given string as the command delimiter
-d1	Print statements as they are executed [command-prompt mode only]
-datasource <i>dsn-name</i>	Specify an ODBC data source to connect to
-host <i>hostname</i>	Specify the hostname or IP address of the machine running a database server
-jConnect	Use jConnect to connect to the database
-nogui	Run in command-prompt mode
-ODBC	Use the JDBC/ODBC bridge to connect to the database
-onerror { continue exit }	Override the ON_ERROR option for all users
-port <i>portnumber</i>	Look on the specified port number for the database server
-q	Quiet mode—no windows or messages
-x	Syntax check only—no commands executed

Description

Interactive SQL allows you to type SQL commands, or run command files. It also provides feedback about the number of rows affected, the time required for each command, the execution plan of queries, and any error messages.

If *dbisql-command* is specified, Interactive SQL executes the command. You can also specify a command file name. If no *dbisql-command* is specified, Interactive SQL enters interactive mode, where you can type a command into a command window.

Exit codes are 0 (success) or non-zero (failure).

Interactive SQL requires that the QUOTED_IDENTIFIER database option be set to ON since a number of database functions, including some statements, rely on this setting to function properly. Interactive SQL automatically sets it ON when connecting to a database.

Examples

The following command, entered at a command prompt, runs the command file *mycom.sql* against the current default server, using the user ID DBA and the password SQL. If there is an error in the command file, the process terminates.

```
dbisql -c "uid=DBA;pwd=SQL" -onerror exit mycom.sql
```

The following command, when entered on a single line at a system prompt, adds a user to the current default database:

```
dbisql -c "uid=DBA;pwd=SQL" grant connect to joe identified by passwd
```


Interactive SQL utility options

-c Specify connection parameters. See "Connection parameters" on page 70 for a description of the connection parameters. If Interactive SQL cannot connect, you are presented with a dialog where you can enter the connection parameters.

-codepage Specify the codepage to use when reading or writing files. The default code page is the default code page for the platform you are running on.

For example, on an English Windows NT machine, Interactive SQL uses the 1252 (ANSI) code page. If you want Interactive SQL to read files created using the 297 (IBM France) code page, specify the following option.

```
-codepage 297
```

 For a list of supported code pages, see "Supported code pages" on page 256.

-d Specify a command delimiter. Quotation marks around the delimiter are optional, but required when the command shell itself interprets the delimiter in some special way.

Command delimiters are used for all connections in that Interactive SQL session, regardless of the setting stored in the database (for the user, or the PUBLIC setting).

-d1 (The final character is a number 1, not a lower-case L). Interactive SQL echoes all statements it executes to the Command window (STDOUT). This can provide useful feedback for debugging SQL scripts, or when Interactive SQL is processing a long SQL script.

-datasource Specify an ODBC data source to connect to. You do not need to be using the JDBC/ODBC bridge to use this option. However, if the data source to which you are connecting is not configured to use TCP/IP, you must use the JDBC/ODBC bridge to connect. Adaptive Server Anywhere data sources are configured to use TCP/IP by default.

-host Specify the hostname or IP address of the computer on which the database server is running. You can use the name **localhost** to represent the current machine.

-jConnect Use the Sybase jConnect JDBC driver to connect to the database. This is the default method, and is recommended in most circumstances.

-nogui Run Interactive SQL in a command-prompt mode, with no windowed user interface. This is useful for batch operations.

In this mode, Interactive SQL sets the program exit code to indicate success or failure. On Windows operating systems, the environment variable `ERRORLEVEL` is set to the program exit code. The exit codes are as follows:

Program exit code	Description
0	Success.
1	General failure. At some point, a SQL or Interactive SQL statement did not execute successfully and the user chose to stop executing SQL statements. Alternatively, Interactive SQL noted an internal error.
5	The user terminated interactive SQL. When an error occurs during execution, the user is prompted to ignore it, stop, or exit Interactive SQL. If the user opts to exit, the program returns code 5. Code 5 is also returned if an error occurs and the Interactive SQL option <code>ON_ERROR</code> is set to <code>EXIT</code> .
9	Unable to connect.
255	Bad command. The command contained incomplete or invalid options.

-ODBC Connect using the JDBC/ODBC bridge. If you do not specify this option, Interactive SQL connects using jConnect.

-onerror Controls what happens if an error is encountered while reading statements from a command file. This option overrides the `ON_ERROR` setting. It is useful when using Interactive SQL in batch operations.

For more information, see "ON_ERROR option" on page 586.

-port Specify the port number on which the database server is running. The default port number for Adaptive Server Anywhere is **2638**.

-q Do not display output messages. This is useful only if you start Interactive SQL with a command or command file.

-x Scan commands but do not execute them. This is useful for checking long command files for syntax errors.

☞ For detailed descriptions of SQL statements and Interactive SQL commands, see "SQL Language Elements" on page 3 of the book *ASA SQL Reference Manual*.

The Language utility

The Language utility reports and changes the registry settings that control the languages used by Adaptive Server Anywhere and Sybase Central.

Managing languages using the `dblang` command-line utility

Syntax `dblang { options } [language-code]`

Option	Description
<code>-d</code>	Change deployment settings only
<code>-q</code>	Operate quietly—do not print messages

Language code	Language
DE	German
EN	English
ES	Spanish
FR	French
IT	Italian
JA	Japanese
KO	Korean
PL	Polish
PT	Portuguese
RU	Russian
TW	Traditional Chinese
UK	Ukrainian
ZH	Simplified Chinese

Description The utility is installed only when the International Resources Deployment Kit (IRDK) is selected.

Running `dblang` without a language code reports the current settings. These settings are as follows:

- ◆ **Adaptive Server Anywhere** A key from `HKEY_LOCAL_MACHINE` that holds a two-letter language code from the table above.

This setting controls which language resource library is used to deliver informational and error messages from the Adaptive Server Anywhere database server. The language resource library is a DLL with a name of the form *dblgXX8.dll*, where XX is a two-letter language code.

Ensure that you have the appropriate language resource library on your machine when you change the settings.

- ◆ **Sybase Central** A key from HKEY_LOCAL_MACHINE that holds a two-letter language code from the table above.

This setting controls the resources used to display user interface elements for Sybase Central and Interactive SQL. You must have purchased the appropriate localized version of SQL Anywhere Studio for this setting to take effect.

Exit codes are 0 (success) or non-zero (failure).

Examples

The following command displays a dialog box containing the current settings:

```
dblang
```



The following command changes the settings to French, and displays a dialog containing the previous and new settings:

```
dblang FR
```

Language utility options

Deployment only (-d) Changes the registry setting only for Adaptive Server Anywhere. Do not change the Sybase Central setting.

Operate quietly (-q) Do not display output messages.

The License utility

The license utility adds licensed users to your network database server.

Managing licenses using the dblic command-line utility

Syntax

dblic { *options* } *executable-name* *user-name* *company-name*

Option	Description
-l <i>type</i>	License type. Allowed values are perseat or concurrent
-o <i>filename</i>	Log output messages to a file
-p <i>operating-system</i>	Target operating system. Allowed values are NetWare , UNIX , WIN32 , WINNT
-q	Operate quietly—do not print messages
-u <i>license-number</i>	Total number of users for license
<i>Executable-name</i>	The executable to be licensed, with path relative to the current directory. Should be either an Adaptive Server Anywhere network database server or a MobiLink synchronization server.

Description

The License utility adds licensed users to your network database server. You must only use this utility in accordance with your Sybase license agreement to license the number of users to which you are entitled. Running this command does *not* grant you license. You can also use this utility to view the current license information for a server executable without starting the server.

Repeated running of the *dblic* command updates the license information.

Exit codes are 0 (success) or non-zero (failure).

Example

The following command, executed in the same directory as the database server executable, applies a license for 50 concurrent users, in the name of *Sys Admin*, for company *My Co*, to a Windows NT network database server. The command must be entered all on one line:

```
dblic -l concurrent -p WINNT -u 50 dbsrv8.exe "Sys
Admin" "My Co"
```

The following message appears on the screen to indicate the success of the license:

Licensed nodes: 50
User: Sys Admin
Company: My Co

License utility options

License type (-l) Enter the license type that matches the licensing model described in your software license agreement. Valid entries are **perseat** and **concurrent**.

Log output messages to file (-o) Write output messages to the named file.

Operating system (-p) Enter the operating system for which you are licensed. Allowed values are as follows:

- ◆ **NetWare** Novell NetWare operating system.
- ◆ **UNIX** A UNIX operating system.
- ◆ **WIN32** Windows operating systems.
- ◆ **WINNT** This has the same meaning as WIN32.

Operate quietly (-q) Do not display output messages.

License number (-u) The total number of licensed seats or concurrent connections. If you are adding extra licenses, this is the total, not the number of additional licenses.

Executable name Enter the path and file name of the network database server executable (*dbsrv8.exe*) or MobiLink synchronization server (*dbmlsrv8.exe*) you are licensing. The path must be either absolute or relative to the current working directory.

You can view the current license information for a server executable without starting the server by entering only the executable name.

User name A user name for the license. This name appears on the database server window on startup. If there are spaces in the name, enclose it in double quotes.

Company name The company name for the license. This name appears on the database server window on startup. If there are spaces in the name, enclose it in double quotes.

The Log Transfer Manager

The Log Transfer Manager (LTM) is also known as a **replication agent**.

The LTM is required for any Adaptive Server Anywhere database that participates in a Replication Server installation as a primary site.

Using the dbltm command-line utility

Syntax	dbltm [<i>options</i>]	
Parameters	Option	Description
	-A	Do not filter updates
	-C <i>config_file</i>	Use the named configuration file
	-I <i>interface_file</i>	Use the named interface file
	-M	Recovery mode
	-S <i>LTM_name</i>	Specify an LTM name
	-dl	Display log messages on screen
	-ek <i>key</i>	Specify encryption key
	-ep	Prompt for encryption key
	-o <i>filename</i>	Log output messages to a file
	-os <i>size</i>	Maximum size of output file
	-ot <i>file</i>	Truncate file, and log output messages to file
	-q	Run in minimized window
	-s	Show Log Transfer Language (LTL) commands
	-ud	Run as a daemon [UNIX]
	-v	Verbose mode

Description The Adaptive Server Anywhere LTM reads a database transaction log and sends committed changes to Replication Server. The LTM is not required at replicate sites.

The LTM sends messages to Replication Server in a language named Log Transfer Language (LTL).

By default, the LTM uses a log file named *DBLTM.LOG* to hold status and other messages. You can use options to change the name of this file and to change the volume and type of messages that are sent to it.

Exit codes are 0 (success) or non-zero (failure).

Log Transfer Manager utility options

Use configuration file (-C) Use the configuration file *config_file* to determine the LTM settings. The default configuration file is *dbltm.cfg*.

☞ For a description of the configuration file, see "The LTM configuration file" on page 483.

Use interface file (-I) (Upper case i.) Use the named interfaces file. The interfaces file is the file created by SQLEDIT which holds the connection information for Open Servers. The default interfaces file is *SQL.ini* in the *ini* subdirectory of your *Sybase* directory.

Specify LTM name (-S) Provides the server name for this LTM. The default LTM_name is DBLTM_LTM. The LTM_name must correspond to the Open Server name for the LTM that was entered in SQLEDIT.

Do not filter updates (-A) Do not filter updates. By default, all changes made by the maintenance user are not replicated. If the -A option is set, these changes are replicated. This may be useful in non-hierarchical Replication Server installations, where a database acts as both a replicate site and as a primary site.

Recover mode (-M) This is used to initiate recovery actions. The LTM starts reading logs from the earliest available position. If the offline directory is specified in the configuration file, the LTM reads from the oldest offline log file.

Specify encryption key (-ek) This option allows you to specify the encryption key for strongly encrypted databases directly in the command. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way, including offline transaction logs. For strongly encrypted databases, you must specify either -ek or -ep, but not both. The command fails if you do not specify a key for a strongly encrypted database.

Prompt for encryption key (-ep) This option allows you to specify that you want to be prompted for the encryption key. This option causes a dialog box to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text. For strongly encrypted databases, you must specify either -ek or -ep, but not both. The command fails if you do not specify a key for a strongly encrypted database.

Display Log Transfer Language messages (-dl) Display Log Transfer Language (LTL) messages in the LTM window or at the command prompt and also in the log file. This should be used only to diagnose problems, and is not recommended in a production environment. It carries a significant performance penalty.

Log output messages to file (-o) Use a log file different from the default (*dbltn.log*). Write output messages from log transfer operations to this file.

Limit size of output file (-os) Specify the maximum size of the output file, in bytes. The minimum value is 10000 (ten thousand). If the log file grows to the point where it would exceed this limit, it is renamed to *yymmddAA.log*, and a new file named *yymmddAB.log* is started. The value of *xx* in *yymmddxx.log* is incremented for each file created on a given day.

Truncate log file and use named log file (-ot) Use a log file different from the default (*dbltn.log*), and truncate the log file (all existing content is deleted) when the LTM starts. Output messages from log transfer operations are sent to this file for later review.

Operate quietly (-q) Minimize the window when the LTM is started.

Log all LTL commands (-s) Log all LTL commands that are generated by the LTM. This should be used only to diagnose problems, and is not recommended in a production environment. It carries a significant performance penalty.

Run as a daemon (-ud) You can run the LTM as a daemon on UNIX operating systems. If you run in this manner, output is logged to the log file.

Operate in verbose mode (-v) Displays messages, other than LTL messages, for debugging purposes.

The LTM configuration file

The Adaptive Server Anywhere and Adaptive Server Enterprise LTM configuration files are very similar. This section describes the entries in the Adaptive Server Anywhere LTM configuration file, and the differences from the Adaptive Server Enterprise LTM configuration file.

The configuration file that an LTM uses is specified using the *-C* option.

LTM configuration file parameters

The following table describes each of the configuration parameters that the LTM recognizes. Parameters that are used by the Adaptive Server Enterprise LTM but not by the Adaptive Server Anywhere LTM are included in this list, and marked as either ignored (in which case they may be present in the configuration file, but have no effect) or as unsupported (in which case they will cause an error if present in the configuration file).

Parameter	Description
APC_pw	The password for the APC_user login name. This entry is present only in Adaptive Server Anywhere LTM configuration files.
APC_user	A user ID that is used when executing asynchronous procedures at the primary site. This user ID must have permissions appropriate for all asynchronous procedures at the primary site. This entry is present only in Adaptive Server Anywhere LTM configuration files.
backup_only	By default, this is off. If set to on, the LTM will replicate only backed-up transactions.
batch_ltl_cmds	Set to on (the default) to use batch mode. Batch mode can increase overall throughput, but may lead to longer response times.
batch_ltl_sz	The number of commands that are saved in the buffer before being sent to Replication Server, when batch_ltl_cmds is on . The default is 200.
batch_ltl_mem	The amount of memory that the buffer can use before its contents are sent to Replication Server, when batch_ltl_cmds is on . The default is 256 kb.
Continuous	By default, this is on. When set to off, the LTM automatically shuts down as soon as all committed data has been replicated.
LTM_admin_pw	The password for the LTM_admin_user login name.
LTM_admin_user	The system administrator LTM login name that is used to log in to the LTM. This parameter is required so that the LTM can check whether a user logging on to the LTM to shut it down has the correct login name.
LTM_charset	The Open Client/Open Server character set for the LTM to use.
LTM_language	The Open Client/Open Server language for the LTM to use.

Parameter	Description
LTM_sortorder	<p>The Open Client/Open Server sort order for the LTM to use to compare user names. You can specify any Adaptive Server Enterprise-supported sort order that is compatible with the LTM's character set. All sort orders in your replication system should be the same.</p> <p>The default sort order is a binary sort.</p>
maint_cmds_to_skip	[ignored]
qualify_table_owner	<p>Set to on for the LTM to end LTLs with table names and columns names, as well as table owners to Rep Server. The setting applies to all replicating tables, and the create replication definition statements must match this setting. The default is off.</p>
rep_func	<p>Set to on to use asynchronous procedure calls (APCs). The default is off.</p>
Retry	<p>The number of seconds to wait before retrying a failed connection to an Adaptive Server Anywhere database server or Replication Server. The default is 10 seconds.</p>
RS	<p>The name of the Replication Server to which the LTM is transferring the log.</p>
RS_pw	<p>The password for the RS_user login name.</p>
RS_source_db	<p>The name of the database whose log the LTM transfers to the Replication Server. This name must match the name of the database as defined within the Replication Server connection definitions. Most configurations use the same setting for both RS_Source_db and SQL_database configuration options.</p>
RS_source_ds	<p>The name of the server whose log the LTM transfers to the Replication Server. This name must match the name of the server as defined within the Replication Server connection definitions. Most configurations use the same setting for both RS_Source_ds and SQL_server configuration options.</p>
RS_user	<p>A login name for the LTM to use to log into Replication Server. The login name must have been granted connect source permission in the Replication Server.</p>
scan_retry	<p>The number of seconds that the LTM waits between scans of the transaction log. This</p>

Parameter	Description
	parameter is somewhat different in meaning to that of the Adaptive Server Enterprise LTM. The Adaptive Server Anywhere server does not <i>wake up</i> and scan the log when records arrive in the log. For this reason, you may wish to set the <i>scan_retry</i> value to a smaller number than that for an Adaptive Server Enterprise LTM.
skip_ltl_cmd_err	[ignored]
SQL_database	The primary site database name on the server SQL_server to which the LTM connects. For Adaptive Server Enterprise during recovery, this is the temporary database whose logs the LTM will transfer to Replication Server. The Adaptive Server Anywhere LTM uses the SQL_log_files parameter to locate offline transaction logs.
SQL_log_files	A directory that holds off-line transaction logs. The directory must exist when the LTM starts up. This entry is present only in Adaptive Server Anywhere LTM configuration files.
SQL_pw	The password for the SQL_user user ID.
SQL_server	The name of the primary site Adaptive Server Anywhere server to which the LTM connects. For Adaptive Server Enterprise during recovery, this is a data server with a temporary database whose logs the LTM will transfer to Replication Server. The LTM uses the SQL_log_files parameter to locate offline transaction logs.
SQL_user	The login name that the LTM uses to connect to the database specified by RS_source_ds and RS_source_db.

Example configuration file

- ◆ The following is a sample LTM configuration file.


```
# This is a comment line
# Names are case sensitive.
SQL_user=sa
SQL_pw=sysadmin
SQL_server=PRIMESV
SQL_database=primedb
RS_source_ds=PRIMEOS
RS_source_db=primedb
RS=MY_REPSERVER
RS_user=sa
RS_pw=sysadmin
LTM_admin_user=DBA
LTM_admin_pw=SQL
LTM_charset=cp850
scan_retry=2
SQL_log_files=e:\logs\backup
APC_user=sa
APC_pw=sysadmin
```

The Log Translation utility

With the Log Translation utility, you can translate a transaction log into a SQL command file.

You can access the Log Translation utility in the following ways:

- ◆ From Sybase Central, using the Translate Log File wizard.
- ◆ At the command prompt, using the *dbtran* command. This is useful for incorporating into batch or command files.

Translating a transaction log using the Translate Log File wizard

❖ **To translate a transaction log into a command file:**

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Translate Log File in the right pane.

The Translate Log File wizard appears.

- 3 Follow the instructions in the wizard.

Tip

You can also access the Translate Log File wizard by choosing Tools►Adaptive Server Anywhere 8►Translate Log File.

Translating a transaction log using the dbtran command-line utility

Syntax

Running against a transaction log.

```
dbtran [ options ] transaction-log [ SQL-file ]
```

Running against a database server.

```
dbtran [ options ]
```

Option	Description
-a	Include uncommitted transactions
-c <i>"keyword=value; ..."</i>	Supply database connection parameters—cannot be used with a transaction log name
-d	Display output in chronological order
-ek <i>key</i>	Specify encryption key
-ep	Prompt for encryption key
-f	Output only since the last checkpoint
-g	Include audit records in output
-ir <i>offset1,offset2</i>	Include only the portion of the log between the two specified offsets
-is <i>source,...</i>	Include only rows originating from the specified sources
-it <i>user.table,...</i>	Include only operations on specified tables by specifying a comma-delimited list of user.table
-j <i>date/time</i>	Output from the last checkpoint prior to the given time
-m	Specify transaction logs directory (requires -n option)
-n <i>filename</i>	Output SQL file, when used against a database server
-o <i>filename</i>	Log output messages to a file
-q	Run quietly, do not print messages
-r	Remove uncommitted transactions (default)
-rsu <i>username,...</i>	Override default Replication Server user names
-s	Produce ANSI standard SQL UPDATE transactions
-sr	Generate SQL Remote comments
-t	Include trigger-generated transactions in output
-u <i>userid,...</i>	Translate transactions for listed users only
-x <i>userid,...</i>	Exclude transactions for listed users
-y	Replace file without confirmation
-z	Include trigger-generated transactions as comments only
<i>Transaction-log</i>	Log file to be translated—cannot be used together with -c or -n
<i>SQL-file</i>	Output file containing the translated information—for use with <i>transaction-log</i> only

Description

The *dbtran* utility takes the information in a transaction log and places it as a set of SQL statements and comments into an output file. The utility can be run in the following ways:

- ◆ **Against a database server** Run in this way, the utility is a standard client application. It connects to the database server using the connection string specified following the *-c* option, and places output in a file specified with the *-n* option. DBA authority is required to run in this way.

The following command translates log information from the server **asademo** and places the output in a file named **asademo.SQL**.

```
dbtran -c "eng=asademo;dbn=asademo;uid=DBA;pwd=SQL" -n asademo.sql
```

- ◆ **Against a transaction log file** Run in this way, the utility acts directly against a transaction log file. You should protect your transaction log file from general access if you wish to prevent users from having the capability of running this statement.

```
dbtran asademo.log asademo.sql
```

When the *dbtran* utility runs, it displays the earliest log offset in the transaction log. This can be an effective method for determining the order in which multiple log files were generated.

Exit codes are 0 (success) or non-zero (failure).

🔗 For more information about the Log translation utility options, see "Log translation utility options" on page 490

Log translation utility options

Include uncommitted transactions (-a) The transaction log contains any changes made before the most recent COMMIT by any transaction. Changes made after the most recent commit are not present in the transaction log.

Connection string (-c) When running the utility against a database server, this parameter specifies the connection string.

DBA authority is required to run *dbtran*.

🔗 For a description of the connection parameters, see "Connection parameters" on page 70.

Display output in chronological order (-d) Transactions are displayed in order from earliest to latest. This feature is provided primarily for use when auditing database activity: the output of this command should not be applied against a database.

Specify encryption key (-ek) This option allows you to specify the encryption key for strongly encrypted databases directly in the command. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way.

For strongly encrypted databases, you must specify either `-ek` or `-ep`, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

If you are running against a database server (using the `-c` option), make sure you specify the key using a connection parameter, not using the `-ek` option. For example, the following command gets the transaction log information about database *enc.db* from engine **sample**, and saves its output in *log.sql*.

```
dbtran -n log.sql -c eng=sample;dbf=enc.db;uid=dba;pwd=sql;dbkey=mykey
```

Prompt for encryption key (-ep) This option allows you to specify in the command that you want to be prompted for the encryption key. This option causes a dialog box to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text.


For strongly encrypted databases, you must specify either `-ek` or `-ep`, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

If you are running against a database server (using the `-c` option), make sure you specify the key using a connection parameter, not using the `-ep` option. For example, the following command gets the transaction log information about database *enc.db* from engine **sample**, and saves its output in *log.sql*.

```
dbtran -n log.sql -c eng=sample;dbf=enc.db;uid=dba;pwd=sql;dbkey=mykey
```

Output from last checkpoint only (-f) Only transactions that were completed since the last checkpoint are output.

Include audit information (-g) If the AUDITING database option is turned on, auditing information is added to the transaction log. You can include this information as comments in the output file using this option.

 For more information, see "AUDITING option" on page 553.

Include rows from specified sources (-is) Isolate operations on rows that have been modified by operations from one or more of the following sources, specified as a comma-separated list:

- ◆ **All** All rows. This is the default setting.
- ◆ **SQLRemote** Include only rows that were modified using SQL Remote. You can also use the short form **SR**.

◆ **RepServer** Include only rows that were modified using the Replication Agent (LTM) and Replication Server. You can also use the short form **RS**.

◆ **Local** Include only rows that are not replicated.

Include offset range (-ir) Isolate a portion of the transaction log between two specified offsets.

Include specified tables (-it) Isolate those operations on the specified, comma-separated list of tables. Each table should be specified as *owner.table*.

Output from the last checkpoint prior to a given date (-j) Only transactions from the most recent checkpoint prior to the given date and/or time are translated. The user-provided argument can be a date, time or date and time enclosed in quotes. If the time is omitted, the time is assumed to be the beginning of the day. If the date is omitted, the current day is assumed. The following is an acceptable format for the date and time: "YY/MM/DD HH:NN".

Transaction logs directory (-m) Use this option to specify a directory that contains transaction logs. This option must be used in conjunction with the -n option.

Output file (-n) When you run the *dbtran* utility against a database server, use this option to specify the output file that holds the SQL statements.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Do not include uncommitted transactions (-r) Remove any transactions that were not committed. This is the default behavior.

Override Replication Server user names (-rsu) By default, the -is option assumes the default Replication Server user names of *dbmaint* and *sa*. You can override this assumption using the -rsu option with a comma-separated list of user names.

Generate ANSI standard SQL UPDATE (-s) If the option is not used, and there is no primary key or unique index on a table, the Translation utility generates UPDATE statements with a non-standard FIRST keyword in case of duplicate rows. If the option is used, the FIRST keyword is omitted for compatibility with the SQL standard.

Generate SQL Remote comments (-sr) Place generated comments in the output file describing how SQL Remote distributes operations to remote sites.

Include transactions generated by triggers (-t) By default, actions carried out by triggers are not included in the command file. If the matching trigger is in place in the database, when the command file is run against the database, the trigger will carry out the actions automatically. Trigger actions should be included if the matching trigger does not exist in the database against which the command file is to be run.

Output transactions for listed users only (-u) This option allows you to limit the output from the transaction log to include only specified users.

Output transactions except for listed users (-x) This option allows you to limit the output from the transaction log to exclude specified users.

Operate without confirming actions (-y) Choosing this option automatically replaces existing command file(s) without prompting you for confirmation.

Include transactions generated by triggers as comments only (-z) Transactions that were generated by triggers will be included only as comments in the output file.

The Ping utility

The Ping utility is provided to assist in troubleshooting connection problems.

Troubleshooting connections using the dbping command-line utility

Syntax

dbping [*options*]

Option	Description
-c " <i>keyword=value; ...</i> "	Database connection parameters
-d	Make a database connection if the server is found
-l <i>library</i>	Load specified ODBC driver or library
-m	Use ODBC driver manager
-o <i>filename</i>	Log output messages to a file
-pc <i>property,...</i>	Report specified connection properties
-pd <i>property,...</i>	Report specified database properties
-ps <i>property,...</i>	Report specified database server properties
-q	Operate quietly—do not print messages
-z	Display debugging information

Description

The *dbping* utility is a tool to help debug connection problems. It takes a full or partial connection string and returns a message indicating whether the attempt to locate a server or database, or to connect, was successful.

The utility can be used for embedded SQL or ODBC connections. It cannot be used for jConnect (TDS) connections.

Exit codes are 0 (success) or non-zero (failure).

Ping utility options

Connection parameters (-c) For a description of the connection parameters, see "Connection parameters" on page 70. If the connection parameters are not specified, connection parameters from the SQLCONNECT environment variable are used, if set.

Make database connection (-d) Ping the database, not just the server.

If you do not supply the `-d` option, then *dbping* reports success if it finds the server specified by the `-c` option. If you do supply the `-d` option, then *dbping* reports success only if connects to the server and also connect to a database.

For example, if you have a server named **blair** running the database *sample*, the following succeeds:

```
dbping -c "eng=blair;dbn=sample"
```

The following command fails, with the message Ping database failed -- specified database not found:

```
dbping -d -c "eng=blair;dbn=sample"
```

Load specified library (-l) Specify the library to use (without its file extension). This option avoids the use of the ODBC driver manager, and so is particularly useful on UNIX operating systems.

For example, the following command loads the ODBC driver directly:

```
dbping -m -c "dsn=ASA 8.0 Sample" -l dbodbc8
```

On UNIX, if you wish to use a threaded connection library, you must use the threaded version of the Ping utility, *dbping_r*.

Use ODBC to connect (-m) Establish a connection using ODBC. By default, the utility connects using the embedded SQL interface.

Report connection properties (-pc) Upon connection, display the specified connection properties. Supply the properties in a comma-separated list. You must specify enough connection information to establish a database connection if you use this option.

☞ For a list of connection properties, see "Connection-level properties" on page 618.

For example, the following command displays the `DIVIDE_BY_ZERO_ERROR` option setting, which is available as a connection property.

```
dbping -c ... -pc Divide_by_zero_error
```


Report database properties (-pd) Upon connection, display the specified database properties. Supply the properties in a comma-separated list. You must specify enough connection information to establish a database connection if you use this option.

☞ For a list of database properties, see "Database-level properties" on page 630.

For example, the following command displays the Java version in use by the database:

```
dbping -c ... -pd JDKVersion
```

Report database server properties (-ps) Upon connection, display the specified database server properties. Supply the properties in a comma-separated list.

 For a list of database server properties, see "Server-level properties" on page 625.

For example, the following command displays the number of licensed seats for the database server:

```
dbping -c ... -ps LicenseCount
```


Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) If *dbping* fails, a message is always displayed. If *dbping* succeeds, no message appears if *-q* is specified.

Display debugging information (-z) This option is available only when an embedded SQL connection is being attempted. That is, it cannot be combined with *-m* or *-l*. It displays the network communication protocols used to attempt connection, and other diagnostic messages.

The Rebuild utility

Databases can be rebuilt using the Rebuild batch file, command file, or shell script, which invokes a series of utilities to rebuild a database, or as part of the unload process using the Unload Database wizard in Sybase Central.

 For more information about the Unload Database wizard, see "The Unload utility" on page 513.

Rebuilding a database using the rebuild batch or command file

Syntax `rebuild old-database new-database [DBA-password]`

See also "Unloading a database using the dbunload command-line utility" on page 514
"Creating a database using the dbinit command-line utility" on page 466
"The Interactive SQL utility" on page 472

Description This batch file, command file, or shell script uses *dbunload* to rebuild *old-database* into *new-database*. This is a simple script, but it helps document the rebuilding process, and provides a basis for customization. Both database names should be specified without extensions. An extension of *.db* is automatically added.

You can use *dbunload* with the `-ar` option to carry out unloading and reloading without using the rebuild batch file.

The *DBA-password* must be specified if the password to the DBA user ID in the *old-database* is not the initial password SQL.

Rebuild runs the *dbunload*, *dbinit*, and Interactive SQL commands with the default options. If you need different options, you will need to run the three commands separately.

Exit codes are 0 (success) or non-zero (failure).

The Server Location utility

The server location utility is provided to assist in diagnosing connection problems by locating database servers on the immediate TCP/IP network.

Locating servers using the dblocate command-line utility

Syntax

dblocate [*options*]

Option	Description
-o <i>filename</i>	Log output messages to a file
-q	Operate quietly—do not print messages

Description

The *dblocate* utility locates any Adaptive Server Anywhere database servers running over TCP/IP on the immediate network. It prints a list of database servers and their addresses.

Depending on your network, the utility may take several seconds before printing its results.

Exit codes are 0 (success) or non-zero (failure).

Server location utility options

- Log output messages to file (-o)** Write output messages to the named file.
- Operate quietly (-q)** Do not display output messages.

The Service Creation utility

The Service Creation utility is a tool used to create, delete, and modify Adaptive Server Anywhere services.

Managing services using the dbsvc command-line utility

Syntax

dbsvc [options] <svc>

dbsvc [-q] [-y] -d <svc>

dbsvc [-q] -g <svc>

dbsvc [-q] -l

dbsvc [-q] [-y] <creation options> -w <svc> <details>


Option	Description
-a acct	Account name to use (used with -p)
-as	Use local system account
-d	Delete a service
-I	Allow service to interact with desktop
-g	Get details of a service
-l	List all Adaptive Server Anywhere services
-p	Password for account (used with -a)
-q	Do not print banner
-rg dependency,...	Specify group dependencies when creating a service
-rs dependency,...	Specify service dependencies when creating a service
-s startup	Startup option (default manual). You must specify automatic, manual, or disabled
-t type	Type of service
-y	Delete or overwrite service without confirmation
-w executable parameters	Create write service

Description

A service runs a database server or other application with a set of options. This utility provides a comprehensive way of managing Adaptive Server Anywhere services on Windows. The Service Creation utility provides the same functionality as the Service Creation wizard in Sybase Central.

You must be a member of the Administrators group on the local machine to use the Service Creation utility.

Exit codes are 0 (success) or non-zero (failure).

 For more information about services, see "Understanding Windows services" on page 23.

Examples

Create a personal server service called **myserv**, which starts the specified engine with the specified parameters. The engine runs as the **LocalSystem** user:

```
dbsvc -as -w myserv "C:\Program Files\Sybase\SQL  
Anywhere 8\win32\dbeng8.exe" -n myeng -c 8m "C:\Program  
Files\Sybase\SQL Anywhere 8\sample.db"
```

Create a network server service called **mynetworkserv**. The server runs under the local account, and starts automatically when the machine is booted:

```
dbsvc -as -s auto -t network -w mynetworkserv  
"C:\Program Files\Sybase\SQL Anywhere  
8\win32\dsrv8.exe" -x tcpip -c 8m "C:\Program  
Files\Sybase\SQL Anywhere 8\sample.db"
```

List all details about service **myserv**:

```
dbsvc -g myserv
```

Delete the service called **myserv**, without prompting for confirmation:

```
dbsvc -y -d myserv
```

Create a service dependent on the Workstation service and the TDI group:

```
dbsvc -rs Workstation -rg TDI -w ...
```

Create a service called **Adaptive Server Anywhere - mysyncservice**:

```
dbsvc -as -s manual -t dbmlsync -w mysyncservice  
"C:\Program Files\Sybase\SQL Anywhere  
8\win32\dbmlsync.exe" -c "dsn=ultralite 8.0 sample"
```

Service creation utility options

Account name (-a) All services run under a Windows account. If you run under an account you've created, you must name the account with the **-a** option and supply a password with the **-p** option.

Use local system account (-as) All services run under a Windows account. Using the `-as` option, the service will run under a Windows account. No password is required.

Delete a service (-d) Removes the server name from the list of services. If you supply `-y`, any service is deleted without confirmation.

Allow service to interact with desktop (-i) Displays an icon that you can double-click to display the server window.

Get details of a service (-g) Lists the definition of the service, not including the password.

List all Adaptive Server Anywhere services (-l) Lists the available Adaptive Server Anywhere services.

Password for account (-p) Use this option with the `-a` option to specify the password for the account the service runs under.

Do not print banner (-q) Suppress the informational banner. The `-q` option can be used with any of the `-d`, `-g`, `-l`, or `-w` options.

Set group dependencies (-rg) At least one service from each of the groups in the list must be started before the service being created is allowed to start.

Set service dependencies (-rs) All the services in the list must have started before the service being created is allowed to start.

Startup option (-s) Sets startup behavior for Adaptive Server Anywhere services. You can set startup behavior to Automatic, Manual, or Disabled.

Type of service (-t type) Specifies the executable to run for this service. You can choose from the following types:

Type	Description
Network	Adaptive Server Anywhere network database server (<i>dbsrv8</i>)
Standalone	Adaptive Server Anywhere personal database server (<i>dbeng8</i>)
DBRemote	SQL Remote Message Agent (<i>dbremote</i>)
MobiLink	MobiLink synchronization server (<i>dbmlsrv8</i>)
DBMLSync	MobiLink synchronization client (<i>dbmlsync</i>)

The default setting is **Standalone**. If creating a service, you must specify options for the appropriate executable along with the type.

Create service (-w) Creates a new service, or overwrites one if one of the same name exists. If you supply -y, any existing service is overwritten without confirmation.

You must supply the full path to the executable that you wish to use as a service, as the account under which the service is running may not have the appropriate SQL Anywhere directory in its path.

You must supply parameters appropriate for the service you are creating. For more information, see the following locations:

- ◆ **dbsrv8 and dbeng8** "The database server" on page 120.
- ◆ **dbmlsrv8** "MobiLink Synchronization Server Options" on page 379 of the book *MobiLink Synchronization User's Guide*.
- ◆ **dbremote** "The Message Agent" on page 302 of the book *SQL Remote User's Guide*.

Delete or overwrite service without confirmation (-y) Automatically carries out the action without prompting for confirmation. This option can be used with the -w or -d options.

The Spawn utility

This utility is provided to start a database server in the background.

Running a server in the background using the dbspawn command-line utility

Syntax

dbspawn [**-q**] *server command*

Option	Description
-f	Do not check for a running server
-p	Report operating system process ID
-q	Quiet mode—do not print messages

Description

The *dbspawn* utility is provided to start a server in the background. *dbspawn* starts the server in the background and returns with an exit code of 0 (success) or non-zero (failure). If the server specified in *server-command* is already running, *dbspawn* reports failure.

The *dbspawn* utility is useful for starting a server from a batch file, especially when subsequent commands in the batch file require a server that is accepting requests.

If the specified path includes at least one space, and if you are executing the spawn utility from a batch file, you must enclose the path in one set of double quotes. For example,

```
dbspawn dbeng8 "C:\Program Files\Sybase\SQL Anywhere
8\asademo.db"
```

If the specified path includes at least one space, and if you are executing the spawn utility at the command prompt, you must provide an extra set of quotes, and escape the quotes around the database file. For example,

```
dbspawn dbeng8 "\"C:\Program Files\Sybase\SQL Anywhere
8\asademo.db\""
```

If the specified path does not contain spaces, then quotes are not required.

For a description of the server commands, see "The database server" on page 120.

Spawn utility options

Do not check for a running server (-f) If a database server is already running, the *dbspawn* command may use it. This option forces *dbspawn* to start a new database server each time it is run.

Report process ID (-p) The operating system process ID of the database server process. For example:

```
dbspawn -p dbeng8 -n newserver
```

reports a message of the following form to the command prompt:

```
New process id is 306
```

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

The Stop utility

The Stop utility stops a database server. You can use the `-d` option to stop a specified database.

The Stop utility can only be run at the command prompt. In windowed environments, you can stop a database server by clicking Shutdown on the server window.

Stopping a database server using the dbstop command-line utility

Syntax

dbstop [*options*]

Option	Description
<i>server-name</i>	Name of a local database server to stop.
<code>-c "keyword=value; ..."</code>	Connection parameters
<code>-d</code>	Stop specified database only.
<code>-o filename</code>	Log output messages to a file
<code>-q</code>	Quiet mode—do not print messages
<code>-x</code>	Do not stop if there are active connections
<code>-y</code>	Stop without prompting even if there are active connections

Description

Options let you control whether a server is stopped even if there are active connections, and whether to stop a server or only a database.

Exit codes are 0 (success) or non-zero (failure).

For more information about the Stop utility options, see "Stop utility options" on page 505.

Stop utility options

Server name The name of a database server running on the current machine. The database server must be started so that no permissions are required to shut it down. The personal database server starts in this mode by default. For the network database server, you must supply the `-gk all` option.

If you supply a server name, do not supply connection parameters as well.

For more information, see "`-gk server option`" on page 140.

Connection parameters (-c) When stopping a network server, you must supply a connection string with a user ID that has permissions to stop the server. By default, DBA permission is required on the network server, and all users can shut down a personal server, but the `-gk` server option can be used to change this.

The behavior of *dbstop* can be controlled if there are active connections on a server. If there are active connections, *dbstop* provides a prompt asking if you wish to shut down the server. If you specify **unconditional=true**, the server is shut down without prompting, even if there are active connections.

If you supply connection parameters, do not supply a server name as well.

🔗 For more information, see "Connection parameters" on page 70, "Unconditional connection parameter" on page 187, and "`-gk` server option" on page 140.

Stop database only (-d) Do not stop the database server. Instead, only stop the database specified in the connection string.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not print a message if the database was not running.

Do not stop if there are active connections (-x) Do not stop the server if there are still active connections to the server.

Stop without prompting (-y) Stop the server even if there are still active connections to the server.

Example 1

You are running the server named **myserver** without a database. To stop the server, specify the utility database as a **DatabaseName (DBN)** connection parameter:

```
dbstop -c "uid=DBA;pwd=SQL;eng=myserver;dbn=utility_db"
```

Example 2

You are running the server named **myserver** with the database **asademo.db** started. To stop the server and database:

```
dbstop -c "uid=DBA;pwd=SQL;eng=myserver "
```

The Transaction Log utility

With the Transaction Log utility, you can display or change the name of the transaction log or transaction log mirror associated with a database. You can also stop a database from maintaining a transaction log or mirror, or start maintaining a transaction log or mirror.

A transaction log mirror is a duplicate copy of a transaction log, maintained by the database in tandem.

The name of the transaction log is first set when the database is initialized. The Transaction Log utility works with database files or with write files. The database server must not be running on that database when the transaction log filename is changed (or an error message appears).

You can access the Transaction Log utility in the following ways:

- ◆ From Sybase Central, using the Change Log File Settings wizard.
- ◆ At the command prompt, using the *dblog* command.

Managing log files using the Change Log File Settings wizard

❖ To change a transaction log file name:

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Change Log File Settings in the right pane.
The Change Log File Settings wizard appears.
- 3 Follow the instructions in the wizard.

Tip

You can also access the Change Log File Settings wizard by choosing Tools►Adaptive Server Anywhere 8►Change Log File Settings.

Managing log files using the dblog command-line utility

Syntax

dblog [*options*] *database-file*

Option	Description
-ek <i>key</i>	Specify encryption key
-ep	Prompt for encryption key
-g <i>n</i>	Sets the LTM generation number to <i>n</i>
-il	Ignores the LTM truncation offset stored in the database
-ir	Ignores the SQL Remote truncation offset stored in the database
-m <i>mirror-name</i>	Set transaction log mirror name
-n	No longer use a transaction log or mirror log
-o <i>filename</i>	Log output messages to a file
-q	Quiet mode—do not print messages
-r	No longer use a transaction log mirror
-t <i>log-name</i>	Set the transaction log name
-x <i>n</i>	Set the transaction log current relative offset to <i>n</i>
-z <i>n</i>	Set the transaction log starting offset to <i>n</i>


Description

The dblog utility allows you to display or change the name of the transaction log or transaction log mirror associated with a database. You can also stop a database from maintaining a transaction log or mirror, or start maintaining a transaction log or mirror.

The utility displays additional information about the transaction log, including the following:

- ◆ Version number
- ◆ Starting offset, for use in replication
- ◆ Ending offset, for use in replication
- ◆ Page size
- ◆ Total number of pages
- ◆ Number of empty pages
- ◆ Percentage of the log file in use

Exit codes are 0 (success) or non-zero (failure).

 For more information about the Transaction log utility options, see "Transaction log utility options" on page 509.

Transaction log utility options

Specify encryption key (-ek) This option allows you to specify the encryption key for strongly encrypted databases directly in the command. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way. For strongly encrypted databases, you must specify either `-ek` or `-ep`, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Prompt for encryption key (-ep) This option allows you to specify that you want to be prompted for the encryption key. This option causes a dialog box to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text. For strongly encrypted databases, you must specify either `-ek` or `-ep`, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Set the generation number (-g) Use this option if you are using the Log Transfer Manager to participate in a Replication Server installation. It can be used after a backup is restored, to set the generation number. It performs the same function as the following Replication Server function:

```
dbcc settrunc( 'ltm', 'gen_id', n )
```

☞ For information on generation numbers and *dbcc*, see your Replication Server documentation.

Ignore the LTM truncation offset (-il) Use this option if you are using the Log Transfer Manager to participate in a Replication Server installation. It performs the same function as the following Replication Server function:

```
dbcc settrunc( 'ltm', 'ignore' )
```

☞ For information on *dbcc*, see your Replication Server documentation.

Ignore the SQL Remote truncation offset (-ir) Use this option if you are using the Log Transfer Manager to participate in a Replication Server installation and a SQL Remote installation. It resets the offset that is kept for the `DELETE_OLD_LOGS` option, allowing transaction logs to be deleted when they are no longer needed.

Set the name of the transaction log mirror file (-m) This option sets a filename for a new transaction log mirror. If the database is not currently using a transaction log mirror, it starts using one. If the database is already using a transaction log mirror, it changes to using the new file as its transaction log mirror.

No longer use a transaction log (-n) Stop using a transaction log, and stop using a mirror log. Without a transaction log, the database can no longer participate in data replication or use the transaction log in data recovery.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

No longer use a transaction log mirror (-r) For databases that maintain a mirrored transaction log, this option changes their behavior to maintain only a single transaction log.

Set the name of the transaction log file (-t) This option sets a filename for a new transaction log. If the database is not currently using a transaction log, it starts using one. If the database is already using a transaction log, it changes to using the new file as its transaction log.

Set the current log offset (-x) For use when reloading a SQL Remote consolidated database. This option resets the current log offset so that the database can take part in replication.

☞ For information on how to use this option, see "Unloading and reloading a database participating in replication" on page 265 of the book *SQL Remote User's Guide*.

Set the starting log offset (-z) For use when reloading a SQL Remote consolidated database. This option resets the starting log offset so that the database can take part in replication.

☞ For information on how to use this option, see "Unloading and reloading a database participating in replication" on page 265 of the book *SQL Remote User's Guide*.

The Uncompression utility

With the Uncompression utility, you can expand a compressed database file created by the Compression utility. The Uncompression utility reads the compressed file and restores the database file to its uncompressed state.

You can access the Uncompression utility in the following ways:

- ◆ From Sybase Central, using the Uncompress Database wizard.
- ◆ At the command prompt, using the *dbexpand* command. This is useful for incorporating into batch or command files.

Uncompressing a database using the Uncompress Database wizard

❖ **To uncompress a compressed database file:**

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Uncompress Database in the right pane.

The Uncompress Database wizard appears.

- 3 Follow the instructions in the wizard.

Tip

You can also access the Uncompress Database wizard by choosing Tools►Adaptive Server Anywhere 8►Uncompress Database.

Uncompressing a database using the dbexpand command-line utility

Syntax

dbexpand [*options*] *compressed-database-file* [*database-file*]

Option	Description
-ek <i>key</i>	Specify encryption key
-ep	Prompt for encryption key
-o <i>filename</i>	Log output messages to a file
-q	Operate quietly—do not print messages
-y	Erase an existing output file without confirmation

Description

The input filename extension defaults to *cdb*. The output filename (with extension) must not have the same name as the input filename (with extension).

Exit codes are 0 (success) or non-zero (failure).

☞ For more information about the Uncompression utility options, see "Uncompression utility options" on page 512.

Uncompression utility options

Specify encryption key (-ek) This option allows you to specify the encryption key for strongly encrypted databases directly in the command. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way. For strongly encrypted databases, you must specify either *-ek* or *-ep*, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Prompt for encryption key (-ep) This option allows you to specify that you want to be prompted for the encryption key. This option causes a dialog box to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text. For strongly encrypted databases, you must specify either *-ek* or *-ep*, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Operate without confirming actions (-y) Choosing this option automatically replaces an existing database file without prompting you for confirmation.

The Unload utility

With the Unload utility, you can unload a database and put a set of data files in a named directory. The Unload utility creates an Interactive SQL command file to rebuild your database. It also unloads all of the data in each of your tables, into files in the specified directory in comma-delimited format. Binary data is properly represented with escape sequences.

You can also use the Unload utility to directly create a new database from an existing one. This avoids potential security problems with the database contents being written to ordinary disk files.

Accessing the Unload utility

You can access the Unload utility in the following ways:

- ◆ From Sybase Central, using the Unload Database wizard.
- ◆ At the command prompt, using the *dbunload* command. This is useful for incorporating into batch or command files.

*The Unload utility should be run from a user ID with DBA authority. This is the only way you can be sure of having the necessary privileges to unload all the data. In addition, the *reload.sql* file should be run from the DBA user ID. (Usually, it will be run on a new database where the only user ID is **DBA** with password **SQL**.)*

The database server `-gl` option controls the permissions required to unload data from the database. For information, see "`-gl` server option" on page 141.


Objects owned by dbo

The **dbo** user ID owns a set of Adaptive Server Enterprise-compatible system objects in a database.

The Unload utility does not unload the objects that were created for the **dbo** user ID during database creation. Changes made to these objects, such as redefining a system procedure, are lost when the data is unloaded. Any objects that were created by the **dbo** user ID since the initialization of the database are unloaded by the Unload utility, and so these objects are preserved.

Unloading and replication

There are special considerations for unloading databases involved in replication.

 For information, see "Unloading and reloading a database participating in replication" on page 265 of the book *SQL Remote User's Guide*.

Unloading a database using the Unload Database wizard utility


- ❖ **To unload a database file or a running database:**
 - 1 Open the Utilities folder in the left pane.
 - 2 Double-click Uncompress Database in the right pane.

The Uncompress Database wizard appears.
 - 3 Follow the instructions in the wizard.
 - ◆ Right-clicking the database and then clicking Unload Database from the popup menu

Tip

You can also access the Unload Database wizard from within Sybase Central using any of the following methods:

- ◆ choosing Tools>Adaptive Server Anywhere 8>Unload Database.
- ◆ Selecting a database in the left pane, and choosing Unload Database from the File menu.
- ◆ Right-clicking a database, and choosing Unload Database from the popup menu.

 For full information on unloading a database from Sybase Central, see "Exporting a database" on page 437 of the book *ASA SQL User's Guide*.

Unloading a database using the dbunload command-line utility

Syntax

```
dbunload [ options ] [directory]
```

Option	Description
-ac "keyword=value; ..."	Supply connection parameters for the reload
-an database	Creates a database file with the same settings as the database being unloaded, and automatically reloads it
-ar directory	Rebuild and replace database
-c "keyword=value; ..."	Supply database connection parameters for unload
-d	Unload data only
-e table, ...	Do not unload listed tables
-ea algorithm	Specify which strong encryption algorithm to encrypt your database with: you can choose AES or

Option	Description
	MDSR
-ek <i>key</i>	Specify encryption key
-ep	Prompt for encryption key
-ii	Internal unload, internal reload (default)
-ix	Internal unload, external reload
-j <i>nnn</i>	Repeated unload of view creation statements
-n	No data—schema definition only
-o <i>filename</i>	Log output messages to a file
-p <i>char</i>	Escape character for external unloads (default "\")
-q	Quiet mode—no windows or messages
-r <i>reload-file</i>	Specify name and directory of generated reload Interactive SQL command file (default <i>reload.sql</i>)
-t <i>table,...</i>	Unload only the listed tables
-u	Unordered data. Do not use indexes to unload data.
-v	Verbose messages
-xi	External unload, internal reload
-xx	External unload, external reload
-y	Replace the command file without confirmation

Description

The *directory* is the destination directory where the unloaded data is to be placed.

In the default mode, or if **-ii** or **-ix** is used, the directory used by *dbunload* to hold the data is relative to the database server, not to the current directory of the user.

☞ For details of how to supply a filename and path in this mode, see "UNLOAD TABLE statement" on page 573 of the book *ASA SQL Reference Manual*.

If the **-xi** or **-xx** is used, the directory is relative to the current directory of the user.

The *reload.sql* command file is always relative to the current directory of the user.

If no list of tables is supplied, the whole database is unloaded. If a list of tables is supplied, only those tables are unloaded.

If you want to unload a strongly encrypted database, you will be required to provide the encryption KEY. You can use the **EncryptionKey (DBKEY)** connection parameter to provide the KEY in the command. Alternatively, if you want to be prompted for the encryption key rather than entering it in plain view, you can use the `-ep` option as follows:

```
dbunload -c "dbf=enc.db;start=dbeng8 -ep"
```

If you are using `dbunload -an` to unload a database and reload into a new one, and you want to use the `-ek` or `-ep` options to set the encryption password for the new database, keep the following in mind:

- ◆ If the original database is strongly encrypted, you will need to specify the key for the original database using the **EncryptionKey (DBKEY)** connection parameter in the `-c` option, not using `-ek` or `-ep`.
- ◆ Using the `-ek` and `-ep` options, it is possible to unload an unencrypted database and reload into a new, strongly encrypted database. When you use `-ep` and `-an`, you must confirm the key correctly or the unload fails.
- ◆ If the original database is strongly encrypted, but the `-ek` and `-ep` options are not used, then the new database will be encrypted without requiring a key.
- ◆ The `-ek` and `-ep` options are ignored if `-an` is not specified.

Exit codes are 0 (success) or non-zero (failure).

🔗 For more information about the Unload utility options, see "Unload utility options" on page 516.

🔗 There are special considerations for unloading databases involved in replication. For information, see "Unloading and reloading a database participating in replication" on page 265 of the book *SQL Remote User's Guide*.

For more information about encryption, see "`-ep` server option" on page 137 or the "Encryption Key connection parameter" on page 179.

Unload utility options

Connection parameters for reload database (`-ac`) This option causes the Unload utility to connect to an existing database and reload the data directly into it. You can combine the operation of unloading a database and reloading the results into an existing database using this option.

Typically, you would create a new database using the Initialization utility, and then reload it using this option. This method is useful when you want to change the initialization options, such as page size, or collation.

For example, the following command (which should be entered all on one line) loads a copy of the *asademo.db* database into an existing database file named *newdemo.db*:

```
dbunload -c "uid=DBA;pwd=SQL;dbf=asademo.db" -ac  
"uid=DBA;pwd=SQL;dbf=newdemo.db"
```

If you use this option, no interim copy of the data is created on disk, so you do not specify an unload directory in the command. This provides greater security for your data, but at some cost for performance.

You do not need to specify a directory for this option

Create a database for reloading (–an) You can combine the operations of unloading a database, creating a new database, and loading the data using this option. This option applies to personal server connections, and network server connections over shared memory.

Typically, you would use this option when you do not want to change the initialization option of your database. The options specified when you created the source database are used to create the new database.

For example, the following command (which should be entered all on one line) creates a new database file named *asacopy.db* and copies the schema and data of *asademo.db* into it:

```
dbunload -c "uid=DBA;pwd=SQL;dbf=asademo.db" -an  
asacopy.db
```

If you use this option, no interim copy of the data is created on disk, so you do not specify an unload directory in the command. This provides greater security for your data, but at some cost for performance.

You do not need to specify a directory for this option

Rebuild and replace database (–ar) This option creates a new database with the same settings as the old database, reloads it, and replaces the old database. If you use this option, there can be no other connections to the database, and the database connection must be local, not over a network.

If you specify *directory*, the transaction log offsets are reset for replication purposes, and the transaction log from the old database is moved to the specified directory. The named directory should be the directory that holds the old transaction logs used by the Message Agent and the Replication Agent. The transaction log management is handled only if the database is used in replication: if there is no SQL Remote publisher or LTM check, then the old transaction log is not needed and is deleted instead of being copied to the specified directory.

☞ For more information on transaction log management, see "Backup methods for remote databases in replication installations" on page 314.

Connection parameters for source database (-c) For a description of the connection parameters, see "Connection parameters" on page 70. The user ID should have DBA authority, to ensure that the user has permissions on all the tables in the database.

For example, the following statement unloads the **asademo** database running on the **sample_server** server, connecting as user ID **DBA** with password **SQL**. The data is unloaded into the *c:\unload* directory.

```
dbunload -c "eng=sample_server;dbn=asademo;uid=DBA;pwd=SQL" c:\unload
```

Unload data only (-d) With this option, none of the database definition commands are generated (CREATE TABLE, CREATE INDEX, and so on); *reload.sql* contains statements to reload the data only.

No data output for listed tables (-e) This option is accessible only when you run this utility at the command prompt. If you wish to unload almost all of the tables in the database, the *-e* option unloads all tables except the specified tables. A *reload.sql* file created with the *-e* option should not be used to rebuild a database because the file will not include all the database tables.

Specify encryption algorithm (-ea) This option allows you to choose which strong encryption algorithm to encrypt your new database with. You can choose either AES or MDSR. Algorithm names are case-insensitive. If you specify the *-ea* option, you must also specify *-ep* or *-ek*. If you specify *-ea* without specifying *-an*, the *-ea* option is ignored.

Specify encryption key (-ek) This option allows you to specify an encryption key for the new database created if you unload and reload a database (using the *-an* option). If you create a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way. If you specify the *-ek* option without specifying an algorithm (using the *-ea* option), the algorithm used will be AES. If you specify *-ek* without specifying *-an*, the *-ek* option is ignored.

Prompt for encryption key (-ep) This option prompts you to specify an encryption key for the new database created if you unload and reload your database (using the *-an* option). It provides an extra measure of security by never allowing the encryption key to be seen in clear text. If you specify the *-ep* option without specifying an algorithm (*-ea* option), the algorithm used will be AES. If you specify *-ep* without specifying *-an*, the *-ep* option is ignored. If you specify *-ep* and *-an*, you must input the encryption key twice to confirm that it was entered correctly. If the keys don't match, the unload fails.

Internal versus external unloads and reloads

The following options offer combinations of internal and external unloads and reloads: `-ii`, `-ix`, `-xi` and `-xx`. A significant performance gain can be realized using internal commands (UNLOAD/LOAD) versus external commands (Interactive SQL's INPUT and OUTPUT statement). However, internal commands are executed by the server so that file and directory paths are relative to the location of the database server. Using external commands, file and directory paths are relative to the current directory of the user.

In Sybase Central, you can specify whether to unload relative to the server or client.

For more information on filenames and paths for the Unload utility, see "UNLOAD TABLE statement" on page 573 of the book *ASA SQL Reference Manual*.

Use internal unload, internal reload (-ii) This option uses the UNLOAD statement to extract data from the database, and uses the LOAD statement in the *reload.sql* file to repopulate the database with data. This is the default.

Use internal unload, external reload (-ix) This option uses the UNLOAD statement to extract data from the database, and uses the Interactive SQL INPUT statement in the *reload.sql* file to repopulate the database with data.

Repeated unload of view creation statements (-j) If your database contains view definitions that are dependent on each other, you can use this option to avoid failure when reloading the views into a database. This option causes view creation statements to be unloaded multiple times, as specified by the count entered. This count should be small, and should correspond to the number of levels of view dependency.

Unload schema definition only (-n) With this option, none of the data in the database is unloaded; *reload.sql* contains SQL statements to build the structure of the database only.

Log output messages to file (-o) Write output messages to the named file.

Escape character (-p) The default escape character (`\`) for external unloads (`dbunload -x` option) can be replaced by another character, using this option. This option is available only when you run this utility from the command prompt.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Specify reload filename (-r) Modify the name and directory of the generated reload Interactive SQL command file. The default is *reload.sql* in the current directory. The directory is relative to the current directory of the client application, not the server.

Unload only listed tables (-t) Provide a list of tables to be unloaded. By default, all tables are unloaded. Together with the *-n* option, this allows you to unload a set of table definitions only.

Output unordered data (-u) Normally, the data in each table is ordered by the primary key. Use this option if you are unloading a database with a corrupt index, so that the corrupt index is not used to order the data.

Enable verbose mode (-v) The table name of the table currently being unloaded, and how many rows have been unloaded, appears. This option is available only when you run this utility from the command prompt.

Use external unloading, internal reload (-xi) This option uses the Interactive SQL OUTPUT statement to extract data from the database, and uses the LOAD statement in the generated reload command file, *reload.sql*, to repopulate the database with data.

Use external unloading, external reload (-xx) This option uses the Interactive SQL OUTPUT statement to extract data from the database, and uses the Interactive SQL INPUT statement in the generated reload command file, *reload.sql*, to repopulate the database with data.

Operate without confirming actions (-y) Choosing this option replaces existing command files without prompting you for confirmation.

Rebuilding a database

To unload a database, start the database server with your database, and run the Unload utility with the DBA user ID and password.

To reload a database, create a new database and then run the generated *reload.sql* command file through Interactive SQL.

In Windows 95/98/Me and NT/2000/XP, and UNIX, there is a file (*rebuild.bat*, *rebuild.cmd*, or *rebuild*) that automates the unload and reload process.

The Upgrade utility

While you can run a database created with an older version of Adaptive Server Anywhere on a newer release of Adaptive Server Anywhere, the new features are only available if the database is upgraded.

The Upgrade utility upgrades a database from an older version of Adaptive Server Anywhere to a newer version of Adaptive Server Anywhere, and allows you to take advantage of the newer release's full list of features.

Upgrading a database does not require you to unload and reload your database.

If you wish to use replication on an upgraded database, you must also archive your transaction log and start a new one on the upgraded database.

You can access the Upgrade utility in the following ways:

- ◆ From Sybase Central, using the Upgrade Database wizard.
- ◆ At the command prompt, using the *dbupgrad* command.

Back up before upgrading

As with any software, it is recommended that you make a backup of your database before upgrading.

Upgrading a database using the Upgrade Database wizard

❖ To upgrade a database:

- 1 Connect to the database.
- 2 Open the Utilities folder in the left pane, and double-click Unload Database in the right pane.

The Upgrade Database wizard appears.

- 3 Follow the instructions in the wizard.

Tip

You can also access the Upgrade Database wizard by:

- ◆ Choosing Tools➤Adaptive Server Anywhere 8➤Upgrade Database.
- ◆ Right-clicking the database, and choosing Upgrade Database from the popup menu.
- ◆ Selecting a database in the left pane, and choosing Upgrade Database from the File menu.

Upgrading databases that are too old for the Upgrade utility

- ❖ **To upgrade a database created with a version of Adaptive Server Anywhere that is too old to be upgraded:**
- 1 Unload the database using the Unload utility.
 - 2 Create a database with the name you wish to use for the upgraded version using the Initialization utility.
 - 3 Connect to the new database from Interactive SQL as the DBA user ID, and read the *reload.sql* command file to build the new database.

You can also use the *dbunload* utility to directly rebuild.

Upgrading a database using the dbupgrad command-line utility

Syntax

dbupgrad [*options*]

Option	Description
-c "keyword=value; ..."	Supply database connection parameters
-I	Do not install Sybase jConnect support
-j	Do not install runtime Java classes
-ja	Add runtime Java classes
-jdk <i>version</i>	Install entries for the named version of the Java Development Kit
-jr	Remove runtime Java classes
-o <i>filename</i>	Log output messages to a file
-q	Quiet mode—no windows or messages

Description

The *dbupgrad* utility upgrades a database created with earlier versions of the software to enable features from the current version of the software. The earliest version that can be upgraded is Watcom SQL 3.2. While later versions of the database server do run against databases were created with earlier releases of the software, some of the features introduced since the version that created the database are unavailable unless the database is upgraded.

Exit codes are 0 (success) or non-zero (failure).

☞ For more information about the Upgrade utility options, see "Upgrade utility options" on page 523.

Not all features made available

Features that require a physical reorganization of the database file are not made available by *dbupgrad*. Such features include index enhancements and changes in data storage. To obtain the benefits of these enhancements, you must unload and reload your database.

☞ For more information, see "Upgrading the database file format" on page 144 of the book *What's New in SQL Anywhere Studio*.

Java in the database

By default, Java in the database is not included in new databases. Java in the database is removed from the database during an upgrade if:

- ◆ the database is Java-enabled, but Java in the database is not being used, and
- ◆ the Java VM cannot be loaded by the database server, and
- ◆ you did not specify any Java-related options to the upgrade.

Upgrade utility options

Connection parameters (-c) For a description of the connection parameters, see "Connection parameters" on page 70. The user ID must have DBA authority.

For example, the following command upgrades a database called sample40 to a version 8.0 format, connecting as user **DBA** with password **SQL**:

```
dbupgrad -c "uid=DBA;pwd=SQL;dbf=c:\wsql40\sample40.db"
```

The *dbupgrad* utility must be run by a user with DBA authority.

Do not install Sybase jConnect support (-i) If you wish to use the Sybase jConnect JDBC driver to access system catalog information, you need to install jConnect support. Use this option if you wish to exclude the jConnect system objects. You can still use JDBC, as long as you do not access system information. If you want, you can add Sybase jConnect support later using Sybase Central.

🔗 For more information, see "Installing jConnect system objects into a database" on page 137 of the book *ASA Programming Guide*.

Do not install runtime Java classes (-j) If you do not intend to use Java classes, you can specify either no Java option, or the -j option to avoid including these classes in the database.

The runtime classes add several megabytes to the size of a database. If you do not intend to use Java classes, you can specify the -j option to avoid including these classes in the upgraded database. If you want, you can add Sybase runtime Java classes later using Sybase Central or ALTER DATABASE.

🔗 For more information, see "Java-enabling a database" on page 89 of the book *ASA Programming Guide*.

Install runtime Java classes (-ja) Add Java in the database to the upgraded database. Specifying -ja installs version 1.3 of the JDK into the database.

Java in the database is a separately-licensable component. These classes are installed during upgrade only if you have installed the Java-in-the-database option or if the Java classes were in use in the database being upgraded.

The runtime classes add several megabytes to the size of a database, so if you do not intend to use Java classes, you can omit the -ja option to avoid installing them.

You can add Sybase runtime Java classes later using the Upgrade Wizard, or the ALTER DATABASE statement.

🔗 For more information, see "Java-enabling a database" on page 89 of the book *ASA Programming Guide*.

Remove runtime Java classes (-jr) Remove Java in the database from the upgraded database.

Install support for the named version of the JDK (-jdk) If you want to install a version of Java, use the -jdk option followed by the version number. Valid values are 1.1.8, and 1.3.


For example, the following command creates a database that supports JDK 1.1.8 applications in the database.

```
dbupgrad -jdk 1.1.8 java2.db
```

The `-jdk` option implies the `-ja` option.

The runtime classes add several megabytes to the size of a database, so if you do not intend to use Java classes, you can omit the `-jdk` option to avoid installing them.

You can add Sybase runtime Java classes later using the Upgrade Wizard, or the ALTER DATABASE statement.

 For more information, see "Java-enabling a database" on page 89 of the book *ASA Programming Guide*.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

The Validation utility

With the Validation utility, you can validate the indexes and keys on some or all of the tables in the database. The Validation utility scans the entire table, and looks up each record in every index and key defined on the table.

This utility can be used in combination with regular backups (see "Backup and Data Recovery" on page 299) to give you confidence in the integrity of the data in your database.

You can access the Validation utility in the following ways:

- ◆ From Sybase Central, using the Validate Database wizard
- ◆ At the command prompt, using the *dbvalid* command. This is useful for incorporating into batch or command files.

 For more information on validating tables, see "VALIDATE TABLE statement" on page 586 of the book *ASA SQL Reference Manual*.

Validating a database using the Validate Database wizard

You can validate a database from Sybase Central, and validate individual tables.

❖ To validate a database:

- 1 Open the Utilities folder.
- 2 Double-click Validate Database.
The Validate Database wizard appears.
- 3 Follow the instructions in the wizard.

❖ To validate a running database:

- 1 Connect to the database.
- 2 Right-click the database and choose Validate Database from the popup menu.
The Validate Database wizard appears.
- 3 Follow the instructions in the wizard.

❖ **To validate an individual table:**

- 1 Connect to the database.
- 2 Locate the table you wish to validate.
- 3 Right-click the table and choose Validate from the popup menu.

Tip

You can also access the Validate Database wizard by

- ◆ Choosing Tools►Adaptive Server Anywhere 8►Validate Database.
- ◆ Right-clicking the database, and choosing Validate Database from the popup menu.
- ◆ Selecting a database in the left pane, and choosing ValidateDatabase from the File menu.

Validating a database using the dbvalid command-line utility

Syntax

dbvalid [*options*] [*object-name*,...]

Option	Description
<i>object-name</i>	The name of a table or (if -i is used) an index to validate
-c " <i>keyword=value; ...</i> "	Supply database connection parameters
-o <i>filename</i>	Log output messages to a file
-f	Validate tables with full check
-fd	Validate tables with data check
-fi	Validate tables with index check
-fx	Validate tables with express check
-i	Each <i>object-name</i> is an index
-q	Quiet mode—do not print messages
-t	Each <i>object-name</i> is a table

Description

With the Validation utility, you can validate the indexes and keys on some or all of the tables in the database. This utility scans the entire table, and confirms that each row in the table exists in the appropriate indexes. It is equivalent to running the VALIDATE TABLE statement on each table.

Exit codes are 0 (success) or non-zero (failure).

🔗 For information on specific checks made during validation, see "VALIDATE TABLE statement" on page 586 of the book *ASA SQL Reference Manual*.

🔗 For more information about the Validation utility options, see "Validation utility options" on page 528.

Validation utility options

Connection parameters (-c) For a description of the connection parameters, see "Connection parameters" on page 70. The user ID must have DBA authority or REMOTE DBA authority.

For example, the following validates the sample database, connecting as user **DBA** with password **SQL**:

```
dbvalid -c "uid=DBA;pwd=SQL;dbf=c:\asa\asademo.db"
```

Full check for each table (-f) In addition to the default validation checks, carry out both data checks (-fd) and index checks (-fi). This corresponds to the **WITH FULL CHECK** option on the **VALIDATE TABLE** statement. Depending on the contents of your database, this option may significantly extend the time required to validate.

Data check for each table (-fd) In addition to the default validation checks, check that all of each **LONG BINARY**, **LONG VARCHAR**, **TEXT**, or **IMAGE** data type can be read. This corresponds to the **WITH DATA CHECK** option on the **VALIDATE TABLE** statement. Depending on the contents of your database, this option may significantly extend the time required to validate.

Index check for each table (-fi) In addition to the default validation checks, validate each index on the table. This corresponds to the **WITH INDEX CHECK** option on the **VALIDATE TABLE** statement. Depending on the contents of your database, this option may significantly extend the time required to validate.

Express check for each table (-fx) In addition to the default and data checks (-fd), check that the number of rows in the table matches the number of entries in the index. This corresponds to the **WITH EXPRESS CHECK** option on the **VALIDATE TABLE STATEMENT**. This option does not perform individual index lookups for each row. Using this option can significantly improve performance when validating large databases with a small cache. This option is only available for databases created with version 8.0 or later of Adaptive Server Anywhere.

Validate specified indexes (-i) Instead of validating tables, validate indexes. In this case, for *dbvalid*, each of the *object-name* values supplied represents an index rather than a table, and has a name of the following form:

[[*owner.*] *table-name.*] *index-name*

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Validate tables (-t) The list of *object-name* values represents a list of tables. This is also the default behavior.

The Write File utility

The Write File utility is used to manage database write files. A **write file** is a file associated with a particular database. All changes are written into the write file, leaving the database file unchanged.

Using write files for development

Write files can be used effectively for testing when you do not wish to modify the production database. They can also be used in network and student environments where read-only access to a database is desired, or when you distribute on CD-ROM a database that you wish users to be able to modify.

Compressed databases

If you are using a compressed database, you must use a write file; compressed database files cannot be accessed directly. The write file name is then used in place of the database name when connecting to the database, or when loading a database on the database server command line.

You can access the Write File utility in the following ways:

- ◆ From Sybase Central, using the Create Write File wizard.
- ◆ At the command prompt, using the *dbwrite* command. This is useful for incorporating into batch or command files.

The Write File utility runs against a database file. The database must not be running on a server when you run the Write File utility.

The database file cannot be modified after the write file is created, otherwise the write file becomes invalid.

Creating a write file using the Create Write File wizard

❖ To create a write file for a database:

- 1 Open the Utilities folder in the left pane.
- 2 Double-click Create Write File in the right pane.
The Create Write File wizard appears.
- 3 Follow the instructions in the wizard.

Tip

You can also access the Create Write File wizard by

- ◆ Choosing Tools►Adaptive Server Anywhere 8►Create Write File.
- ◆ Right-clicking the server, and choosing Create Write File from the popup menu.
- ◆ Selecting a server in the left pane, and choosing Create Write File from the File menu.

Creating a write file using the dbwrite command-line utility

Syntax

dbwrite [*options*] *database-file* [*write-name*]

Option	Description
-c	Create a new write file
-d <i>database-file</i>	Point a write file to a different database
-ek <i>key</i>	Specify encryption key
-ep	Prompt for encryption key
-f <i>database-file</i>	Force the write file to point at a file
-m <i>mirror-name</i>	Set the transaction log mirror name
-o <i>filename</i>	Log output messages to a file
-q	Quiet mode—do not print messages
-s	Report the write file status only
-t <i>log-name</i>	Set the transaction log name
-y	Erase old files without confirmation

Description

If any changes are made to the original database (not using the write file), the write file will no longer be valid. This happens if you start the server using the original database file, so you should create your write file from an archived copy of a database.

If a write file becomes invalid, you can discard all of your changes and create a new write file with the following command.

```
dbwrite -c db-name write-name
```

The *log-name* and *mirror-name* parameters are used only when creating a new write file. The *write-name* parameter is used only with the `-c` and `-d` parameters. Note that the *database_file* parameter must be specified before the *write-name* parameter.

Exit codes are 0 (success) or non-zero (failure).

☞ For more information about the Write file utility options, see "Write file utility options" on page 532.

Write file utility options

Create a new write file (`-c`) If an existing write already exists, any information in the old write file will be lost. If no write file name is specified in the command, the write filename defaults to the database name with the extension *wrt*. If no transaction log name is specified, the log file name will default to the database name with the extension *wlg*.

Change the database file to which an existing write file points (`-d`) If a database file is moved to another directory, or renamed, this option allows you to maintain the link between the write file and the database file. This option is available only when you run this utility from the command prompt.

Specify encryption key (`-ek`) This option allows you to specify the encryption key for strongly encrypted databases directly in the command. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way. For strongly encrypted databases, you must specify either `-ek` or `-ep`, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Prompt for encryption key (`-ep`) This option allows you to specify that you want to be prompted for the encryption key. This option causes a dialog box to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text. For strongly encrypted databases, you must specify either `-ek` or `-ep`, but not both. The command will fail if you do not specify a key for a strongly encrypted database.

Force a write file to point to a file (`-f`) This option is available only when you run this utility from the command prompt. This option is for use when a write file is being created and the database file is held on a Novell NetWare or other network path, for operating systems on which they cannot be entered directly. By providing the full Novell path name for the database file (for example: *SYS:lasademo.db*), you can avoid dependencies on local mappings of the NetWare path. Unlike with the `-d` option, no checking is done on the specified path.

Set mirror name (-m) Set the name of the transaction log mirror file. This can only be used in conjunction with -c.

Log output messages to file (-o) Write output messages to the named file.

Operate quietly (-q) Do not display output messages. This option is available only when you run this utility from the command prompt.

Report the write file status only (-s) This displays the name of the database to which the write file points. This option is available only when you run this utility from the command prompt.

Set the transaction log filename (-t) If you are creating a new write file (-c), you can use this option to set the name of the transaction log file.

Operate without confirming actions (-y) Choosing this option replaces the existing write file without prompting you for confirmation.

C H A P T E R 1 6

Database Options

About this chapter This chapter describes the database, replication, compatibility and Interactive SQL options that can be set to customize and modify database behavior.

Contents	Topic	Page
	Introduction to database options	536
	Database options	541
	Compatibility options	544
	Replication options	547
	Interactive SQL options	548
	Alphabetical list of options	550

Introduction to database options

Database options control many aspects of database behavior. For example, you can use database options for the following purposes:

- ◆ **Compatibility** You can control how much like Adaptive Server Enterprise your Adaptive Server Anywhere database operates, and whether SQL that does not conform to SQL/92 generates errors.
- ◆ **Error handling** You can control what happens when errors such as dividing by zero, or overflow errors, occur.
- ◆ **Concurrency and transactions** You can control the degree of concurrency, and details of COMMIT behavior.

Setting options

You set options with the SET OPTION statement. It has the following general syntax:

```
SET [ EXISTING ] [ TEMPORARY ] OPTION  
    [ userid. | PUBLIC. ]option-name = [ option-value ]
```

Specify a user ID or group name to set the option for that user or group only. Every user belongs to the PUBLIC group. If no user ID or group is specified, the option change is applied to the currently logged on user ID that issued the SET OPTION statement.

For example, the following statement applies an option change to the user DBA, if DBA is the user that issues it:

```
SET OPTION login_mode = mixed
```

The following statement applies a change to the **PUBLIC** user ID, a user group to which all users belong.

```
SET OPTION Public.login_mode = standard
```

Caution

Changing option settings while fetching rows from a cursor is not supported because it can lead to ill-defined behavior. For example, changing the DATE_FORMAT setting while fetching from a cursor would lead to different date formats among the rows in the result set. Do not change option settings while fetching rows.

🔗 For more information, see "SET OPTION statement" on page 539 of the book *ASA SQL Reference Manual*.

Finding option settings

You can obtain a list of option settings, or the values of individual options, in a variety of ways.

Getting a list of option values

- ◆ Current option settings for your connection are available as a subset of **connection properties**. You can list all connection properties using the *sa_conn_properties* system procedure.

```
call sa_conn_properties
```

To order this list, you can call *sa_conn_properties_by_name*.

🔗 For more information, see "sa_conn_properties_by_name system procedure" on page 690 of the book *ASA SQL Reference Manual*, and "sa_conn_properties system procedure" on page 689 of the book *ASA SQL Reference Manual*.

- ◆ In Interactive SQL, the SET statement with no arguments lists the current setting of options.

```
SET
```

- ◆ In Sybase Central, right-click on a database, and choose Options from the popup menu.
- ◆ Use the following query on the SYSOPTIONS system view:

```
SELECT *
FROM SYSOPTIONS
```

This displays all PUBLIC values, and those USER values that have been explicitly set.

Getting individual option values

You can obtain a single setting using the *connection_property* system function. For example, the following statement reports the value of the ANSI_INTEGER_OVERFLOW option:

```
SELECT CONNECTION_PROPERTY ( 'ANSI_INTEGER_OVERFLOW' )
```

🔗 For more information, see "CONNECTION_PROPERTY function" on page 113 of the book *ASA SQL Reference Manual*.

Scope and duration of database options

You can set options at 3 levels of scope: public, user, and temporary.

Temporary options take precedence over user and public settings. User level options take precedence over public settings. If you set a user level option for the current user, the corresponding temporary option is set as well.

Some options (such as COMMIT behavior) are database-wide in scope. Setting these options requires DBA permissions. Other options (such as ISOLATION_LEVEL) can also be applied to just the current connection, and need no special permissions.

Changes to option settings take place at different times, depending on the option. Changing a global option such as RECOVERY_TIME takes place the next time the server is started.

Options that affect the current connection only generally take place immediately. You can change option settings in the middle of a transaction, for example. One exception to this is that *changing options when a cursor is open can lead to unreliable results*. For example, changing DATE_FORMAT may not change the format for the next row when a cursor is opened. Depending on the way the cursor is being retrieved, it may take several rows before the change works its way to the user.

Setting temporary options

Adding the TEMPORARY keyword to the SET OPTION statement changes the duration of the change. Ordinarily an option change is permanent. It does not change until it is explicitly changed using the SET OPTION statement.

When the SET TEMPORARY OPTION statement is executed, the new option value takes effect only for the current connection, and for the duration of the connection.

When the SET TEMPORARY OPTION is used to set a **PUBLIC** option, the change is in place for as long as the database is running. When the database is shut down, temporary options for the **PUBLIC** user ID revert back to their permanent value.

Setting an option for the **PUBLIC** user ID temporarily offers a security advantage. For example, when the LOGIN_MODE option is enabled, the database relies on the login security of the system on which it is running. Enabling it temporarily means that a database relying on the security of a Windows NT/2000/XP domain will not be compromised if the database is shut down and copied to a local machine. In this case, the LOGIN_MODE option will revert to its permanent value, which could be Standard, a mode where integrated logins are not permitted.

Setting public options

DBA authority is required to set an option for the **PUBLIC** user ID.


Changing the value of an option for the **PUBLIC** user ID sets the value of the option for all users who have not SET their own value. An option value cannot be set for an individual user ID unless there is already a **PUBLIC** user ID setting for that option.

Deleting option settings

If *option-value* is omitted, the specified option setting will be deleted from the database. If it was a personal option setting, the value reverts back to the PUBLIC setting. If a TEMPORARY option is deleted, the option setting reverts back to the permanent setting.

For example, the following statement resets the ANSI_INTEGER_OVERFLOW option to its default value:

```
SET OPTION ANSI_INTEGER_OVERFLOW =
```

 For more information, see "SET OPTION statement" on page 539 of the book *ASA SQL Reference Manual*.

Option classification

Adaptive Server Anywhere provides many options. It is convenient to divide them into a few general classes. The classes of options are:

- ◆ "Database options" on page 541
- ◆ "Compatibility options" on page 544
- ◆ "Replication options" on page 547
- ◆ "Interactive SQL options" on page 548

Initial option settings

Connections to Adaptive Server Anywhere can be made through the TDS protocol (Open Client and jConnect JDBC connections) or through the Adaptive Server Anywhere protocol (ODBC and Embedded SQL).

If you have users who use both TDS and the Adaptive Server Anywhere-specific protocol, you can configure their initial settings using stored procedures. As it is shipped, Adaptive Server Anywhere uses this method to set Open Client connections and jConnect connections to reflect default Adaptive Server Enterprise behavior.

The initial settings are controlled using the LOGIN_PROCEDURE option. This option names a stored procedure to run when users connect. The default setting is to use the *sp_login_environment* system procedure. If you wish to change this behavior you can do so.

In turn, *sp_login_environment* checks to see if the connection is being made over TDS. If it is, it calls the *sp_tsq_environment* procedure, which sets several options to new "default" values for the current connection.

☞ For more information, see "LOGIN_PROCEDURE option" on page 578, "sp_login_environment system procedure" on page 721 of the book *ASA SQL Reference Manual*, and "sp_tsql_environment system procedure" on page 726 of the book *ASA SQL Reference Manual*.

Database options

This section lists all database options.

OPTION	VALUES	DEFAULT
"AUDITING option" on page 553	ON, OFF	OFF
"BACKGROUND_PRIORITY option" on page 555	ON, OFF	OFF
"BLOCKING option" on page 556	ON, OFF	ON
"BLOCKING_TIMEOUT option" on page 556	Integer	0
"CHECKPOINT_TIME option" on page 557	Number of minutes	60
"CIS_ROWSET_SIZE option" on page 557	Integer	50
"COOPERATIVE_COMMIT_TIMEOUT option" on page 561	Integer	250
"COOPERATIVE_COMMITS option" on page 561	ON, OFF	ON
"DATE_FORMAT option" on page 562	String	'YYYY-MM-DD'
"DATE_ORDER option" on page 564	'YMD', 'DMY', 'MDY'	'YMD'
"DEFAULT_TIMESTAMP_INCREMENT option" on page 564	Integer	1
"DELAYED_COMMIT_TIMEOUT option" on page 564	Integer	500
"DELAYED_COMMITS option" on page 565	ON, OFF	OFF
"EXTENDED_JOIN_SYNTAX option" on page 567	ON/OFF	ON
"FIRST_DAY_OF_WEEK option" on page 568	1, 2, 3, 4, 5, 6, 7	7
"GLOBAL_DATABASE_ID option" on page 569	Integer	2147483647
"ISOLATION_LEVEL option" on page 571	0, 1, 2, 3	0
"JAVA_HEAP_SIZE option" on page 576	Number of bytes	1 000 000
"JAVA_NAMESPACE_SIZE option" on page 577	Number of bytes	4 000 000

OPTION	VALUES	DEFAULT
"LOGIN_MODE option" on page 577	STANDARD, MIXED, INTEGRATED	STANDARD
"LOGIN_PROCEDURE option" on page 578	String	sp_login_environment
"MAX_CURSOR_COUNT option" on page 581	Integer	50
"MAX_HASH_SIZE option" on page 581	Integer	10
"MAX_PLANS_CACHED option" on page 581	Integer	20
"MAX_STATEMENT_COUNT option" on page 583	Integer >=0	50
"MAX_WORK_TABLE_HASH_SIZE option" on page 582	Integer	20
"MIN_PASSWORD_LENGTH option" on page 583	Integer	0
"MIN_TABLE_SIZE_FOR_HISTOGRAM option" on page 583	Integer	1 000 rows
"ON_CHARSET_CONVERSION_FAILURE option" on page 585	IGNORE, WARNING, ERROR	IGNORE
"OPTIMIZATION_GOAL option" on page 587	First-row or all-rows	first-row
"PINNED_CURSOR_PERCENT_OF_CACHE option" on page 591	0-100	10
"PRECISION option" on page 591	Number of digits	30
"PREFETCH option" on page 592	ON, OFF	ON
"PRESERVE_SOURCE_FORMAT option" on page 593	ON, OFF	ON
	ON, OFF	OFF
"RECOVERY_TIME option" on page 595	Number of minutes	2
"RETURN_DATE_TIME_AS_STRING option" on page 596	OFF	ON, OFF
"RETURN_JAVA_AS_STRING option" on page 597	ON, OFF	OFF
"ROW_COUNTS option" on page 597	ON, OFF	OFF

OPTION	VALUES	DEFAULT
"SCALE option" on page 598	Number of digits	6
"SORT_COLLATION option" on page 598	any valid collation_name or collation_ID	Internal
"SUPPRESS_TDS_DEBUGGING option" on page 601	ON, OFF	OFF
"TDS_EMPTY_STRING_IS_NULL option" on page 601	ON, OFF	OFF
"TIME_FORMAT option" on page 601	String	'HH:NN:SS.SSS'
"TIME_ZONE_ADJUSTMENT option" on page 602	Integer, or Negative integer enclosed in quotation marks, or String representing time in hours and minutes, preceded by + or -	Set by either the client's or the server's time zone, depending on the client's connection type
"TIMESTAMP_FORMAT option" on page 602	String	'YYYY-MM-DD HH:NN:SS.SSS'
"TRUNCATE_DATE_VALUES option" on page 604	ON, OFF	ON (post version 7 databases) OFF (pre version 7 databases)
"TRUNCATE_TIMESTAMP_VALUES option" on page 604	ON, OFF	OFF
"TRUNCATE_WITH_AUTO_COMMIT option" on page 605	ON, OFF	OFF
"USER_ESTIMATES option" on page 606	ENABLED, OVERRIDE-MAGIC, DISABLED	OVERRIDE-MAGIC
"WAIT_FOR_COMMIT option" on page 608	ON, OFF	OFF

Compatibility options

The following options allow Adaptive Server Anywhere behavior to be made compatible with Adaptive Server Enterprise, or to support both old behavior and allow ISO SQL/92 behavior.

For further compatibility with Adaptive Server Enterprise, some of these options can be set for the duration of the current connection using the Transact-SQL SET statement instead of the Adaptive Server Anywhere SET OPTION statement.

☞ For a listing, see "SET statement [T-SQL]" on page 533 of the book *ASA SQL Reference Manual*.

Default settings

The default setting for some of these options differs from the Adaptive Server Enterprise default setting. For compatibility, you should explicitly set the options to ensure compatible behavior.

When a connection is made using the Open Client or JDBC interfaces, some option settings are explicitly set for the current connection to be compatible with Adaptive Server Enterprise. These options are listed in the following table.

☞ For information on how the settings are made, see "System and catalog stored procedures" on page 685 of the book *ASA SQL Reference Manual*.

Options for Open Client and JDBC connection compatibility with Adaptive Server Enterprise

Option	Setting
ALLOW_NULLS_BY_DEFAULT	OFF
ANSI_BLANKS	ON
ANSINULL	OFF
CHAINED	OFF
CONTINUE_AFTER_RAISERROR	ON
DATE_FORMAT	YYYY-MM-DD
DATE_ORDER	MDY
ESCAPE_CHARACTER	OFF
FLOAT_AS_DOUBLE	ON
ISOLATION_LEVEL	1
ON_TSQL_ERROR	CONDITIONAL

QUOTED_IDENTIFIER	OFF
TSQL_HEX_CONSTANT	ON
TSQL_VARIABLES	ON
TIME_FORMAT	HH:NN:SS.SSS
TIMESTAMP_FORMAT	YYYY-MM-DD HH:NN:SS.SSS

Transact-SQL and SQL/92 compatibility options

The following table lists the compatibility options, their allowed values, and their default settings.

OPTION	VALUES	DEFAULT
"ALLOW_NULLS_BY_DEFAULT option" on page 550	ON, OFF	ON
"ANSI_BLANKS option" on page 550	ON, OFF	OFF
"ANSI_CLOSE_CURSORS_ON_ROLLBACK option" on page 551	ON, OFF	OFF
"ANSI_INTEGER_OVERFLOW option" on page 551	ON, OFF	OFF
"ANSI_PERMISSIONS option" on page 551	ON, OFF	ON
"ANSI_UPDATE_CONSTRAINTS option" on page 552	OFF, CURSORS, STRICT	CURSORS
"ANSINULL option" on page 553	ON, OFF	ON
"AUTOMATIC_TIMESTAMP option" on page 554	ON, OFF	OFF
"CHAINED option" on page 557	ON, OFF	ON
"CLOSE_ON_ENDTRANS option" on page 558	ON, OFF	ON
"CONTINUE_AFTER_RAISERROR option" on page 560	ON, OFF	ON
"CONVERSION_ERROR option" on page 560	ON, OFF	ON
"DIVIDE_BY_ZERO_ERROR option" on page 566	ON, OFF	ON
"ESCAPE_CHARACTER option" on page 567	System use	OFF
"FIRE_TRIGGERS option" on page 568	ON, OFF	ON

OPTION	VALUES	DEFAULT
"FLOAT_AS_DOUBLE option" on page 569	ON, OFF	OFF
"NEAREST_CENTURY option" on page 584	Integer	50
"NON_KEYWORDS option" on page 584	comma-separated keywords list	No keywords turned off
"ON_TSQL_ERROR option" on page 587	STOP, CONDITIONAL, CONTINUE	CONDITIONAL
"PERCENT_AS_COMMENT option" on page 590	ON, OFF	ON
"QUERY_PLAN_ON_OPEN option" on page 594	ON, OFF	OFF
"QUOTED_IDENTIFIER option" on page 594	ON, OFF	ON
"RI_TRIGGER_TIME option" on page 597	BEFORE, AFTER	AFTER
"SQL_FLAGGER_ERROR_LEVEL option" on page 599	E, I, F, W	W
"SQL_FLAGGER_WARNING_LEVEL option" on page 599	E, I, F, W	W
"STRING_RTRUNCATION option" on page 600	ON, OFF	OFF
"TSQL_HEX_CONSTANT option" on page 606	ON, OFF	ON
"TSQL_VARIABLES option" on page 606	ON, OFF	OFF

Replication options

The following options are included to provide control over replication behavior. Some replication options are useful for SQL Remote replication, and some are relevant for use with Sybase Replication Server.

OPTION	VALUES	DEFAULT
"BLOB_THRESHOLD option" on page 555	Integer	256
"COMPRESSION option" on page 559	Integer, from -1 to 9	6
"DELETE_OLD_LOGS option" on page 566	ON, OFF	OFF
"QUALIFY_OWNERS option" on page 593	ON, OFF	ON
"QUOTE_ALL_IDENTIFIERS option" on page 594	ON, OFF	OFF
"REPLICATE_ALL option" on page 595	ON, OFF	OFF
"REPLICATION_ERROR option" on page 596	Procedure-name	(no procedure)
"SUBSCRIBE_BY_REMOTE option" on page 600	ON, OFF	ON
"VERIFY_ALL_COLUMNS option" on page 607	ON, OFF	OFF
"VERIFY_THRESHOLD option" on page 608	Integer	1 000

Interactive SQL options

Syntax 1	SET [TEMPORARY] OPTION [<i>userid.</i> PUBLIC.] <i>option-name</i> = [<i>option-value</i>]
Syntax 2	SET PERMANENT
Syntax 3	SET
Parameters	<i>userid:</i> <i>identifier, string or host-variable</i> <i>option-name:</i> <i>identifier, string or host-variable</i> <i>option-value:</i> <i>host-variable (indicator allowed), string, identifier, or number</i>
Description	<p>SET PERMANENT (syntax 2) stores all current Interactive SQL options in the SYSOPTION system table. These settings are automatically established every time Interactive SQL is started for the current user ID.</p> <p>Syntax 3 is used to display all of the current option settings. If any options have temporary settings for Interactive SQL or the database server, these are displayed; otherwise, the permanent option settings are displayed.</p>


OPTION	VALUES	DEFAULT
"AUTO_COMMIT option" on page 554	ON, OFF	OFF
"AUTO_REFETCH option" on page 554	ON, OFF	ON
"BELL option" on page 555	ON, OFF	ON
"COMMAND_DELIMITER option" on page 558	String	' ; '
"COMMIT_ON_EXIT option" on page 559	ON, OFF	ON
"DESCRIBE_JAVA_FORMAT option" on page 566	Varchar, binary	Varchar
"ECHO option" on page 567	ON, OFF	ON
"INPUT_FORMAT option" on page 570	ASCII, DBASE, DBASEII, DBASEIII, EXCEL, FIXED, FOXPRO, LOTUS	ASCII
"ISQL_COMMAND_TIMING option" on page 572	ON, OFF	ON
"ISQL_ESCAPE_CHARACTER option" on page 572	Character	' \ '

OPTION	VALUES	DEFAULT
"ISQL_FIELD_SEPARATOR option" on page 573	String	' , '
"ISQL_LOG option" on page 574	File-name	(empty string)
"ISQL_PLAN option" on page 574	NONE, SHORT, LONG, GRAPHICAL, GraphicalWithStatistics	SHORT
"ISQL_PRINT_RESULT_SET option" on page 575	LAST, ALL, NONE	LAST
"ISQL_QUOTE option" on page 576	String	'
"NULLS option" on page 585	String	'(NULL)'
"ON_ERROR option" on page 586	STOP, CONTINUE, PROMPT, EXIT, NOTIFY_CONTINUE, NOTIFY_STOP, NOTIFY_EXIT	PROMPT
"OUTPUT_FORMAT option" on page 588	ASCII, DBASEII, DBASEIII, EXCEL, FIXED, FOXPRO, HTML, LOTUS, SQL, XML,	ASCII
"OUTPUT_LENGTH option" on page 590	Integer	0
"OUTPUT_NULLS option" on page 590	String	'NULL'
"STATISTICS option" on page 600	0,3,4,5,6	3
"TRUNCATION_LENGTH option" on page 605	Integer	30

Alphabetical list of options

This section lists options alphabetically.

ALLOW_NULLS_BY_DEFAULT option [compatibility]

Function	Controls whether new columns that are created without specifying either NULL or NOT NULL are allowed to contain NULL values.
Allowed values	ON, OFF
Default	ON OFF for Open Client and JDBC connections
Description	<p>The <code>ALLOW_NULLS_BY_DEFAULT</code> option is included for Transact-SQL compatibility.</p> <p> For more information, see "Setting options for Transact-SQL compatibility" on page 395 of the book <i>ASA SQL User's Guide</i>.</p>

ANSI_BLANKS option [compatibility]

Function	Controls behavior when character data is truncated at the client side.
Allowed values	ON, OFF
Default	OFF ON for Open Client and JDBC connections
Description	<p>The <code>ANSI_BLANKS</code> option has no effect unless the database was created with the <code>-b</code> command-line option. It forces a truncation error whenever a value of data type <code>CHAR(N)</code> is read into a C <code>char(M)</code> variable for values of <i>N</i> greater than or equal to <i>M</i>. With <code>ANSI_BLANKS</code> set to OFF, a truncation error occurs only when at least one non-blank character is truncated.</p> <p>For Embedded SQL, if <code>ANSI_BLANKS</code> is ON when you supply a value of data type <code>DT_STRING</code>, you must set the sqlen field to the length of the value, including space for the terminating null character. With <code>ANSI_BLANKS</code> off, the length is determined solely by the position of the null character.</p> <p>The <code>ANSI_BLANKS</code> setting persists for the life of a connection. Changing it after a connection has been established does not affect that connection.</p>

ANSI_CLOSE_CURSORS_ON_ROLLBACK option [compatibility]

Function	Controls whether cursors that were opened WITH HOLD are closed when a ROLLBACK is performed.
Allowed values	ON, OFF
Default	OFF
Description	<p>The draft SQL/3 standard requires all cursors be closed when a transaction is rolled back. By default, on a rollback Adaptive Server Anywhere closes only those cursors that were opened without a WITH HOLD clause. This option allows you to force closure of all cursors.</p> <p>The CLOSE_ON_ENDTRANS option overrides the ANSI_CLOSE_CURSORS_ON_ROLLBACK option.</p>

ANSI_INTEGER_OVERFLOW option [compatibility]

Function	Controls what happens when an arithmetic expression causes an integer overflow error.
Allowed values	ON, OFF
Default	OFF
Description	<p>The ISO SQL/92 standard requires integer overflow should result in a SQLSTATE = 22003 - overflow error. Adaptive Server Anywhere behavior was previously different from this. The ANSI_INTEGER_OVERFLOW option can be set to OFF to maintain compatibility with previous releases of the software.</p>

ANSI_PERMISSIONS option [compatibility]

Function	Controls permissions checking for DELETE and UPDATE statements.
Allowed values	ON, OFF
Scope	Can be set for the PUBLIC group only. Takes effect immediately. DBA authority is required to set this option.
Default	ON
Description	<p>With ANSI_PERMISSIONS ON, the SQL/92 permissions requirements for DELETE and UPDATE statements are checked. The default value is OFF in Adaptive Server Enterprise. The following table outlines the differences.</p>

SQL Statement	Permissions required with ANSI_PERMISSIONS off	Permissions required with ANSI_PERMISSIONS on
UPDATE	UPDATE permission on the columns where values are being set	UPDATE permission on the columns where values are being set SELECT permission on all columns appearing in the WHERE clause SELECT permission on all columns on the right side of the set clause
DELETE	DELETE permission on the table	DELETE permission on the table SELECT permission on all columns appearing in the WHERE clause

The ANSI_PERMISSIONS option can be set only for the PUBLIC group. No private settings are allowed.

ANSI_UPDATE_CONSTRAINTS option [compatibility]

Function	Controls the range of updates that are permitted.
Allowed values	OFF , CURSORS , STRICT
Default	CURSORS in new databases. OFF in databases created before version 7.0.
See also	"UPDATE statement" on page 575 of the book <i>ASA SQL Reference Manual</i>
Description	<p>Adaptive Server Anywhere provides several extensions that allow updates which are not permitted by the ANSI SQL standard. These extensions provide powerful, efficient mechanisms for performing updates. However, in some cases, they cause behavior that is not intuitive. This behavior can produce anomalies such as lost updates if the user application is not designed to expect the behavior of these extensions.</p> <p>The ANSI_UPDATE_CONSTRAINTS option controls whether updates are restricted to those permitted by the SQL/92 standard.</p> <p>If the option is set to STRICT, the following updates are prevented:</p> <ul style="list-style-type: none">◆ Updates of cursors containing JOINS

- ◆ Updates of columns that appear in an ORDER BY clause
- ◆ The FROM clauses is not allowed in UPDATE statements

If the option is set to **CURSORS**, these same restrictions are in place, but only for cursors. If a cursor is not opened with **FOR UPDATE** or **FOR READ ONLY**, the database server chooses updatability based on the SQL/92 standard. If the **ANSI_UPDATE_CONSTRAINTS** option is **CURSORS** or **STRICT**, cursors containing an ORDER BY clause default to **FOR READ ONLY**; otherwise, they continue to default to **FOR UPDATE**.

ANSINULL option [compatibility]

Function	Controls the interpretation of = and != with NULL values.
Allowed values	ON, OFF
Default	ON
Description	<p>With ANSINULL ON, any comparisons with NULL using = or != are unknown.</p> <p>Also, aggregate functions on columns that contain NULL values cause the warning 'null value eliminated in aggregate function' (SQLSTATE=01003).</p> <p>Setting ANSINULL to OFF allows comparisons with NULL to yield results that are not unknown, for compatibility with Adaptive Server Enterprise, and turns off the warning.</p>

AUDITING option [database]

Function	Enables and disables auditing in the database.
Allowed values	ON, OFF
Scope	Can be set for the PUBLIC group only. Takes effect immediately. DBA permissions are required to set this option.
Default	OFF
Description	<p>This option turns auditing on and off.</p> <p>Auditing is the recording of detailed information about many events in the database in the transaction log. Auditing provides some security features, at the cost of some performance.</p>
Example	<ul style="list-style-type: none">◆ Turn on auditing

```
SET OPTION PUBLIC.AUDITING = 'ON'
```

AUTO_COMMIT option [ISQL]

Function	Controls whether a COMMIT is performed after each statement.
Allowed values	ON, OFF
Default	OFF
Description	<p>If AUTO_COMMIT is ON, a database COMMIT is performed after each successful statement. If the COMMIT fails, you have the option to execute additional SQL statements and perform the COMMIT again, or execute a ROLLBACK statement.</p> <p>By default, a COMMIT or ROLLBACK is performed only when the user issues a COMMIT or ROLLBACK statement or a SQL statement that causes an automatic commit (such as the CREATE TABLE statement).</p>

AUTO_REFETCH option [ISQL]

Function	Controls whether query results are fetched again after deletes, updates, and inserts.
Allowed values	ON, OFF
Default	ON
Description	If AUTO_REFETCH is on, the current query results that appear on the Results tab in the Interactive SQL Results pane will be refetched from the database after any INSERT, UPDATE, or DELETE statement. Depending on how complicated the query is, this may take some time. For this reason, it can be turned off.

AUTOMATIC_TIMESTAMP option [compatibility]

Function	Controls interpretation of new columns with the TIMESTAMP data type.
Allowed values	ON, OFF
Default	OFF
Description	<p>Controls whether any new columns with the TIMESTAMP data type that do not have an explicit default value defined are given a default value of the Transact-SQL timestamp value as a default. The AUTOMATIC_TIMESTAMP option is included for Transact-SQL compatibility. The default is OFF.</p>

🔗 For more information, see "Setting options for Transact-SQL compatibility" on page 395 of the book *ASA SQL User's Guide*.

BACKGROUND_PRIORITY option [database]

Function	To limit impact on the performance of connections other than the current connection.
Allowed values	ON, OFF
Scope	<p>Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.</p> <p>If you set this option temporarily, that setting applies to the current connection only. Different connections under the same user ID can have different settings for this option.</p>
Default	OFF
Description	When set to ON , it requests that the current connection have minimal impact on the performance of other connections. This option allows tasks for which responsiveness is critical to coexist with other tasks for which performance is not as important.

BELL option [ISQL]


Function	Controls whether the bell sounds when an error occurs.
Allowed values	ON, OFF
Default	ON
Description	Set this option according to your preference.

BLOB_THRESHOLD option [replication]

Function	Controls the size of value that the Message Agent treats as a long object (blob).
Allowed values	<i>Integer</i> , in kilobytes
Default	256
Description	Any value longer than the BLOB_THRESHOLD option is replicated as a blob. That is, it is broken into pieces and replicated in chunks, before being reconstituted using a SQL variable and concatenating the pieces at the recipient site.

If you set BLOB_THRESHOLD to a high value in remote Adaptive Server Anywhere databases, blobs are not broken into pieces, and operations can be applied to Adaptive Server Enterprise by the Message Agent. Each SQL statement must fit within a message, so this only allows replication of small blobs.

BLOCKING option [database]

Function	Controls the behavior in response to locking conflicts.
Allowed values	ON, OFF
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	ON
Description	<p>If BLOCKING is ON, any transaction attempting to obtain a lock that conflicts with an existing lock held by another transaction waits until every conflicting lock is released. At that time, the write goes through. If BLOCKING is OFF, the transaction that attempts to obtain a conflicting lock receives an error.</p> <p> For more information, see "Two-phase locking" on page 131 of the book <i>ASA SQL User's Guide</i>.</p>


BLOCKING_TIMEOUT option [database]

Function	To control how long a transaction waits to obtain a lock.
Allowed values	<i>Integer</i> , in milliseconds
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	<i>0</i>
See also	"BLOCKING option" on page 556
Description	<p>When BLOCKING is ON, any transaction attempting to obtain a lock that conflicts with an existing lock waits for BLOCKING_TIMEOUT milliseconds for the conflicting lock to be released. If the lock is not released within BLOCKING_TIMEOUT milliseconds, then an error is returned for the waiting transaction.</p> <p>Setting this option to 0 forces all transactions attempting to obtain a lock to wait until all conflicting transactions release their locks.</p>

CHAINED option [compatibility]

Function	Controls transaction mode in the absence of a BEGIN TRANSACTION statement.
Allowed values	ON, OFF
Default	ON OFF for Open Client and JDBC connections
Description	Controls the Transact-SQL transaction mode. In Unchained mode (CHAINED = OFF), each statement is committed individually unless an explicit BEGIN TRANSACTION statement is executed to start a transaction. In chained mode (CHAINED = ON) a transaction is implicitly started before any data retrieval or modification statement.

CHECKPOINT_TIME option [database]

Function	Set the maximum number of minutes that the database server will run without doing a checkpoint.
Allowed values	<i>Integer</i>
Scope	Can be set only for the PUBLIC group. DBA authority is required to set the option. You must shut down and restart the database server for the change to take effect.
Default	<i>60</i>
Description	This option is used with the "RECOVERY_TIME option" on page 595 to decide when checkpoints should be done.  For information on checkpoints, see "Checkpoints and the checkpoint log" on page 322.

CIS_ROWSET_SIZE option [database]

Function	Set the number of rows that are returned from remote servers for each fetch.
Allowed values	<i>Integer</i>
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect when a new connection is made to a remote server.
Default	<i>50</i>

Description	This option sets the ODBC FetchArraySize value when using ODBC to connect to a remote database server.
--------------------	---

CLOSE_ON_ENDTRANS option [compatibility]

Function	Controls the closing of cursors at the end of a transaction.
-----------------	--

Allowed values	ON, OFF
-----------------------	----------------

Default	ON
----------------	-----------

Description	When CLOSE_ON_ENDTRANS is set to ON , cursors are closed whenever a transaction is committed unless the cursor was opened WITH HOLD. The behavior when a transaction is rolled back is governed by the ANSI_CLOSE_CURSORS_AT_ROLLBACK option.
--------------------	--

When CLOSE_ON_ENDTRANS is set to **OFF**, cursors are not closed at either a commit or a rollback, regardless of the ANSI_CLOSE_CURSORS_AT_ROLLBACK option setting or whether the cursor was opened WITH HOLD or not.

Setting this to **OFF** provides Adaptive Server Enterprise compatible behavior.

COMMAND_DELIMITER option [ISQL]

Function	Sets the string that indicates the end of a statement in Interactive SQL.
-----------------	---

Allowed values	<i>String</i>
-----------------------	---------------

Default	Semi-colon (;)
----------------	----------------

See also	"The Interactive SQL utility" on page 472
-----------------	---

Description	You can specify a command delimiter to indicate the end of a SQL statement. This command delimiter can be any sequence of characters (including numbers, letters, and punctuation), but it cannot contain embedded blanks. As well, it can contain a semicolon, but only as the first character.
--------------------	--

If the command delimiter is set to a string beginning with a character that is valid in identifiers, the command delimiter must be preceded by a space. This command delimiter is case sensitive. You must enclose the new command delimiter in single quotation marks.

Note that if the command delimiter is a semicolon (the default), no space is required before the semicolon.

Example 1

```
SET OPTION COMMAND_DELIMITER='~'
```

Example 2

```
set temporary option command_delimiter = ' select';
message 'hello' select
```

Example 3

You can also use Interactive SQL's `-d` command line option to set the command delimiter without including a `SET TEMPORARY OPTION COMMAND_DELIMITER` statement in a .SQL file. For example, if you have a script file named `test.sql` which uses tildes (`~`) as command delimiters, you could run:

```
dbisql -d "~" test.sql
```

COMMIT_ON_EXIT option [ISQL]

Function	Controls behavior when Interactive SQL disconnects or terminates.
Allowed values	ON, OFF
Default	ON
Description	Controls whether a COMMIT or ROLLBACK is done when you leave Interactive SQL. When COMMIT_ON_EXIT is set to ON , a COMMIT is done; otherwise a ROLLBACK is done.

COMPRESSION option [replication]

Function	Set the level of compression for SQL Remote messages.
Allowed values	<i>Integer</i> , from -1 to 9
Default	6
Description	<p>The values have the following meanings:</p> <ul style="list-style-type: none"> ◆ -1 Send messages in Version 5 format. Message Agents (both <i>dbremote</i> and <i>ssremote</i>) from previous versions of SQL Remote cannot read messages sent in Version 6 format. You should ensure that COMPRESSION is set to -1 until all Message Agents in your system are upgraded to Version 6. ◆ 0 No compression.

- ◆ **1 to 9** Increasing degrees of compression. Creating messages with high compression can take longer than creating messages with low compression.

CONTINUE_AFTER_RAISERROR option [compatibility]

Function	Controls behavior following a RAISERROR statement.
Allowed values	ON, OFF
Default	OFF for databases created with Version 5.x. ON for databases created with Version 6 or later.
See also	"ON_TSQL_ERROR option" on page 587
Description	<p>The RAISERROR statement is used within procedures and triggers to generate an error. When this option is set to OFF, the execution of the procedure or trigger is stopped whenever the RAISERROR statement is encountered.</p> <p>If you set the CONTINUE_AFTER_RAISERROR option to ON, the RAISERROR statement no longer signals an execution-ending error. Instead, the RAISERROR status code and message are stored and the most recent RAISERROR is returned when the procedure completes. If the procedure that caused the RAISERROR was called from another procedure, the RAISERROR is not returned until the outermost calling procedure terminates.</p> <p>Intermediate RAISERROR statuses and codes are lost after the procedure terminates. If, at return time, an error occurs along with the RAISERROR, then the information for the new error is returned and the RAISERROR information is lost. The application can query intermediate RAISERROR statuses by examining the @@error global variable at different execution points.</p> <p>The setting of the CONTINUE_AFTER_RAISERROR option is used to control behavior following a RAISERROR statement <i>only</i> if the ON_TSQL_ERROR option is set to CONDITIONAL (the default). If you set the ON_TSQL_ERROR option to STOP or CONTINUE, the ON_TSQL_ERROR setting takes precedence over the CONTINUE_AFTER_RAISERROR setting.</p>

CONVERSION_ERROR option [compatibility]

Function	Controls the reporting of data type conversion failures on fetching information from the database.
----------	--

Allowed values	ON, OFF
Default	ON
Description	<p>This option controls whether data type conversion failures, when data is fetched from the database or inserted into the database, are reported by the database as errors (CONVERSION_ERROR set to ON) or as a warning (CONVERSION_ERROR set to OFF).</p> <p>When CONVERSION_ERROR is set to ON, the SQL_CONVERSION_ERROR error is generated. If the option is set to OFF, the warning SQL_CANNOT_CONVERT is produced.</p> <p>If conversion errors are reported as warnings only, the NULL value is used in place of the value that could not be converted. In Embedded SQL, an indicator variable is set to -2 for the column or columns that cause the error.</p>

COOPERATIVE_COMMIT_TIMEOUT option [database]

Function	Governs when a COMMIT entry in the transaction log is written to disk.
Allowed values	<i>Integer</i> , in milliseconds
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	<i>250</i>
Description	This option has meaning only when COOPERATIVE_COMMITS is set to ON . The database server waits for the specified number of milliseconds for other connections to fill a page of the log before writing to disk. The default setting is 250 milliseconds.

COOPERATIVE_COMMITS option [database]

Function	Controls when commits are written to disk.
Allowed values	ON or OFF
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	ON
Description	If COOPERATIVE_COMMITS is set to OFF , a COMMIT is written to disk as soon as the database server receives it, and the application is then allowed to continue.

If COOPERATIVE_COMMITS is set to **ON** (the default) and if there are other active connections, the database server does not immediately write the COMMIT to the disk. Instead, the application waits for up to the maximum length set by the COOPERATIVE_COMMIT_TIMEOUT option for something else to put on the pages before they are written to disk.

Setting COOPERATIVE_COMMITS to **ON**, and increasing the COOPERATIVE_COMMIT_TIMEOUT setting, increases overall database server throughput by cutting down the number of disk I/Os, but at the expense of a longer turnaround time for each individual connection.

If both COOPERATIVE_COMMITS and DELAYED_COMMITS are set to **ON**, and the COOPERATIVE_COMMIT_TIMEOUT interval passes without the pages getting written, the application is resumed (as if the commit had worked), and the remaining interval (DELAYED_COMMIT_TIMEOUT - COOPERATIVE_COMMIT_TIMEOUT) is used as a DELAYED_COMMIT interval. The pages will then be written, even if they are not full.

DATE_FORMAT option [compatibility]

Function	Sets the format for dates retrieved from the database.
Allowed values	String
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	'YYYY-MM-DD'. This corresponds to ISO date format specifications.
Description	The format is a string using the following symbols:

Symbol	Description
yy	Two digit year
yyyy	Four digit year
mm	Two digit month, or two digit minutes if following a colon(as in hh:mm)
mmm[m...]	Character short form for months—as many characters as there are "m"s
d	Single digit day of week, (0 = Sunday, 6 = Saturday)
dd	Two digit day of month
ddd[d...]	Character short form for day of the week
hh	Two digit hours

Symbol	Description
nn	Two digit minutes
ss[.ss..]	Seconds and parts of a second
aa	AM or PM (12 hour clock)
pp	PM if needed (12 hour clock)
jjj	Day of the year, from 1 to 366

Each symbol is substituted with the appropriate data for the date that is being formatted.

For symbols that represent character data (such as `mmm`), you can control the case of the output as follows:

- ◆ Type the symbol in all upper case to have the format appear in all upper case. For example, `MMM` produces `JAN`.
- ◆ Type the symbol in all lower case to have the format appear in all lower case. For example, `mmm` produces `jan`.
- ◆ Type the symbol in mixed case to have Adaptive Server Anywhere choose the appropriate case for the language that is being used. For example, in English, typing `Mmm` produces `May`, while in French it produces `mai`.

For symbols that represent numeric data, you can control zero-padding with the case of the symbols:

- ◆ Type the symbol in same-case (such as `MM` or `mm`) to allow zero padding. For example, `yyyy/mm/dd` could produce `2002/01/01`.
- ◆ Type the symbol in mixed case (such as `Mm`) to suppress zero padding. For example, `yyyy/Mm/Dd` could produce `2002/1/1`.
- ◆ The following table illustrates `DATE_FORMAT` settings, together with the output from the following statement, executed on Thursday May 21, 1998.

```
SELECT CURRENT DATE
```

DATE_FORMAT	SELECT CURRENT DATE
yyyy/mm/dd/ddd	2001/05/21/thu
jjj	141
mmm yyyy	May 1998
mm-yyyy	05-1998

Examples

DATE_ORDER option [compatibility]


Function	Controls the interpretation of date formats.
Allowed values	'MDY', 'YMD', or 'DMY'
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	YMD . This corresponds to ISO date format specifications. For Open Client and JDBC connections, the default is set to MDY .
Description	The database option DATE_ORDER is used to determine whether 10/11/12 is Oct 11 1912, Nov 12 1910, or Nov 10 1912. The option can have the value 'MDY', 'YMD', or 'DMY'.

DEFAULT_TIMESTAMP_INCREMENT option [database]

Function	Specifies the number of microseconds to add to a column of type TIMESTAMP in order to keep values in the column unique.
Allowed values	<i>Integer</i> , between 1 and 60 000 000 inclusive.
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	<i>1</i>
See also	"TRUNCATE_TIMESTAMP_VALUES option" on page 604
Description	<p>Since a TIMESTAMP value is precise to six decimal places in Adaptive Server Anywhere, by default 1 microsecond (0.000001 of a second) is added to differentiate between two identical TIMESTAMP values.</p> <p>Some software, such as Microsoft Access, truncates TIMESTAMP values to three decimal places, making valid comparisons a problem. You can set the TRUNCATE_TIMESTAMP_VALUES option to ON to specify the number of decimal place values Adaptive Server Anywhere stores to maintain compatibility.</p>

DELAYED_COMMIT_TIMEOUT option [database]

Function	Determines when the server returns control to an application following a COMMIT.
Allowed values	<i>Integer</i> , in milliseconds.

Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	500
Description	<p>This option has meaning only when DELAYED_COMMITS is set to ON; it governs when a COMMIT entry in the transaction log is written to disk. With DELAYED_COMMITS set to ON, the database engine waits for the number of milliseconds set in the DELAYED_COMMIT_TIMEOUT option for other connections to fill a page of the log before writing the current page contents to disk.</p> <p> For more information, see "DELAYED_COMMITS option" on page 565.</p>

DELAYED_COMMITS option [database]

Function	Determines when the server returns control to an application following a COMMIT.
Allowed values	ON, OFF
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	OFF . This corresponds to ISO COMMIT behavior.
Description	<p>When set to ON, the database server replies to a COMMIT statement immediately instead of waiting until the transaction log entry for the COMMIT has been written to disk. When set to OFF, the application must wait until the COMMIT is written to disk.</p> <p>When this option is ON, the log is written to disk when the log page is full or according to the DELAYED_COMMIT_TIMEOUT option setting, whichever is first. There is a slight chance that a transaction may be lost even though committed if a system failure occurs after the server replies to a COMMIT, but before the page is written to disk. Setting DELAYED_COMMITS to ON, and the DELAYED_COMMIT_TIMEOUT option to a high value, promotes a quick response time at the cost of security.</p> <p>If both COOPERATIVE_COMMITS and DELAYED_COMMITS are set to ON, and if the COOPERATIVE_COMMIT_TIMEOUT interval passes without the pages getting written, the application is resumed (as if the commit had worked), and the remaining interval (DELAYED_COMMIT_TIMEOUT - COOPERATIVE_COMMIT_TIMEOUT) is used as a DELAYED_COMMIT interval after which the pages will be written, even if they are not full.</p>

DELETE_OLD_LOGS option [replication]

Function	Controls whether transaction logs are deleted when their messages have been replicated.
Allowed values	ON, OFF, DELAY
Default	OFF
Description	<p>This option is used at MobiLink clients, by SQL Remote, and by the Adaptive Server Anywhere Replication Agent. The default setting is OFF. When it is set to ON, each old transaction log is deleted when all the changes it contains have been sent and confirmed as received. When it is set to DELAY, each old transaction log with a file name indicating that it was created on the current day is not deleted.</p> <p>For more information about how to use the <code>DELETE_OLD_LOGS</code> option in conjunction with the <code>Backup</code> statement to delete old copies of transaction logs, see the "BACKUP statement" on page 245 of the book <i>ASA SQL Reference Manual</i>.</p>

DESCRIBE_JAVA_FORMAT option [ISQL]

Function	Controls whether Java objects are interpreted as strings (for display) or as binary (for loading and unloading).
Allowed values	varchar, binary
Default	Varchar
Description	<p>When set to varchar, Interactive SQL converts all data fetched from the database to strings. The database server calls the <code>toString()</code> method on Java columns to provide formatted output for the Results tab in the Results pane.</p> <p>When set to binary, Interactive SQL does no data conversion.</p>

DIVIDE_BY_ZERO_ERROR option [compatibility]

Function	Controls the reporting of division by zero.
Allowed values	ON, OFF
Default	ON
Description	This option indicates whether division by zero is reported as an error. If the option is set ON , then division by zero results in an error with SQLSTATE 22012.

If the option is set **OFF**, division by zero is not an error. Instead, a **NULL** is returned.

ECHO option [ISQL]

Function	Controls whether statements are echoed before they are executed.
Allowed values	ON, OFF
Default	ON
Description	This option is most useful when you use the READ statement to execute a Interactive SQL command file.

ESCAPE_CHARACTER option [compatibility]

This option is reserved for system use. Do not change the setting of this option.

EXTENDED_JOIN_SYNTAX option [database]

Function	Controls whether queries with duplicate correlation names syntax for multi-table joins are allowed, or reported as an error.
Allowed values	ON, OFF
Default	ON
See also	" REWRITE function" on page 172 of the book <i>ASA SQL Reference Manual</i>
Description	If this option is set to ON then Adaptive Server Anywhere allows duplicate correlation names to be used in the null-supplying side of the outer joins. The tables or views specified with the same correlation name are interpreted as the same instance of the table or view.

The following **FROM** clause illustrates the Adaptive Server Anywhere interpretation of a join using duplicate correlation names:

```
( R left outer join T on (C1), T join S on ( C2 ) )
```

where **C1** and **C2** are search conditions.

If the option is set to **ON**, this join is interpreted as follows:

```
( R left outer join T on ( C1 ) ) join S on ( C2 )
```

If the option is set to **OFF**, the following error is generated:

ASA Error -137: Table 'T' requires a unique correlation name.

Note: To see the result of eliminating duplicate correlation names, you can view the rewritten statement using the REWRITE function with the second argument set to ANSI.

FIRE_TRIGGERS option [compatibility]

Function	Controls whether triggers are fired in the database.
Allowed values	ON, OFF
Default	ON
Description	<p>When set to ON, triggers are fired. When set to OFF, no triggers are fired, including referential integrity triggers (such as cascading updates and deletes). Only a user with DBA authority can set this option. The option is overridden by the <code>-gf</code> command-line option, which turns off all trigger firing regardless of the FIRE_TRIGGERS setting.</p> <p>This option is relevant when replicating data from Adaptive Server Enterprise to Adaptive Server Anywhere because all actions from Adaptive Server Enterprise transaction logs are replicated to Adaptive Server Anywhere, including actions carried out by triggers.</p>

FIRST_DAY_OF_WEEK option [database]

Function	Sets the numbering of the days of the week.
Allowed values	1 2 3 4 5 6 7
Default	7 (Sunday is the first day of the week)
Description	The values have the following meaning:

Value	Meaning
1	Monday
2	Tuesday
3	Wednesday
4	Thursday
5	Friday
6	Saturday
7	Sunday

FLOAT_AS_DOUBLE option [compatibility]

Function	Controls the interpretation of the FLOAT keyword.
Allowed values	ON, OFF
Default	OFF ON for Open Client and JDBC connections
Description	<p>The <code>FLOAT_AS_DOUBLE</code> option makes the <code>FLOAT</code> keyword behave like Adaptive Server Enterprise's <code>FLOAT</code> keyword when a precision is not specified.</p> <p>When enabled (set to ON), all occurrences of the keyword <code>FLOAT</code> are interpreted as equivalent to the keyword <code>DOUBLE</code> within SQL statements.</p> <p>By default, Adaptive Server Anywhere <code>FLOAT</code> values are interpreted by Adaptive Server Enterprise as <code>REAL</code> values. Since Adaptive Server Enterprise treats its own <code>FLOAT</code> values as <code>DOUBLE</code>, enabling this option makes Adaptive Server Anywhere treat <code>FLOAT</code> values in the same way Enterprise treats <code>FLOAT</code> values.</p> <p><code>REAL</code> values are four bytes; <code>DOUBLE</code> values are eight bytes. According to the ANSI SQL/92 specification, <code>FLOAT</code> can be interpreted based on the platform. It is up to the database to decide what size the value is, so long as it can handle the necessary precision. Adaptive Server Enterprise and Adaptive Server Anywhere exhibit different default behavior.</p> <p>The <code>FLOAT_AS_DOUBLE</code> option takes effect only when no precision is specified. For example, the following statement is not affected by the option setting:</p>

```
create table t1(  
    c1 float(5)  
)
```

The following statement is affected by the option setting:

```
create table t2(  
    c1 float)  
// affected by option setting
```

GLOBAL_DATABASE_ID option [database]

Function	For use in generating unique primary keys in a replication environment. Controls the beginning of the range of values for columns created with <code>DEFAULT GLOBAL AUTOINCREMENT</code> .
Allowed values	<i>Integer</i>

Default	2147483647
Scope	PUBLIC setting only.
See also	"CREATE TABLE statement" on page 350 of the book <i>ASA SQL Reference Manual</i> "Database-level properties" on page 630 "Ensuring unique primary keys" on page 129 of the book <i>SQL Remote User's Guide</i>
Description	<p>If you create a column with DEFAULT GLOBAL AUTOINCREMENT, its value is incremented. The initial value is determined by the GLOBAL_DATABASE_ID value and the partition size. For more information on the partition size, see "CREATE TABLE statement" on page 350 of the book <i>ASA SQL Reference Manual</i>.</p> <p>Setting GLOBAL_DATABASE_ID to the default value indicates that GLOBAL DEFAULT AUTOINCREMENT is disabled. In this case NULL is generated as a default.</p> <p>You can find the value of the option in the current database using the following statement:</p> <pre>select db_property('GlobalDBId')</pre> <p>This feature is of particular use in replication environments to ensure unique primary keys.</p> <p>☞ For more information, see "Ensuring unique primary keys" on page 129 of the book <i>SQL Remote User's Guide</i> and "Maintaining unique primary keys using global autoincrement" on page 96 of the book <i>MobiLink Synchronization User's Guide</i>.</p>

INPUT_FORMAT option [ISQL]

Function	Sets the default data format expected by the INPUT statement.
Allowed values	<i>String</i> . See below.
Default	<i>ASCII</i>
Description	<p>Certain file formats contain information about column names and types. Using this information, the INPUT statement will create the database table if it does not already exist. This is a very easy way to load data into the database. The formats that have enough information to create the table are: DBASEII, DBASEIII, FOXPRO, and LOTUS.</p> <p>Allowable input formats are:</p>


- ◆ **ASCII** Input lines are assumed to be ASCII characters, one row per line, with values separated by commas. Alphabetic strings may be enclosed in apostrophes (single quotes) or quotation marks (double quotes). Strings containing commas must be enclosed in either single or double quotes. If single or double quotes are used, double the quote character to use it within the string. Optionally, you can use the **DELIMITED BY** clause to specify a different delimiter string than the default, which is a comma (.).

Three other special sequences are also recognized. The two characters `\n` represent a newline character, `\\` represents a single backslash character, and the sequence `\xDD` represents the character with hexadecimal code `DD`.

- ◆ **DBASE** The file is in dBASE II or dBASE III format. Interactive SQL will attempt to determine which format, based on information in the file. If the table doesn't exist, it will be created.
- ◆ **DBASEII** The file is in dBASE II format. If the table doesn't exist, it will be created.
- ◆ **DBASEIII** The file is in dBASE III format. If the table doesn't exist, it will be created.
- ◆ **EXCEL** Input file is in the format of Microsoft Excel 2.1. If the table doesn't exist, it will be created.
- ◆ **FIXED** Input lines are in fixed format. The width of the columns can be specified using the **COLUMN WIDTHS** clause. If they are not specified, column widths in the file must be the same as the maximum number of characters required by any value of the corresponding database column's type.
- ◆ **FOXPRO** The file is in FoxPro format. The FoxPro memo field is different than the dBASE memo field. If the table doesn't exist, it will be created.
- ◆ **LOTUS** The file is a Lotus WKS format worksheet. **INPUT** assumes that the first row in the Lotus WKS format worksheet consists of column names. If the table doesn't exist, it will be created. In this case, the types and sizes of the columns created may not be correct because the information in the file pertains to a cell, not to a column.

ISOLATION_LEVEL option [compatibility]

Function	Controls the locking isolation level.
Allowed values	0, 1, 2, or 3

Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	<i>0</i> <i>1</i> for Open Client and JDBC connections
Description	<p>This option controls the locking isolation level as follows.</p> <ul style="list-style-type: none">◆ 0 Allow dirty reads, non-repeatable reads, and phantom rows.◆ 1 Prevent dirty reads. Allow non-repeatable reads and phantom rows.◆ 2 Prevent dirty reads and guarantee repeatable reads. Allow phantom rows.◆ 3 Serializable. Do not allow dirty reads, guarantee repeatable reads, and do not allow phantom rows. <p> For more information, see "Isolation levels and consistency" on page 94 of the book <i>ASA SQL User's Guide</i>.</p>

ISQL_COMMAND_TIMING option [ISQL]

Function	Controls whether SQL statements are timed or not.
Allowed values	ON, OFF
Default	ON
Description	<p>This boolean option controls whether SQL statements are timed or not. If you set the option to ON, YES, or 1, the time of execution appears in the Messages pane after you execute a statement. If you set the option to OFF, NO, or 0, the time does not appear.</p> <p>You can also set this option on the Messages tab of the Options dialog.</p>

ISQL_ESCAPE_CHARACTER option [ISQL]

Function	Controls the escape character used in place of unprintable characters in data exported to ASCII files.
Allowed values	Any single character
Default	A backslash (\)

Description When Interactive SQL exports strings that contain unprintable characters (such as a carriage return), it converts each unprintable character into a hexadecimal format and precedes it with an escape character. The character you specify for this setting is used in the output if your OUTPUT statement does not contain an ESCAPE CHARACTER clause. This setting is used only if you are exporting to an ASCII file.

Example ♦ Create a table that contains one string value with an embedded carriage return (denoted by the "\n" in the INSERT statement). Then export the data to *c:\escape.txt* with a # sign as the escape character.

```
CREATE TABLE escape_test( TEXT varchar(10 ) );
INSERT INTO escape_test VALUES( 'one\ntwo' );
SET TEMPORARY OPTION ISQL_ESCAPE_CHARACTER='#';
SELECT * FROM escape_test;
OUTPUT TO c:\escape.txt FORMAT ASCII
```

This code places the following data in *escape.txt*:

```
'one#x0Atwo'
```

where # is the escape character and *x0A* is the hexadecimal equivalent of the "\n" character.

The start and end characters (in this case, single quotation marks) depend on the ISQL_QUOTE setting.

ISQL_FIELD_SEPARATOR option [ISQL]

Function Controls the default string used for separating values in data exported to ASCII files.

Allowed values *String*

Default A comma (,)

Description Controls the default string used for separating (or delimiting) values in data exported to ASCII files. If an OUTPUT statement does not contain a DELIMITED BY clause, the value of this setting is used.

Example ♦ Set the field separator to a colon in the data exported to *c:\employee.txt*.

```
SET TEMPORARY OPTION ISQL_FIELD_SEPARATOR=': ';
SELECT emp_lname, emp_fname FROM employee WHERE
emp_id < 150;
OUTPUT TO c:\employee.txt FORMAT ASCII
```

This code places the following data in *employee.txt*:

```
'Whitney': 'Fran'  
'Cobb': 'Matthew'  
'Chin': 'Philip'  
'Jordan': 'Julie'
```

The start and end characters (in this case, single quotation marks) depend on the ISQL_QUOTE setting.

ISQL_LOG option [ISQL]

Function	Controls logging behavior.
Allowed values	<i>String</i> , containing a file name.
Default	<i>Empty string</i> .
Description	If ISQL_LOG is set to a non-empty string, all Interactive SQL statements are added to the end of the named file. Otherwise, if ISQL_LOG is set to the empty string, Interactive SQL statements are not logged.

Individual session only
This option logs an individual Interactive SQL session only. See "Backup and Data Recovery" on page 299 for a description of the transaction log that logs all changes to the database by all users.

ISQL_PLAN option [ISQL]

Function	Controls the amount of information that appears on the Plan tab and the UltraLite Plan tab in the Results pane after you execute statements.
Allowed values	SHORT, LONG, GRAPHICAL, GRAPHICALWITHSTATISTICS, NONE
Default	SHORT
Description	After you execute a SQL statement in Interactive SQL, the optimizer provides information about how it optimized the execution on the Plan tab. To set the level of detail for the plan the optimizer provides, you can choose one of the following values: SHORT provides basic information about the optimizer plan. LONG provides detailed information about the optimizer plan.

GRAPHICAL provides a tree diagram of the query. You can click on a node in the plan diagram to see details about that part of the query.

GRAPHICALWITHSTATISTICS provides a tree diagram of the query, as well as statistics which indicate the resources used by the part of the query that is selected. The UltraLite plan treats this value the same as the **GRAPHICAL** value.

NONE means no optimizer information can be accessed. This parameter is deprecated.

The plan is computed only when you click the Plan tab.

You can also specify the plan type on the Plan tab of the Options dialog.

ISQL_PRINT_RESULT_SET option [ISQL]

Function	<p>This option specifies which result set(s) are printed when a .SQL file is run.</p> <p>The ISQL_PRINT_RESULT_SET option takes effect only when you run DBISQL within a command window (for example, using the -nogui option).</p>
Allowed values	LAST, ALL, NONE
Default	LAST
Description	<p>This option allows you to specify which result set(s) are printed when a .SQL file is run. It is remembered on a per-machine basis (not per-user), and is case-insensitive.</p> <p>You can choose one of the following print options:</p> <p>LAST prints the result set from the last statement in the file.</p> <p>ALL prints result sets from each statement in the file which returns a result set.</p> <p>NONE does not print any result sets.</p> <p>Although this option has no affect when Interactive SQL is running in windowed mode, you can still view and set the option in windowed mode. Choose Tools►Options, then select the appropriate action in the Console Mode box on the Results page. Remember to choose the Make Permanent option if you want to set this option for all Interactive SQL sessions. Otherwise the change will be discarded when you exit Interactive SQL.</p>

ISQL_QUOTE option [ISQL]

Function	Controls the default string that begins and ends all strings in data exported to ASCII files.
Allowed values	<i>String</i>
Default	A single apostrophe (')
Description	Controls the default string that begins and ends all strings in data exported to ASCII files. If an OUTPUT statement does not contain a QUOTE clause, this value is used by default.
Example	<p>◆ To change the default string that begins and ends all strings to a double quote character.</p> <pre>SET TEMPORARY OPTION ISQL_QUOTE=' " ' ; SELECT emp_lname, emp_fname FROM employee WHERE emp_id < 150; OUTPUT TO c:\employee.txt FORMAT ASCII</pre> <p>This code places the following data in <i>employee.txt</i>:</p> <pre>"Whitney", "Fran" "Cobb", "Matthew" "Chin", "Philip" "Jordan", "Julie"</pre> <p>The separator characters (in this case, commas) depend on the ISQL_FIELD_SEPARATOR setting.</p>

JAVA_HEAP_SIZE option [database]

Function	To limit the memory used by Java applications for a connection.
Allowed values	<i>Integer</i>
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately. DBA permissions are required to set this option for <i>any</i> connection.
Default	<i>1 000 000</i>
Description	This option sets the maximum size (in bytes) of the memory that is allocated to Java applications on a per-connection basis. Per-connection memory allocations typically consist of the user's working set of allocated Java variables and Java application stack space.

While a Java application is executing on a connection, the per-connection allocations come out of the fixed cache of the database server, so it is important that a run-away Java application is prevented from using up too much memory.

JAVA_INPUT_OUTPUT option [database]

Function	To enable file access from Java in the database
Allowed values	ON, OFF Supported on NT only.
Scope	DBA authority required.
Default	OFF
See also	"Security management for Java" on page 115 of the book <i>ASA Programming Guide</i>
Description	By default, the java.io package is only partially supported. In particular, classes that handle file access are disabled. If JAVA_INPUT_OUTPUT is set to ON , Java file access classes in java.io are enabled.

JAVA_NAMESPACE_SIZE option [database]

Function	To limit the memory used by Java applications for a connection.
Allowed values	<i>Integer</i>
Scope	Can be set only for the PUBLIC group. Takes effect immediately.
Default	<i>4 000 000</i>
Description	<p>This option sets the maximum size (in bytes) of the memory that is allocated to Java applications on a per-database basis.</p> <p>Per-database memory allocations include Java class definitions. Because class definitions are effectively read-only, they are shared between connections. Consequently, their allocations come right out of the fixed cache, and this option sets a limit on the size of these allocations.</p>


LOGIN_MODE option [database]

Function	Controls the use of integrated logins for the database.
Allowed values	Standard, Mixed, or Integrated

Scope	Can be set only for the PUBLIC group. Takes effect immediately.
Default	Standard
Description	<p>This option specifies whether integrated logins are permitted. The following values are accepted (the values are case insensitive):</p> <ul style="list-style-type: none">◆ Standard This is the default setting, which does not permit integrated logins. An error occurs if an integrated login connection is attempted.◆ Mixed With this setting, both integrated logins and standard logins are allowed.◆ Integrated With this setting, all logins to the database must be made using integrated logins.

Caution

Setting the LOGIN_MODE database option to Integrated restricts connections to only those users who have been granted an integrated login mapping. Attempting to connect with a user ID and password generates an error. The only exceptions to this are users with DBA authority (full administrative rights).

 For more information on integrated logins see "Using integrated logins" on page 83.

LOGIN_PROCEDURE option [database]

Function	A login procedure that sets connection compatibility options at startup. By default the procedure calls the <i>sp_login_environment</i> procedure to determine which options to set.
Allowed values	<i>String</i>
Scope	DBA authority required.
Default	<i>sp_login_environment</i>
Description	<p>This login procedure calls the <i>sp_login_environment</i> procedure at run time to determine the database connection settings.</p> <p>You can customize the default database option settings by creating a new procedure and setting LOGIN_PROCEDURE to call the new procedure. You should not edit either <i>sp_login_environment</i> or <i>sp_tsql_environment</i>.</p>
Examples	<ul style="list-style-type: none">◆ The following example shows how you can disallow a connection by signaling the INVALID_LOGON error. <pre>create procedure DBA.login_check()</pre>

```

begin
    declare INVALID_LOGON exception for sqlstate
'28000';
    // Allow a maximum of 3 concurrent connections
    if( db_property('ConnCount') > 3 ) then
        signal INVALID_LOGON;
    else
        call sp_login_environment;
    end if;
end
go
grant execute on DBA.login_check to PUBLIC
go
set option PUBLIC.Login_procedure='DBA.login_check'
go

```

☞ For an alternate way to disallow connections, see "RAISERROR statement [T-SQL]" on page 501 of the book *ASA SQL Reference Manual*.

- ◆ The following example shows how you can block connection attempts if the number of failed connections for a user exceeds 3 within a 30 minute period. All blocked attempts during the block out period receive an invalid password error and are logged as failures. The log is kept long enough for a DBA to analyze it.

```

create table DBA.ConnectionFailure(
    pk int primary key default autoincrement,
    user_name char(128) not null,
    tm timestamp not null default current timestamp
)
go

create index ConnFailTime on DBA.ConnectionFailure(
    user_name, tm )
go

```

```
create event ConnFail type ConnectFailed
handler
begin
  declare usr char(128);
  set usr = event_parameter( 'User' );
  // Put a limit on the number of failures logged.
  if (select count(*) from DBA.ConnectionFailure
      where user_name = usr
      and tm >= dateadd( minute, -30,
          current timestamp ) ) < 20 then
    insert into DBA.ConnectionFailure( user_name )
      values( usr );
    commit;
    // Delete failures older than 7 days
    delete DBA.ConnectionFailure
      where user_name = usr
      and tm < dateadd( day, -7, current timestamp );
    commit;
  end if;
end
go
```

```
create procedure DBA.login_check()
begin
  declare usr char(128);
  declare INVALID_LOGON exception for sqlstate
'28000';
  set usr = connection_property( 'Userid' );
  // Block connection attempts from this user
  // if 3 or more failed connection attempts have occurred
  // within the past 30 minutes.
  if (select count(*) from DBA.ConnectionFailure
      where user_name = usr
      and tm >= dateadd( minute, -30,
          current timestamp ) ) >= 3 then
    signal INVALID_LOGON;
  else
    call sp_login_environment;
  end if;
end
go
```

```
grant execute on DBA.login_check to PUBLIC
go
```

```
set option PUBLIC.Login_procedure='DBA.login_check'
go
```

MAX_CURSOR_COUNT option [database]

Function	A resource governor to limit the maximum number of cursors that a connection can use at once.
Allowed values	<i>Integer</i>
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately. DBA permissions are required to set this option for <i>any</i> connection.
Default	50
Description	<p>This resource governor allows a DBA to limit the number of cursors per connection that a user can use. If an operation would exceed the limit for a connection, an error is generated, indicating that the governor for the resource has been exceeded.</p> <p>If a connection executes a stored procedure, that procedure is executed under the permissions of the procedure owner. However, the resources used by the procedure are assigned to the current connection.</p> <p>You can remove resource limits by setting the option to 0 (zero).</p>

MAX_HASH_SIZE option [database]

Function	Specify the default maximum hash size for new indexes.
Allowed values	<i>Integer</i> , from 2 to 64 inclusive.
Scope	Can be set for the PUBLIC group only. Takes effect immediately. DBA permissions are required to set this option.
Default	10
See also	"CREATE INDEX statement" on page 300 of the book <i>ASA SQL Reference Manual</i> "CREATE TABLE statement" on page 350 of the book <i>ASA SQL Reference Manual</i>
Description	This option controls the default hash size for indexes. This option is deprecated.

MAX_PLANS_CACHED option [database]

Function	Specify the maximum number of execution plans to be stored in a cache.
Allowed values	<i>Integer</i>

Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately. DBA permissions are not required to set this option for any connection.
Default	20
See also	"Access plan caching" on page 322 of the book <i>ASA SQL User's Guide</i> "Connection-level properties" on page 618
Description	<p>This option specifies the maximum number of plans cached for each connection. The optimizer caches the execution plan for queries, INSERT, UPDATE, and DELETE statements that are performed inside stored procedures, functions, and triggers. After a statement in a stored procedure, stored function, or trigger is executed several times by a connection, the optimizer builds a reusable plan for the statement.</p> <p>Reusable plans do not use the values of host variables for selectivity estimation or rewrite optimizations. As a result of this, the reusable plan can have a higher cost than if the statement was re-optimized. When the cost of the reusable plan is close to the best observed cost for a statement, the optimizer adds the plan to the plan cache.</p> <p>The cache is cleared when you execute statements, such as CREATE TABLE and DROP TABLE, that modify the table schema. Statements that reference declared temporary tables are not cached.</p> <p>Setting this option to 0 disables plan caching.</p>

MAX_WORK_TABLE_HASH_SIZE option [database]

Function	Specify the maximum hash size used during query optimization for internal temporary tables.
Allowed values	<i>Integer</i> , from 2 to 64 inclusive.
Scope	Can be set for the PUBLIC group only. Takes effect immediately. DBA permissions are required to set this option.
Default	20
See also	"Use of work tables in query processing" on page 160 of the book <i>ASA SQL User's Guide</i> .
Description	In general, you should not need to use this option as the query optimizer allocates hash sizes for the internal temporary tables based on the data distribution within the table. The option allows you to override the optimizer behavior to either increase the maximum hash size it can use beyond 20, or to restrict it further. This option is deprecated.

MAX_STATEMENT_COUNT option [database]

Function	A resource governor to limit the maximum number of prepared statements that a connection can use at once.
Allowed values	<i>Integer</i> ≥ 0
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately. DBA permissions are required to set this option for <i>any</i> connection.
Default	50
Description	<p>This resource governor allows a DBA to limit the number of prepared statements per connection a user can use. If an operation would exceed the limit for a connection, an error is generated, indicating that the governor for the resource has been exceeded.</p> <p>If a connection executes a stored procedure, that procedure is executed under the permissions of the procedure owner. However, the resources used by the procedure are assigned to the current connection.</p> <p>You can remove resource limits by setting the option to 0 (zero).</p>

MIN_PASSWORD_LENGTH option [database]

Function	Sets the minimum length for new passwords in the database.
Allowed values	<i>Integer</i> , greater than or equal to zero. The value is in bytes. For single-byte character sets, this is the same as the number of characters.
Scope	Can be set for the PUBLIC group. Takes effect immediately. DBA permissions are required to set this option.
Default	0 characters
Description	This option allows the database administrator to impose a minimum length on all new passwords for greater security. Existing passwords are not affected.
Example	<p>◆ Set the minimum length for new passwords to 6 bytes.</p> <pre>SET OPTION PUBLIC.MIN_PASSWORD_LENGTH = 6</pre>

MIN_TABLE_SIZE_FOR_HISTOGRAM option [database]

Function	Specify the minimum table size for which histograms are created.
-----------------	--

Allowed values	<i>Integer</i> , greater than or equal to zero.
Scope	Can be set for the PUBLIC group only.
Default	<i>1 000</i> rows
See also	"Optimizer estimates" on page 315 of the book <i>ASA SQL User's Guide</i>
Description	<p>The optimizer uses histograms to estimate the selectivity of search conditions. For small tables, a histogram does not significantly improve the optimizer's ability to choose an efficient plan. This option lets you specify the minimum number of rows a table must have in order for the optimizer to create a histogram for columns in the table.</p> <p>You can change the setting to build histograms for smaller tables, or to build them only for very large tables. This option is ignored by the CREATE STATISTICS statement: when a CREATE STATISTICS statement is executed, a histogram is created regardless of the number of rows in the table.</p>

NEAREST_CENTURY option [compatibility]

Function	Controls the interpretation of two-digit years in string-to-date conversions.
Allowed values	<i>Integer</i> , between 0 and 100 inclusive.
Default	<i>50</i> for databases created with Version 6 or later. <i>0</i> for databases created with Version 5.5 or earlier.
Description	<p>This option controls the handling of two-digit years when converting from strings to dates or timestamps.</p> <p>The NEAREST_CENTURY setting is a numeric value that acts as a rollover point. Two digit years less than the value are converted to 20yy, while years greater than or equal to the value are converted to 19yy.</p> <p>The historical Adaptive Server Anywhere behavior is to add 1900 to the year. Adaptive Server Enterprise behavior is to use the nearest century, so any year where value yy is less than 50, the year is set to 20yy.</p>

NON_KEYWORDS option [compatibility]

Function	Turns off individual keywords, allowing their use as identifiers.
Allowed values	<i>String</i>
Default	The empty string

Description

This option turns off individual keywords or all keywords introduced since a specific release of the product. This provides a way of ensuring that applications created with older versions of the product are not broken by new keywords. If you have an identifier in your database that is now a keyword, you can either add double quotes around the identifier in all applications or scripts, or turn off the keyword using the `NON_KEYWORDS` option.

In addition to specifying individual keywords, you can turn off all keywords since a specified release, using one of the following special values in the list of keywords:

```
keywords_4_0_d, keywords_4_0_c, keywords_4_0_b,  
keywords_4_0_a, keywords_4_0, keywords_5_0_01,  
keywords_5_0
```

The following statement prevents `TRUNCATE` and `SYNCHRONIZE` from being recognized as keywords:

```
SET OPTION NON_KEYWORDS = 'TRUNCATE, SYNCHRONIZE'
```

The following statement prevents all keywords introduced since release 4.0d from being recognized as keywords:

```
SET OPTION NON_KEYWORDS = 'keywords_4_0_d'
```

Each new setting of this option replaces the previous setting. The following statement clears all previous settings.

```
SET OPTION NON_KEYWORDS =
```

A side-effect of this options is that SQL statements that use a turned off keyword cannot be used: they produce a syntax error.

NULLS option [ISQL]

Function	Specifies how NULL values in the database are displayed.
Allowed values	<i>String</i>
Default	<i>(NULL)</i>
Description	Set this according to your preference.

ON_CHARSET_CONVERSION_FAILURE option [database]

Function	Controls what happens if an error is encountered during character conversion.
Allowed values	String. See below for allowed values.

Default	IGNORE
Description	<p>Controls what happens if an error is encountered during character conversion, as follows:</p> <ul style="list-style-type: none">◆ IGNORE Errors and warnings do not appear.◆ WARNING Reports substitutions and illegal characters as warnings. Illegal characters are not translated.◆ ERROR Reports substitutions and illegal characters as errors. <p>Single-byte to single-byte converters are not able to report substitutions and illegal characters, and must be set to IGNORE.</p>

ON_ERROR option [ISQL]

Function	Controls what happens if an error is encountered while reading statements from a command file.
Allowed values	<i>String</i> . See below for allowed values.
Default	PROMPT
Description	<p>Controls what happens if an error is encountered while reading statements from a command file, as follows:</p> <ul style="list-style-type: none">◆ STOP Interactive SQL stops reading statements from the file and returns to the SQL Statements pane for input.◆ PROMPT Interactive SQL prompts the user to see if the user wishes to continue.◆ CONTINUE The error is ignored and Interactive SQL continues reading statements from the command file. The INPUT statement continues with the next row, skipping the row that caused the error.◆ EXIT Interactive SQL terminates.◆ NOTIFY_CONTINUE The error appears in a message box with a single button. Execution continues once the button is clicked.◆ NOTIFY_STOP The error appears in a message box with a single button. Execution of the script stops once the button is clicked.◆ NOTIFY_EXIT The error appears in a message box with a single button. Interactive SQL terminates once the button is clicked.

ON_TSQL_ERROR option [compatibility]

Function	Controls error-handling in stored procedures.
Allowed values	<i>String</i> . See below for allowed values.
Default	CONDITIONAL
See also	"CREATE PROCEDURE statement" on page 305 of the book <i>ASA SQL Reference Manual</i> "CREATE PROCEDURE statement [T-SQL]" on page 312 of the book <i>ASA SQL Reference Manual</i> "Transact-SQL procedure language overview" on page 406 of the book <i>ASA SQL User's Guide</i>
Description	<p>This option controls error handling in stored procedures.</p> <ul style="list-style-type: none">◆ STOP Stop execution immediately upon finding an error.◆ CONDITIONAL If the procedure uses ON EXCEPTION RESUME, and the statement following the error handles the error, continue, otherwise exit. This setting can be used to simulate Adaptive Server Enterprise behavior for stored procedures.◆ CONTINUE Continue execution, regardless of the following statement. If there are multiple errors, the first error encountered in the stored procedure is returned. This option most closely mirrors Adaptive Server Enterprise behavior. <p>When this option is set to STOP or CONTINUE, it supercedes the setting of the CONTINUE_AFTER_RAISERROR option. However, when this option is set to CONDITIONAL (the default), behavior following a RAISERROR statement is determined by the setting of the CONTINUE_AFTER_RAISERROR option.</p>


OPTIMIZATION_GOAL option [database]

Function	Determines whether query processing is optimized towards returning the first row quickly, or minimizing the cost of returning the complete result set.
Allowed values	first-row, all-rows
Default	all-rows
Description	The OPTIMIZATION_GOAL option controls whether Adaptive Server Anywhere optimizes SQL DML statements for response time or total resource consumption.

If the option is set to **all-rows** (the default), then Adaptive Server Anywhere optimizes a query so as to choose an access plan with the minimal estimated total retrieval time. Setting `OPTIMIZATION_GOAL` to **all-rows** may be appropriate for applications that intend to process the entire result set, such as PowerBuilder DataWindow applications. A setting of all-rows is also appropriate for insensitive (ODBC static) cursors since the entire result is materialized when the cursor is opened. It may also be appropriate for scroll (ODBC keyset-driven) cursors, since the intent of such a cursor is to permit scrolling through the result set.

If the option is set to **first-row**, Adaptive Server Anywhere chooses an access plan that is intended to reduce the time to fetch the first row of the query's result, possibly at the expense of total retrieval time. In particular, the Adaptive Server Anywhere optimizer will typically avoid, if possible, access plans that require the materialization of results in order to reduce the time to return the first row. With this setting, the optimizer favors access plans that utilize an index to satisfy a query's `ORDER BY` clause, rather than plans that require an explicit sorting operation.

You can use the `FASTFIRSTROW` table hint in a query's `FROM` clause to set the optimization goal for a specific query to **first-row**, without having to change the `OPTIMIZATION_GOAL` setting.

 For more information about using the `FASTFIRSTROW` table hint, see "FROM clause" on page 433 of the book *ASA SQL Reference Manual*.

OUTPUT_FORMAT option [ISQL]

Function	Sets the output format for the data retrieved by the <code>SELECT</code> statement and redirected into a file, or output using the <code>OUTPUT</code> statement.
Allowed values	<i>String</i> . See below for allowed values.
Default	ASCII
Description	<p>The valid output formats are:</p> <ul style="list-style-type: none">◆ ASCII The output is an ASCII format file with one row per line in the file. All values are separated by commas, and strings are enclosed in apostrophes (single quotes). The delimiter and quote strings can be changed using the <code>DELIMITED BY</code> and <code>QUOTE</code> clauses. If <code>ALL</code> is specified in the <code>QUOTE</code> clause, then all values (not just strings) will be quoted.

Three other special sequences are also used. The two characters `\n` represent a newline character; `\\` represents a single backslash character, and the sequence `\xDD` represents the character with hexadecimal code `DD`.

- ◆ **DBASEII** The output is a dBASE II format file with the column definitions at the top of the file. Note that a maximum of 32 columns can be output. Also, note that columns longer than 255 characters will be truncated in the file.
- ◆ **DBASEIII** The output is a dBASE III format file with the column definitions at the top of the file. Note that a maximum of 128 columns can be output. Also, note that columns longer than 255 characters will be truncated in the file.
- ◆ **EXCEL** The output is an Excel 2.1 worksheet. The first row of the worksheet contains column labels (or names if there are no labels defined). Subsequent worksheet rows contain the actual table data.
- ◆ **FIXED** The output is fixed format, with each column having a fixed width. The width for each column can be specified using the **COLUMN WIDTH** clause. If this clause is omitted, the width for each column is computed from the data type for the column, and is large enough to hold any value of that data type. No column headings are output in this format.
- ◆ **FOXPRO** The output is a FoxPro format file (the FoxPro memo field is different than the dBASE memo field) with the column definitions at the top of the file. Note that a maximum of 128 columns can be output. Also, note that columns longer than 255 characters will be truncated in the file.
- ◆ **HTML** HTML The output is in the Hyper Text Markup Language format.
- ◆ **LOTUS** The output is a Lotus WKS format worksheet. Column names will be put as the first row in the worksheet. Note that there are certain restrictions on the maximum size of Lotus WKS format worksheets that other software (such as Lotus 1-2-3) can load. There is no limit to the size of file Interactive SQL can produce.
- ◆ **SQL** The output is an Interactive SQL INPUT statement required to recreate the information in the table.
- ◆ **XML** The output is an XML file encoded in UTF-8 and containing an embedded DTD. Binary values are encoded in CDATA blocks with the binary data rendered as 2-hex-digit strings.

OUTPUT_LENGTH option [ISQL]

Function	Controls the length used when Interactive SQL exports information to an external file.
Allowed values	<i>Integer</i>
Default	0 (no truncation)
Description	This option controls the length used when Interactive SQL exports information to an external file (using output redirection with the OUTPUT statement). This option affects only ASCII, HTML, and SQL output formats.

OUTPUT_NULLS option [ISQL]

Function	Controls the way NULL values appear in result sets.
Allowed values	<i>String</i>
Default	'NULL'
Description	This option controls the way NULL values appear in result sets. Every time a NULL value is found in the result set, the string from this option is returned instead. This setting applies to data displayed in Interactive SQL on the Results tab in the Results pane as well as to data in output files generated by the OUTPUT statement. This option affects only ASCII, HTML, and SQL output formats.

PERCENT_AS_COMMENT option [compatibility]

Function	Controls the interpretation of the percent character.
Allowed values	ON, OFF
Default	ON
Description	<p>It is recommended that you not use % as a comment marker.</p> <p>Versions of this product before Version 6 treated the percent character (%) in SQL statements exclusively as a comment delimiter. Since Version 5, alternative comment markers such as //, /* */, and -- (double dash) have been available. The double-dash style is the SQL/92 comment delimiter.</p> <p>Adaptive Server Enterprise treats % as a modulo operator and does not support the Adaptive Server Anywhere mod function. Writing a statement that works in both environments and performs a modulo operation was previously impossible.</p>

The `PERCENT_AS_COMMENT` option controls the meaning of `%`. The default setting is `ON` for backwards compatibility. You can set the option to `OFF` for compatibility with Adaptive Server Enterprise.

Procedures, triggers and views that were created with `%`-style comments are converted to double-dash comments when they are stored in the catalog.

Existing procedures must be recreated before changing option

Any existing procedures that contain `%`-style comments must be recreated before you change the option setting; otherwise, the procedures will fail to load.

The Sybase Central code editor does not highlight `%`-style comments. If you wish to have your comments highlighted in the Sybase Central editor, you should use one of the other comment delimiters.

PINNED_CURSOR_PERCENT_OF_CACHE option [database]

Function	Specifies how much of the cache can be used for pinning cursors.
Allowed values	<i>Integer</i> , between 0-100
Scope	Can be set for the <code>PUBLIC</code> group.
Default	<i>10</i>
Description	<p>The server uses pages of virtual memory for the data structures needed to implement cursors. These pages are kept locked in memory between fetch requests so they are readily available when the next fetch request arrives.</p> <p>To prevent these pages from occupying too much of the cache in low memory environments, a limit is placed on the percentage of the cache allowed to be used for pinning cursors. You can use the <code>PINNED_CURSOR_PERCENT_OF_CACHE</code> option to adjust this limit.</p> <p>The option value is specified as a percentage from 0 to 100, with a default of 10. Setting the option to 0 means that cursor pages will not be pinned between fetch requests.</p>

PRECISION option [database]

Function	Specifies the maximum number of digits in the result of any decimal arithmetic.
Allowed values	<i>Integer</i> , between 0 and 127 inclusive.

Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	30
Description	<p>Precision is the total number of digits to the left and right of the decimal point. The "SCALE option" on page 598 specifies the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum PRECISION.</p> <p>Multiplication, division, addition, subtraction, and aggregate functions can all have results that exceed the maximum precision.</p> <p>For example, when a DECIMAL(8,2) is multiplied with a DECIMAL(9,2), the result could require a DECIMAL(17,4). If PRECISION is 15, only 15 digits will be kept in the result. If SCALE is 4, the result will be a DECIMAL(15,4). If SCALE is 2, the result will be a DECIMAL(15,2). In both cases, there is a possibility of overflow.</p>

PREFETCH option [database]

Function	The PREFETCH option acts as a toggle allowing you to turn fetching on and off.
Allowed values	ON, OFF
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	ON
Description	<p>This option controls whether rows are fetched to the client side in advance of being made available to the client application. Fetching a number of rows at a time, even when the client application requests rows one at a time (for example, when looping over the rows of a cursor) cuts down on response time and improves overall throughput by cutting down the number of requests to the database.</p> <p>The setting of PREFETCH is ignored by Open Client and JDBC connections.</p> <p>🔗 For more information, see "Prefetching rows" on page 40 of the book <i>ASA Programming Guide</i> and "DisableMultiRowFetch connection parameter" on page 175.</p>

PRESERVE_SOURCE_FORMAT option [database]

Function	Controls whether the original source definition of procedures, triggers, views, and event handlers is saved in system files. If saved, it is saved in the column <i>source</i> in SYSTABLE, SYSPROCEDURE, SYSTRIGGER, and SYSEVENT.
Allowed values	ON, OFF
Default	ON
Description	<p>When PRESERVE_SOURCE_FORMAT is ON, the server saves the formatted source from CREATE and ALTER statements on procedures, views, triggers, and events, and puts it in the appropriate system table's <i>source</i> column.</p> <p>Unformatted source text is stored in the same system tables, in the columns <i>proc_defn</i>, <i>trigger_defn</i>, and <i>view_defn</i>. However, these definitions are not easy to read in Sybase Central. The formatted <i>source</i> column allows you to view the definitions with the spacing, comments, and case that you want.</p> <p>This option can be turned off to reduce space used to save object definitions in the database. The option can be set only for the user PUBLIC.</p>

PREVENT_ARTICLE_PKEY_UPDATE option [database]

Function	Controls updates to the primary key columns of tables involved in publications.
Allowed values	ON, OFF
Default	ON
Description	Setting this option to ON disallows updates to the primary key columns of tables that are part of a publication. This option helps ensure data integrity, especially in a replication and synchronization environment

QUALIFY_OWNERS option [replication]

Function	Controls whether SQL statements being replicated by SQL Remote should use qualified object names.
Allowed values	ON, OFF
Default	<p>The default in Adaptive Server Anywhere is ON.</p> <p>The default in Adaptive Server Enterprise is OFF.</p>

Description Qualifying owners in Adaptive Server Enterprise setups is rarely needed because it is common for objects to be owned by **dbo**. When qualification is not needed in Adaptive Server Anywhere setups, messages will be slightly smaller with the option OFF.

QUERY_PLAN_ON_OPEN option [compatibility]

Function Controls whether a plan is returned when a cursor is opened.

Allowed values **ON, OFF**

Default **OFF**

Description In early versions of the software, each time an OPEN was done on a cursor, the server would return in the SQLCA **sqlerrmc** field a string representing the query plan (limited to 70 bytes). A more complete description can be obtained using the EXPLAIN statement or the PLAN function. For this reason, computing and returning the query plan on an OPEN is needed only for compatibility with old applications. The QUERY_PLAN_ON_OPEN option controls whether the plan is returned on an OPEN. By default, the setting is **OFF**.

QUOTE_ALL_IDENTIFIERS option [replication]

Function Controls whether SQL statements being replicated by SQL Remote should use quoted identifiers

Allowed values **ON, OFF**

Default **OFF**

Description When this option is off, *dbremote* quotes identifiers that require quotes by Adaptive Server Anywhere (as it has always done) and *ssremote* does not quote any identifiers.

When the option is on, all identifiers are quoted.


QUOTED_IDENTIFIER option [compatibility]

Function Controls the interpretation of strings that are enclosed in double quotes.


Allowed values **ON, OFF**

Default **ON**

OFF for Open Client and JDBC connections

Description	<p>This option controls whether strings that are enclosed in double quotes are interpreted as identifiers (ON) or as literal strings (OFF). The <code>QUOTED_IDENTIFIER</code> option is included for Transact-SQL compatibility.</p> <p>Note that Interactive SQL automatically executes the equivalent of</p> <pre>SET TEMPORARY OPTION QUOTED_IDENTIFIER = ON</pre> <p>when it connects to a database. <code>TEMPORARY</code> options apply only to the current connection, and last only until the connection is closed.</p> <p> For more information, see "Setting options for Transact-SQL compatibility" on page 395 of the book <i>ASA SQL User's Guide</i>.</p>
--------------------	---

RECOVERY_TIME option [database]

Function	Sets the maximum length of time, in minutes, that the database server will take to recover from system failure.
Allowed values	<i>Integer</i> , in minutes
Scope	Can be set only for the <code>PUBLIC</code> group. Takes effect when server is restarted.
Default	2
Description	<p>This option is used with the "<code>CHECKPOINT_TIME</code>" option on page 557 to decide when checkpoints should be done.</p> <p>Adaptive Server Anywhere uses a heuristic to estimate the recovery time based on the operations that have been performed since the last checkpoint. Thus, the recovery time is not exact.</p> <p> For more information, see "The automatic recovery process" on page 325.</p>

REPLICATE_ALL option [replication]

Function	Allows an entire database to act as a primary site in a Replication Server setup.
Allowed values	ON, OFF
Default	OFF
Description	This option is used by the LTM only. When it is set to ON , the entire database is set to act as a primary site in a Replication Server installation. All changes to the database are sent to Replication Server by the LTM.

REPLICATION_ERROR option [replication]

Function	For SQL Remote, allows you to specify a stored procedure to be called by the Message Agent when a SQL error occurs.
Allowed values	<i>Stored procedure name</i>
Default	<i>No procedure</i>
Description	<p>For SQL Remote, the REPLICATION_ERROR option allows you to specify a stored procedure to be called by the Message Agent when a SQL error occurs. By default no procedure is called.</p> <p>The procedure must have a single argument of type CHAR, VARCHAR, or LONG VARCHAR. The procedure is called once with the SQL error message and once with the SQL statement that causes the error.</p> <p>Although the option allows you to track and monitor SQL errors in replication, you must still design them out of your setup; this option is not intended to resolve such errors.</p>

RETURN_DATE_TIME_AS_STRING option [database]

Function	To control how a date, time, or timestamp value is passed to the client application when queried.
Allowed values	ON, OFF
Scope	Can be set as a temporary option only, for the duration of the current connection.
Default	OFF
See also	"DATE_FORMAT option" on page 562 "TIME_FORMAT option" on page 601 "TIMESTAMP_FORMAT option" on page 602
Description	<p>This option indicates whether date, time, and timestamp values are returned to applications as a date or time datatype or as a string.</p> <p>When this option is set to ON, the engine converts the date, time, or timestamp value to a string before it is sent to the client in order to preserve the TIMESTAMP_FORMAT, DATE_FORMAT, or TIME_FORMAT option setting.</p> <p>Sybase Central and Interactive SQL automatically turn the RETURN_DATE_TIME_AS_STRING option ON.</p>

RETURN_JAVA_AS_STRING option [database]

Function	To control how a Java object is passed to the client application when queried.
Allowed values	ON, OFF
Scope	Can be set as a temporary option only, for the duration of the current connection.
Default	OFF
Description	<p>The option indicates how a Java object is returned to applications connecting over Open Client or jConnect.</p> <p>By default, a Java object is returned over TDS as a Sun serialization of the object. The client, or receiver of the object serialization, is responsible for deserializing the object into an instance.</p> <p>If RETURN_JAVA_AS_STRING is set to ON, the object is first converted to an instance of java.lang.String using the toString() method, and the String object is returned to the client.</p>

RI_TRIGGER_TIME option [compatibility]

Function	Controls the relative timing of referential integrity checks and trigger actions.
Allowed values	BEFORE, AFTER
Scope	Can be set for the PUBLIC option only. DBA authority is required to set this option.
Default	AFTER
Description	<p>The option can be set to either BEFORE or AFTER. When it's set to AFTER, referential integrity actions are executed after the UPDATE or DELETE.</p> <p>Only the PUBLIC setting can be used; any other setting is ignored.</p>

ROW_COUNTS option [database]

Function	Specifies whether the database will always count the number of rows in a query when it is opened.
Allowed values	ON, OFF
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	OFF

Description If this option is set to **OFF**, the row count is usually only an estimate. If this option is set to **ON**, the row count is always accurate.

Warning: When ROW_COUNTS is set to **ON**, it may take significantly longer to execute queries. In fact, it will usually cause Adaptive Server Anywhere to execute the query twice, doubling the execution time.

SCALE option [database]

Function Specifies the minimum number of digits after the decimal point when an arithmetic result is truncated to the maximum PRECISION.

Allowed values *Integer*, between 0 and 127 inclusive.

Scope Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.

Default 6

Description Multiplication, division, addition, subtraction, and aggregate functions can all have results that exceed the maximum precision. See "PRECISION option" on page 591 for an example.

SORT_COLLATION option [database]

Function Allows implicit use of the SORTKEY function on ORDER BY expressions.

Allowed values **INTERNAL**, *collation_name*, or *collation_id*

Default **INTERNAL**

See also "SORTKEY function" on page 179 of the book *ASA SQL Reference Manual*

Description When the value of this option is **INTERNAL**, the ORDERBY clause remains unchanged.

When the value of this option is set to a valid *collation name* or *collation ID*, any string expression in the ORDER BY clause is treated as if the SORTKEY function had been invoked.

Examples Set the sort collation to binary:

```
set temporary option sort_collation='binary'
```

Having the sort collation set to binary transforms


```
SELECT name, id
FROM product
ORDER BY name, id

and

SELECT name, id
FROM product
ORDER BY 1,2

into

SELECT name, id
FROM product
ORDER BY SORTKEY(name, 'binary'), id
```

SQL_FLAGGER_ERROR_LEVEL option [compatibility]

Function	Controls the response to any SQL that is not part of a specified set of SQL/92.
Allowed values	E, I, F, or W
Default	W
Description	<p>This option flags any SQL that is not part of a specified set of SQL/92 as an error.</p> <p>The allowed values of <i>level</i> are as follows:</p> <ul style="list-style-type: none">♦ E Flag syntax that is not entry-level SQL/92 syntax♦ I Flag syntax that is not intermediate-level SQL/92 syntax♦ F Flag syntax that is not full-SQL/92 syntax♦ W Allow all supported syntax

SQL_FLAGGER_WARNING_LEVEL option [compatibility]

Function	Controls the response to any SQL that is not part of a specified set of SQL/92.
Allowed values	E, I, F, or W
Default	W
Description	<p>This option flags any SQL that is not part of a specified set of SQL/92 as a warning.</p> <p>The allowed values of <i>level</i> are as follows:</p>

- ◆ **E** Flag syntax that is not entry-level SQL/92 syntax
- ◆ **I** Flag syntax that is not intermediate-level SQL/92 syntax
- ◆ **F** Flag syntax that is not full-SQL/92 syntax
- ◆ **W** Allow all supported syntax

STATISTICS option [ISQL]

Function	Controls whether the execution time is enabled.
Allowed values	<i>Any non-negative integer</i>
Default	7
See also	"ISQL_COMMAND_TIMING option" on page 572 "ISQL_PLAN option" on page 574
Description	<p>This option is an older version of the ISQL_PLAN option.</p> <p>When you set the STATISTICS option to 0, no information appears in the Messages pane after you execute a SQL statement. When you set the option to any other number, the execution time appears in the Messages pane after you execute a statement. The number you specify also becomes the default height (in lines) of the Messages pane.</p>

STRING_RTRUNCATION option [compatibility]

Function	Determines whether an error is raised when an INSERT or UPDATE truncates a CHAR or VARCHAR string.
Allowed values	ON, OFF
Default	OFF
Description	If the truncated characters consist only of spaces, no exception is raised. The setting of ON corresponds to ANSI/ISO SQL/92 behavior. When it is set to OFF , the exception is not raised and the character string is silently truncated.

SUBSCRIBE_BY_REMOTE option [replication]

Function	Controls interpretation of NULL or empty-string SUBSCRIBE BY values.
Allowed values	ON, OFF
Default	ON

Description	When the option is set to ON , operations from remote databases on rows with a SUBSCRIBE BY value that is NULL or an empty string assume that the remote user is subscribed to the row. When it is set to OFF , the remote user is assumed not to be subscribed to the row.
--------------------	---

SUPPRESS_TDS_DEBUGGING option [database]

Function	Determines whether TDS debugging information appears in the server window.
Allowed values	ON, OFF
Default	OFF
Description	<p>When the server is started with the <code>-z</code> option, debugging information appears in the server window, including debugging information about the TDS protocol.</p> <p>The SUPPRESS_TDS_DEBUGGING option restricts the debugging information about TDS that appears in the server window. When this option is set to OFF (the default) TDS debugging information appears in the server window.</p>

TDS_EMPTY_STRING_IS_NULL option [database]

Function	Controls whether empty strings are returned as NULL or a string containing one blank character for TDS connections.
Allowed values	ON, OFF
Default	OFF
Description	By default, this option is set to OFF and empty strings are returned as a string containing one blank character for TDS connections. When this option is set to ON , empty strings are returned as NULL strings for TDS connections. Non-TDS connections distinguish empty strings from NULL strings.

TIME_FORMAT option [compatibility]

Function	Sets the format for times retrieved from the database.
Allowed values	<i>String</i> , composed of the symbols listed below.

Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	HH:NN:SS.SSS For Open Client and JDBC connections the default is set to HH:NN:SS.SSS.
Description	<p>The format is a string using the following symbols:</p> <ul style="list-style-type: none">♦ hh Two digit hours (24 hour clock)♦ nn Two digit minutes♦ mm Two digit minutes if following a colon (as in hh:mm)♦ ss[s...] Two digit seconds plus optional fraction <p>Each symbol is substituted with the appropriate data for the time that is being formatted. Any format symbol that represents character rather than digit output can be put in uppercase, which causes the substituted characters to also be in uppercase. For numbers, using mixed case in the format string suppresses leading zeros.</p>

TIME_ZONE_ADJUSTMENT option [database]

Function	Allows a connection's time zone adjustment to be modified.
Allowed values	<i>Integer</i> , (for example, 300), or <i>Negative integer</i> , enclosed in quotation marks (for example, '-300'), or <i>String</i> , representing a time in hours and minutes, preceded by + or -, enclosed in quotation marks (for example, '+5:00', or '-5:00').
Default	If the client is connecting via ESQL, ODBC, OLE DB, or ADO, the default value is set according to the client's time zone. If the client is connecting via jConnect or OpenClient, the default is based on the server's time zone.
See also	"Connection-level properties" on page 618
Description	The TIME_ZONE_ADJUSTMENT option value is the same value as that returned by <code>connection_property('TimeZoneAdjustment')</code> . The value represents the number of minutes that must be added to the Coordinated Universal Time (UTC) to display time local to the connection.

TIMESTAMP_FORMAT option [compatibility]

Function	Sets the format for timestamps that are retrieved from the database.
-----------------	--

Allowed values	<i>String</i> , composed of the symbols listed below.
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	<i>YYYY-MM-DD HH:NN:ss.SSS</i> For Open Client and JDBC connections the default is set to <i>YYYY-MM-DD HH:NN:SS.SSS</i> .
Description	The format is a string using the following symbols:

Symbol	Description
yy	Two digit year
yyyy	Four digit year
mm	Two digit month, or two digit minutes if following a colon (as in hh:mm')
mmm[m...]	Character short form for months—as many characters as there are "m"s
dd	Two digit day of month
ddd[d...]	Character short form for day of the week
hh	Two digit hours
nn	Two digit minutes
ss.ssssss	Seconds and fractions of a second, up to six decimal places. Not all platforms support timestamps to a precision of six places.
aa	am or pm (12 hour clock)
pp	pm if needed (12 hour clock)
f	Use French days and months (deprecated)

Each symbol is substituted with the appropriate data for the date that is being formatted.

For symbols that represent character data (such as mmm), you can control the case of the output as follows:

- ◆ Type the symbol in all upper case to have the format appear in all upper case. For example, MMM produces JAN.
- ◆ Type the symbol in all lower case to have the format appear in all lower case. For example, mmm produces jan.

- ◆ Type the symbol in mixed case to have Adaptive Server Anywhere choose the appropriate case for the language that is being used. For example, in English, typing `Mmm` produces `May`, while in French it produces `mai`.

For symbols that represent numeric data, you can control zero-padding with the case of the symbols:

- ◆ Type the symbol in same-case (such as `MM` or `mm`) to allow zero padding. For example, `yyyy/mm/dd` could produce `2002/01/01`.
- ◆ Type the symbol in mixed case (such as `Mm`) to suppress zero padding. For example, `yyyy/Mm/Dd` could produce `2002/1/1`.

TRUNCATE_DATE_VALUES option [database]

Function	To change the storage of time values within DATE data type values.
Allowed values	ON , OFF
Default	ON for databases created with Version 7 or later software, or upgraded to Version 7 or later. OFF for older databases
See also	"DATE_FORMAT option" on page 562 "DATE data type" on page 69 of the book <i>ASA SQL Reference Manual</i>
Description	When the option is set to ON , columns or variables declared as DATE data types have no hours and minutes stored. When the option is set to OFF , DATE columns or variables have hours and minutes stored also. These time components are typically not displayed because of the DATE_FORMAT setting. However, exact comparisons between DATE values may unexpectedly fail because of non-matching time components.

TRUNCATE_TIMESTAMP_VALUES option [database]

Function	Limits the resolution of timestamp values.
Allowed values	ON , OFF
Scope	Can be set for the PUBLIC group. This option should not be enabled for databases already containing timestamp data.
Default	OFF

Description	A TIMESTAMP value is precise to six decimal places in Adaptive Server Anywhere. However, to maintain compatibility with other software, which may truncate the TIMESTAMP value to three decimal places, you can set the TRUNCATE_TIMESTAMP_VALUES option to ON to limit the number of decimal places Adaptive Server Anywhere stores. The DEFAULT_TIMESTAMP_INCREMENT option determines the number of decimal places to which the TIMESTAMP value is truncated.
Example	Setting the DEFAULT_TIMESTAMP_INCREMENT option to 100 000 causes truncation after the first decimal place in the seconds component, allowing a value such as '2000/12/05 10:50:53:700' to be stored.

TRUNCATE_WITH_AUTO_COMMIT option [database]

Function	To speed up TRUNCATE TABLE statements
Allowed values	ON, OFF
Scope	Can be set only for the PUBLIC group.
Default	ON
See also	" TRUNCATE TABLE statement" on page 567 of the book <i>ASA SQL Reference Manual</i>
Description	<p>If TRUNCATE_WITH_AUTO_COMMIT is set to ON, then a COMMIT is executed both before and after the TRUNCATE TABLE statement is executed. The primary purpose of the option is to enable faster table truncation (delete of all rows).</p> <p>There are some cases where a fast TRUNCATE cannot be done:</p> <ul style="list-style-type: none">◆ If there are foreign keys either to or from the table◆ If the TRUNCATE TABLE statement is executed within a trigger◆ If the TRUNCATE TABLE statement is executed within an atomic statement

TRUNCATION_LENGTH option [ISQL]

Function	Controls the truncation of wide columns for displays to fit on a screen.
Allowed values	<i>Integer</i>
Default	30

Description	<p>When SELECT statement results are displayed on the screen, each column of output is limited to the width of the screen. The TRUNCATION_LENGTH option is used to reduce the width of wide columns so that more than one column will fit on the screen. A value of 0 means that columns are not truncated.</p> <p>The default TRUNCATION_LENGTH is 30. For character-mode systems, this is an actual number of characters. For windowing systems, TRUNCATION_LENGTH is used to estimate an area of the screen to be used for display since proportional fonts are used.</p>
--------------------	--

TSQL_HEX_CONSTANT option [compatibility]


Function	Controls whether hexadecimal constants are treated as binary typed constants.
Allowed values	ON, OFF
Default	ON
Description	When this option is set to ON , hexadecimal constants are treated as binary typed constants. To get the historical behavior, set the option to OFF.

TSQL_VARIABLES option [compatibility]

Function	Controls whether the @ sign can be used as a prefix for Embedded SQL host variable names.
Allowed values	ON, OFF
Default	OFF ON for Open Client and JDBC connections
Description	When this option is set to ON , you can use the @ sign instead of the colon as a prefix for host variable names in Embedded SQL. This is implemented primarily for Transact-SQL compatibility.

USER_ESTIMATES option [database]

Function	Controls whether or not user selectivity estimates in query predicates are respected or ignored by the query optimizer.
Allowed values	ENABLED, DISABLED, OVERRIDE-MAGIC

Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	OVERRIDE-MAGIC
Description	<p>Adaptive Server Anywhere allows you to specify user selectivity estimates can improve the optimizer's performance when the server is unable to accurately predict the selectivity of a predicate. However, user selectivity estimates should be used only in appropriate circumstances. For example, it may be useful to supply a selectivity estimate for a predicate that involves one or more functions if the OVERRIDE-MAGIC selectivity estimate used by the optimizer is significantly different from the actual selectivity.</p> <p>If you have used selectivity estimates that are inaccurate as a workaround to performance problems where the software-selected access plan was poor, it is recommended that you set this option to DISABLED. The server may not select an optimal plan if you use inaccurate estimates.</p> <p> For more information about user selectivity estimates, see "Explicit selectivity estimates" on page 31 of the book <i>ASA SQL Reference Manual</i>.</p> <p>When a user selectivity estimate is supplied with a predicate, the estimate is respected or ignored based on the setting of this option. The following values are accepted:</p> <ul style="list-style-type: none">◆ ENABLED All user-supplied selectivity estimates are respected. You can also use ON to turn on this option.◆ OVERRIDE-MAGIC A user selectivity estimate is respected and used only if the optimizer would otherwise choose to use its last-resort, heuristic value (also called the magic value).◆ DISABLED All user estimates are ignored and magic values are used when no other estimate data is available. You can also use OFF to turn off this option.

VERIFY_ALL_COLUMNS option [replication]

Function	Controls whether messages that contain updates published by the local database are sent with all column values included.
Allowed values	ON, OFF
Default	OFF
Description	This option is used by SQL Remote only. When it is set to ON , messages that contain updates published by the local database are sent with all column values included, and a conflict in any column triggers a RESOLVE UPDATE trigger at the subscriber database.

VERIFY_THRESHOLD option [replication]

Function	Controls which columns are verified when updates are replicated.
Allowed values	<i>Integer</i> , in bytes
Default	<i>1 000</i>
Description	This option is used by SQL Remote only. If the data type of a column is longer than the threshold, old values for the column are not verified when an UPDATE is replicated. This keeps the size of SQL Remote messages down, but has the disadvantage that conflicting updates of long values are not detected.

WAIT_FOR_COMMIT option [database]

Function	Determines when foreign key integrity is checked, as data is manipulated.
Allowed values	ON or OFF
Scope	Can be set for an individual connection or for the PUBLIC group. Takes effect immediately.
Default	OFF
Description	If this option is set to ON , the database does not check foreign key integrity until the next COMMIT statement. Otherwise, all foreign keys that are not created with the CHECK_ON_COMMIT option are checked as they are inserted, updated or deleted.

CHAPTER 17

Database Performance and Connection Properties

About this Chapter This chapter contains information and tables relating to database performance monitoring and database connection parameters.

Contents

Topic	Page
Database performance statistics	610
Database properties	618

Database performance statistics

Adaptive Server Anywhere provides a set of statistics that can be used to monitor database performance. These are accessible from Sybase Central, and client applications can access the statistics as functions. In addition, these statistics are made available by the server to the Windows Performance Monitor.


Note:

The Windows Performance Monitor is available in Windows NT/2000/XP.

This section describes how to access performance and related statistics from client applications, how to monitor database performance using Sybase Central, and how to monitor database performance using the Windows Performance Monitor.

Performance Monitor statistics

The Windows Performance Monitor is an application for viewing the behavior of objects such as processors, memory, and applications. Adaptive Server Anywhere provides many statistics for the Performance Monitor to display.

 For information on how to use the Performance Monitor, see "Monitoring database statistics from Windows Performance Monitor" on page 165 of the book *ASA SQL User's Guide*.

Adaptive Server Anywhere makes statistics available for the Performance Monitor. Rates are reported per second. The statistics are grouped into the following areas:

- ◆ "Cache statistics" on page 611
- ◆ "Checkpoint and recovery statistics" on page 611
- ◆ "Communications statistics" on page 612
- ◆ "Disk I/O statistics" on page 613
- ◆ "Disk read statistics" on page 613
- ◆ "Disk write statistics" on page 614
- ◆ "Index statistics" on page 615
- ◆ "Java VM statistics" on page 615
- ◆ "Memory pages statistics" on page 615

- ◆ "Request statistics" on page 616
- ◆ "Miscellaneous statistics" on page 617

Cache statistics

These statistics describe the use of the cache.

Statistic	Scope	Description
Cache Hits/sec	Connection and Database	The rate at which database page lookups are satisfied by finding the page in the cache.
Cache Reads: Index Interior/sec	Connection and Database	The rate at which index internal-node pages are read from the cache.
Cache Reads: Index Leaf/sec	Connection and Database	The rate at which index leaf pages are read from the cache.
Cache Reads: Table/sec	Connection and Database	The rate at which table pages are read from the cache.
Cache Reads: Total Pages/sec	Connection and Database	The rate at which database pages are looked up in the cache.
Cache Size: Current	Engine	The current size of the database server cache, in kilobytes.
Cache Size: Maximum	Engine	The maximum allowed size of the database server cache, in kilobytes.
Cache Size: Minimum	Engine	The minimum allowed size of the database server cache, in kilobytes.
Cache Size: Peak	Engine	The peak size of the database server cache, in kilobytes.

Checkpoint and recovery statistics

These statistics isolate the checkpoint and recovery actions performed when the database is in an idle state.

Statistic	Scope	Description
Checkpoint Flushes/sec	Database	The rate at which ranges of adjacent pages are written out during a checkpoint.
Checkpoint Log	Database	The rate at which the transaction log is checkpointed.
Checkpoint Urgency	Database	Checkpoint Urgency, expressed as a percentage.
Checkpoints/sec	Database	The rate at which checkpoints are performed.
Idle Actives/sec	Database	The rate at which the engine's idle thread becomes active to do idle writes, idle checkpoints, etc.

Statistic	Scope	Description
Idle Checkpoint Time	Database	The total time spent doing idle checkpoints, in seconds.
Idle Checkpoints/sec	Database	The rate at which checkpoints are completed by the engine's idle thread. An idle checkpoint occurs whenever the idle thread writes out the last dirty page in the cache.
Idle Writes/sec	Database	The rate at which disk writes are issued by the engine's idle thread.
Recovery I/O Estimate	Database	The estimated number of I/O operations required to recover the database.
Recovery Urgency	Database	Recovery Urgency expressed as a percentage.
Server Idle Waits/sec	Server	The number of times per second the server goes idle waiting for I/O completion or a new request.

Communications statistics

These statistics describe client/server communications activity.

Statistic	Scope	Description
Comm: Buffer Misses	Server	The total number of network buffer allocations exceeding the connection buffer pool.
Comm: Bytes Received/sec	Connection and Server	The rate at which network data (in bytes) are received.
Comm: Bytes Received Uncompressed/sec	Connection and Server	The rate at which bytes would have been received if compression was disabled.
Comm: Bytes Sent/sec	Connection and Server	The rate at which bytes are transmitted over the network.
Comm: Bytes Sent Uncompressed/sec	Connection and Server	The rate at which bytes would have been sent if compression was disabled.
Comm: Free Buffers	Server	Number of free network buffers.
Comm: Multi-packets Received/sec	Server	The rate at which multi-packet deliveries are received.
Comm: Multi-packets Sent/sec	Server	The rate at which multi-packet deliveries are transmitted.

Statistic	Scope	Description
Comm: Packets Received/sec	Connection and Server	The rate at which network packets are received.
Comm: Packets Received Uncompressed/sec	Connection and Server	The rate at which network packets would have been received if compression was disabled.
Comm: Packets Sent/sec	Connection and Server	The rate at which network packets are transmitted.
Comm: Packets Sent Uncompressed/sec	Connection and Server	The rate at which network packets would have been transmitted if compression was disabled.
Comm: Send Fails/sec	Server	The rate at which the underlying protocol(s) failed to send a packet.
Comm: TotalBuffers	Server	The total number of network buffers.

Disk I/O statistics

These statistics combine disk reads and disk writes to give overall information about the amount of activity devoted to disk I/O.

Statistic	Scope	Description
Disk: Active I/O	Database	The current number of file I/Os issued by the engine which have not yet completed.
Disk: Maximum I/O	Database	The maximum value "Disk Reads: Active I/O" has reached.

Disk read statistics

These statistics describe the amount and type of activity devoted to reading information from disk.

Statistic	Scope	Description
Disk Reads: Total Pages/sec	Connection and Database	The rate at which pages are read from a file.
Disk Reads: Active	Database	The current number of file reads issued by the engine which have not yet completed.
Disk Reads: Index interior/sec	Connection and Database	The rate at which index internal-node pages are being read from disk.
Disk Reads: Index leaf/sec	Connection and Database	The rate at which index leaf pages are being read from disk.
Disk Reads: Table/sec	Connection and Database	The rate at which table pages are being read from disk.
Disk Reads: Maximum Active	Database	The maximum value "Disk Reads: Active" has reached.

Disk write statistics

These statistics describe the amount and type of activity devoted to reading information from disk.

Statistic	Scope	Description
Disk Writes: Active	Database	The current number of file writes issued by the engine which are not yet completed.
Disk Writes: Maximum Active	Database	The maximum value "Disk Writes: Active" has reached.
Disk Writes: Commit Files/sec	Database	The rate at which the engine forces a flush of the disk cache. Windows NT/2000/XP and NetWare platforms use unbuffered (direct) I/O, so the disk cache does not need to be flushed.
Disk Writes: Database Extends/sec	Database	The rate at which the database file is extended, in pages/sec.
Disk Writes: Temp Extends/sec	Database	The rate at which temporary files are extended, in pages/sec.
Disk Writes: Pages/sec	Connection and Database	The rate at which modified pages are being written to disk.
Disk Writes: Transaction Log/sec	Connection and Database	The rate at which pages are written to the transaction log.
Disk Writes:	Connection and	Occurs when a commit of the transaction

Statistic	Scope	Description
Transaction Log Group Commits	Database	log is requested but the log has already been written (so the commit was done for "free").

Index statistics

These statistics describe the use of the index.

Statistic	Scope	Description
Index: Adds/sec	Connection and Database	The rate at which entries are added to indexes.
Index: Lookups/sec	Connection and Database	The rate at which entries are looked up in indexes.
Index: Full Compares/sec	Connection and Database	The rate at which comparisons beyond the hash value in an index must be performed.

Java VM statistics

These statistics describe the memory used by the Java VM.

Statistic	Scope	Description
JVM: Global Fixed Size	Server	The total bytes allocated to the Java VM fixed heap.
JVM: Heap Size	Connection and Database	The total bytes allocated to connection Java VMs
JVM: Namespace Size	Database	The total number of bytes allocated to the Java VM namespace.

Memory pages statistics

These statistics describe the amount and purpose of memory used by the database server.

Statistic	Scope	Description
Mem Pages: Locked Heap	Server	The number of heap pages locked in the cache.
Mem Pages: Lock Table	Database	The number of pages used to store lock information.
Mem Pages: Rollback Log	Connection and Database	The number of pages in the rollback log.
Mem Pages: Main Heap	Server	The number of pages used for global engine data structures.
Mem Pages: Map Pages	Database	The number of map pages used for accessing the lock table, frequency table, and table layout.

Statistic	Scope	Description
Mem Pages: Procedure Definitions	Database	The number of relocatable heap pages used for procedures.
Mem Pages: Relocatable	Database	The number of pages used for relocatable heaps (cursors, statements, procedures, triggers, views, etc.).
Mem Pages: Relocations/sec	Database	The rate at which relocatable heap pages are read from the temporary file.
Mem Pages: Trigger Definitions	Database	The number of relocatable heap pages used for triggers.
Mem Pages: View Definitions	Database	The number of relocatable heap pages used for views.

Request statistics

These statistics describe the database server activity devoted to responding to requests from client applications.

Statistic	Scope	Description
Requests	Server	The rate at which the engine is entered to allow it to handle a new request or continue processing an existing request.
Requests: Active	Server	The number of engine threads currently handling a request.
Requests: Unscheduled	Server	The number of requests that are currently queued up waiting for an available engine thread.
Cursors	Connection	The number of declared cursors currently maintained by the engine.
Cursors Open	Connection	The number of open cursors currently maintained by the engine.
Statements	Connection	The number of prepared statements currently maintained by the engine.
Statement Prepares	Connection	The rate at which statement prepares are being handled by the engine.
Transaction Commits	Connection	The rate at which Commit requests are handled.
Transaction Rollbacks	Connection	The rate at which Rollback requests are handled.

Miscellaneous statistics

Statistic	Scope	Description
Avail IO	Server	Expressed as a count.
Context Switches	Connection	The rate at which the current engine thread changes.
Main Heap Bytes	Server	The number of bytes used for global engine data structures.
Query Low Memory Strategy	Connection and Database	The number of times the server changed its execution plan during execution because of low memory conditions.
Request Queue Waits/sec	Server	The rate at which the engine must wait for room in the request queue.
Temporary Table Pages	Connection and Database	The number of pages in the temporary file used for temporary tables.

Database properties

Accessing properties

Adaptive Server Anywhere provides a set of properties that are made available to client applications. These properties describe aspects of connection, database, and database server behavior.

Each type of property can be accessed by supplying its name as an argument to a system function.

❖ To access connection properties:

- ◆ Use the *connection_property* system function: the following statement returns the number of pages that have been read from file by the current connection.

```
select connection_property ( 'DiskRead' )
```

❖ To access database properties:

- ◆ Use the *db_property* system function. For example, the following statement returns the page size of the current database:

```
select db_property ( 'PageSize' )
```

❖ To access database server properties:

- ◆ Use the *property* system function: the following statement returns the number of pages that have been read from file by the current connection.

```
select property ( 'MainHeapPages' )
```

Connection-level properties

The following table lists properties available for each connection.

Examples

❖ To retrieve the value of a connection property:

- ◆ Use the *connection_property* system function. The following statement returns the number of pages that have been read from file by the current connection.

```
select connection_property ( 'DiskRead' )
```

❖ To retrieve the values of all connection properties:

- ◆ Use the *sa_conn_properties* system procedure:

```
call sa_conn_properties
```




A separate row appears for each connection.

Descriptions

Property	Description
Allow_nulls_by_default	"ALLOW_NULLS_BY_DEFAULT option" on page 550
Ansi_blanks	"ANSI_BLANKS option" on page 550
Ansi_close_cursors_on_rollback	"ANSI_CLOSE_CURSORS_ON_ROLLBACK option" on page 551
Ansi_integer_overflow	"ANSI_INTEGER_OVERFLOW option" on page 551
Ansi_permissions	"ANSI_PERMISSIONS option" on page 551
Ansi_update_constraints	"ANSI_UPDATE_CONSTRAINTS option" on page 552
Ansinull	"ANSINULL option" on page 553
AppInfo	"AppInfo connection parameter" on page 165
Automatic_timestamp	"AUTOMATIC_TIMESTAMP option" on page 554
Background_priority	"BACKGROUND_PRIORITY option" on page 555
BlockedOn	If the current connection is not blocked, this is zero. If it is blocked, the connection number on which the connection is blocked due to a locking conflict.
Blocking	"BLOCKING option" on page 556
BytesReceived	The number of bytes received during client/server communications.
BytesReceivedUncomp	The number of bytes that would have been received during client/server communications if compression was disabled. (This value is the same as the value for BytesReceived if compression is disabled.)
BytesSent	The number of bytes sent during client/server communications.
BytesSentUncomp	The number of bytes that would have been sent during client/server communications if compression was disabled. (This value is the same as the value for BytesSent if compression is disabled.)
CacheHits	The number of successful reads of the cache.
CacheRead	The number of database pages that have been


Property	Description
	looked up in the cache.
CacheReadIndInt	The number of index internal-node pages that have been read from the cache.
CacheReadIndLeaf	The number of index leaf pages that have been read from the cache.
CacheReadTable	The number of table pages that have been read from the cache.
Chained	"CHAINED option" on page 557
CharSet	The character set used by the connection.
Checkpoint_time	"CHECKPOINT_TIME option" on page 557
CIS_option	Reserved
Cis_rowset_size	Reserved
ClientLibrary	jConnect for jConnect connections; CT_Library for Open Client connections; CmdSeq for ODBC and Embedded SQL connections;
Close_on_EndTrans	"CLOSE_ON_ENDTRANS option" on page 558
Commit	The number of Commit requests that have been handled.
CommLink	The communication link for the connection. This is one of the network protocols supported by Adaptive Server Anywhere, or is local for a same-machine connection.
CommProtocol	Returns CmdSeq for Adaptive Server Anywhere protocol (ODBC and Embedded SQL) or TDS for Open Client and jConnect connections.
Compression	Returns ON or OFF, to indicate whether communication compression is enabled on the connection. If a write file is created on a compressed database, the write file is NOT compressed. Starting a write file created on a compressed database and selecting db_property('compression'), returns OFF.
Continue_after_raisererror	"CONTINUE_AFTER_RAISERROR option" on page 560
Conversion_error	"CONVERSION_ERROR option" on page 560
Cooperative_commit_time out	"COOPERATIVE_COMMIT_TIMEOUT option" on page 561

Property	Description
Cooperative_commits	"COOPERATIVE_COMMITS option" on page 561
CurrTaskSwitch	The number of current request context switches.
Cursor	The number of declared cursors that are currently being maintained by the server.
CursorOpen	The number of open cursors that are currently being maintained by the server.
Database_authentication	The authentication string of the client's connected database.
Date_format	"DATE_FORMAT option" on page 562
Date_order	"DATE_ORDER option" on page 564
DBNumber	The ID number of the database.
Default_timestamp_increment	"DEFAULT_TIMESTAMP_INCREMENT option" on page 564
Delayed_commit_timeout	"DELAYED_COMMIT_TIMEOUT option" on page 564
Delayed_commits	"DELAYED_COMMITS option" on page 565
DiskRead	The number of pages that have been read from disk.
DiskReadIndInt	The number of index internal-node pages that have been read from disk.
DiskReadIndLeaf	The number of index leaf pages that have been read from disk.
DiskReadTable	The number of table pages that have been read from disk.
DiskWrite	The number of modified pages that have been written to disk.
Divide_by_zero_error	"DIVIDE_BY_ZERO_ERROR option" on page 566
Encryption	"Encryption connection parameter" on page 177
Escape_character	"ESCAPE_CHARACTER option" on page 567
EventName	The name of the associated event if the connection is running an event handler. Otherwise, the result is NULL.
Extended_join_syntax	"EXTENDED_JOIN_SYNTAX option" on page 567
Fire_triggers	"FIRE_TRIGGERS option" on page 568

Property	Description
Float_as_double	"FLOAT_AS_DOUBLE option" on page 569
FullCompare	The number of comparisons that have been performed beyond the hash value in an index.
IdleTimeout	The idle timeout value of the connection.  For more information, see "Idle connection parameter" on page 181.
IndAdd	The number of entries that have been added to indexes.
IndLookup	The number of entries that have been looked up in indexes.
Isolation_level	"ISOLATION_LEVEL option" on page 571
Java_heap_size	"JAVA_HEAP_SIZE option" on page 576
Java_input_output	"JAVA_INPUT_OUTPUT option" on page 577
Java_namespace_size	"JAVA_NAMESPACE_SIZE option" on page 577
Java_page_buffer_size	The page buffer size used by the Java VM.
JavaHeapSize	The heap size per Java VM.
Language	The locale language.
LastIdle	The number of ticks between requests.
LastReqTime	The time at which the last request for the specified connection started.
LastStatement	The most recently prepared SQL statement for the current connection.  For more information, see "-zl server option" on page 156.
LivenessTimeout	The liveness timeout period for the current connection.  For more information, see "LivenessTimeout connection parameter" on page 182.
Lock_rejected_rows	Reserved
LockName	A 64-bit unsigned integer value representing the lock for which a connection is waiting.
LogFreeCommit	The number of Redo Free Commits. A "Redo Free Commit" occurs when a commit of the transaction log is requested but the log has already been written (so the commit was done for "free").
Login_mode	"LOGIN_MODE option" on page 577

Property	Description
Login_procedure	"LOGIN_PROCEDURE option" on page 578
LogWrite	The number of pages that have been written to the transaction log.
Max_cursor_count	"MAX_CURSOR_COUNT option" on page 581
Max_plans_cached	"MAX_PLANS_CACHED option" on page 581
Max_statement_count	"MAX_STATEMENT_COUNT option" on page 583
Min_password_length	"MIN_PASSWORD_LENGTH option" on page 583
Min_table_size_for_histogram	"MIN_TABLE_SIZE_FOR_HISTOGRAM option" on page 583
Name	The name of the current connection.
Nearest_century	"NEAREST_CENTURY option" on page 584
NodeAddress	The node for the client in a client/server connection.
Non_keywords	"NON_KEYWORDS option" on page 584
Number	The ID number of the connection.
On_tsql_error	"ON_TSQL_ERROR option" on page 587
Optimization_goal	"OPTIMIZATION_GOAL option" on page 587
Optimization_level	Reserved
PacketsReceived	The number of client/server communication packets received.
PacketsReceivedUncomp	The number of packets that would have been received during client/server communications if compression was disabled. (This value is the same as the value for PacketsReceived if compression is disabled.)
PacketsSent	The number of client/server communication packets sent.
PacketsSentUncomp	The number of packets that would have been sent during client/server communications if compression was disabled. (This value is the same as the value for PacketsSent if compression is disabled.)
PacketSize	The packet size used by the connection, in bytes.
Percent_as_comment	"PERCENT_AS_COMMENT option" on page 590
Precision	"PRECISION option" on page 591

Property	Description
Prefetch	"PREFETCH option" on page 592
Prepares	The number of statement preparations carried out.
PrepStmt	The number of prepared statements currently being maintained by the server.
Query_plan_on_open	"QUERY_PLAN_ON_OPEN option" on page 594
QueryLowMemoryStrategy	The number of times the server changed its execution plan during execution as a result of low memory conditions. The strategy can change because less memory is available than the optimizer estimated, or because the execution plan required more memory than the optimizer estimated.
QueryCachePages	The number of pages used to cache execution plans.
Quoted_identifier	"QUOTED_IDENTIFIER option" on page 594
Recovery_time	"RECOVERY_TIME option" on page 595
Replicate_all	"REPLICATE_ALL option" on page 595
ReqType	A string for the type of the last request.
RI_trigger_time	"RI_TRIGGER_TIME option" on page 597
Rlbnk	The number of Rollback requests that have been handled.
RollbackLogPages	The number of pages in the rollback log.
Row_counts	"ROW_COUNTS option" on page 597
Scale	"SCALE option" on page 598
SQL_flagger_error_level	"SQL_FLAGGER_ERROR_LEVEL option" on page 599
SQL_flagger_warning_level	"SQL_FLAGGER_WARNING_LEVEL option" on page 599
String_rtruncation	Returns ON or OFF to indicate whether non-space characters are truncated on the assignment of string data.
Suppress_TDS_debugging	SUPPRESS_TDS_DEBUGGING option
TaskSwitch	The number of times the current server thread has been changed.
TDS_empty_string_is_null	TDS_EMPTY_STRING_IS_NULL option
TempTablePages	The number of pages in the temporary file used for

Property	Description
	temporary tables.
Time_format	"TIME_FORMAT option" on page 601
Timestamp_format	"TIMESTAMP_FORMAT option" on page 602
TimeZoneAdjustment	The number of minutes that must be added to the Coordinated Universal Time (UTC) to display time local to the connection. By default, the value is set according to the client's time zone.
Truncate_timestamp_values option	"TRUNCATE_TIMESTAMP_VALUES option" on page 604
Truncate_with_autocommit	"TRUNCATE_WITH_AUTO_COMMIT option" on page 605
Tsql_hex_constant	"TSQL_HEX_CONSTANT option" on page 606
Tsql_variables	"TSQL_VARIABLES option" on page 606
UncommitOp	The number of uncommitted operations.
User_estimates	USER_ESTIMATES option
Userid	The user ID for the connection.
UtilCmdsPermitted	Returns ON or OFF to indicate whether utility commands such as CREATE DATABASE, DROP DATABASE, and RESTORE DATABASE are permitted for the connection.  For more information, see "-gu server option" on page 143.
Wait_for_commit	"WAIT_FOR_COMMIT option" on page 608

Server-level properties

The following table lists properties that apply across the server as a whole.

Examples

❖ To retrieve the value of a server property:



- ◆ Use the property system function. For example, the following statement returns the number of cache pages being used to hold the main heap:


```
select property ( 'MainHeapPages' )
```

❖ To retrieve the values of all server properties:

- ◆ Use the *sa_eng_properties* system procedure:

call sa_eng_properties

Descriptions	Property	Description
	ActiveReq	The number of server threads that are currently handling a request.
	AvailIO	Reserved
	BuildChange	Reserved
	BuildClient	Reserved
	BuildReproducible	Reserved
	BytesReceived	The number of bytes received during client/server communications.
	BytesReceivedUncomp	The number of bytes that would have been received during client/server communications if compression was disabled. (This value is the same as the value for BytesReceived if compression is disabled.)
	BytesSent	The number of bytes sent during client/server communications.
	BytesSentUncomp	The number of bytes that would have been sent during client/server communications if compression was disabled. (This value is the same as the value for BytesSent if compression is disabled.)
	C2	Returns YES if the <code>-sc</code> option was used when the server was started. Otherwise, returns NO.  For more information, see "-sc server option" on page 150.
	CacheHitsEng	The number of database page lookups.
	CacheReplacements	The number of pages in the cache that have been replaced.
	CharSet	The character set in use by the database server.
	CompanyName	The name of the company owning this software.
	ConnsDisabled	Returns ON or OFF, to indicate the current setting of the server option to disable new connections.  For information, see "sa_server_option system procedure" on page 716 of the book <i>ASA SQL Reference Manual</i> .
	CurrentCacheSize	The current cache size, in kilobytes.
	DefaultCollation	The collation that would be used for new databases, if none is explicitly specified.
	IdleTimeout	The default idle timeout.

Property	Description
	 For more information, see "-ti server option" on page 150.
FreeBuffers	The number of available network buffers.
IsIQ	Returns TRUE if the server is an IQ server, and FALSE otherwise.
IsJavaAvailable	Returns YES if the JavaVM is installed, and NO if the JavaVM is not installed. This property only indicates if the Java VM is available, not whether it is currently being used.
IsNetworkServer	Returns YES if connected to a network database server, and NO if connected to a personal database server.
IsRuntimeServer	Returns YES if connected to the limited desktop runtime database server, and NO otherwise.
JavaGlobFix	Java VM global fixed size.
Language	The locale language for the server.
LegalCopyright	The copyright string for the software.
LegalTrademarks	Trademark information for the software.
LicenseCount	The number of licensed seats.
LicensedCompany	The name of the licensed company.
LicensedUser	The name of the licensed user.
LicenseType	The license type. Can be either <i>concurrent</i> or <i>networked seat (per-seat)</i> .
LivenessTimeout	The client liveness timeout default.
LockedHeapPages	The number of heap pages locked in the cache.
MachineName	The name or IP address of the computer running a database server.
MainHeapBytes	The number of bytes used for global server data structures.
MainHeapPages	The number of pages used for global server data structures.
MaxCacheSize	The maximum allowed cache size, in kilobytes.
MaxMessage	The current maximum line number that can be retrieved from the server's message window. This represents the most recent message displayed in the server's message window.
Message, <i>linenumber</i>	A line from the server's message window, prefixed by

Property	Description
	the date and time the message appeared. The second parameter specifies the line number.
MessageText , <i>linenumber</i>	The text associated with the specified line number in the server's message window, without a date and time prefix. The second parameter specifies the line number.
MessageTime , <i>linenumber</i>	The date and time associated with the specified line number in the server's message window. The second parameter specifies the line number.
MessageWindowSize	The maximum number of lines that can be retrieved from the server's message window.
MinCacheSize	The minimum allowed cache size, in kilobytes.
MultiPacketsReceived	The number of multi-packet deliveries received during client/server communications.
MultiPacketsSent	The number of multi-packet deliveries sent during client/server communications.
Name	The name of the server.
NumProcessorsAvail	The number of processors on the server.
NumProcessorsMax	The maximum number of processors used. Normally this should be 2 for <i>dbeng.exe</i> and 0 for <i>dbsrv.exe</i> .
PacketsReceived	The number of client/server communication packets received.
PacketsReceivedUncomp	The number of packets that would have been received during client/server communications if compression was disabled. (This value is the same as the value for PacketsReceived if compression is disabled.)
PacketsSent	The number of client/server communication packets sent.
PacketsSentUncomp	The number of packets that would have been sent during client/server communications if compression was disabled. (This value is the same as the value for PacketsSent if compression is disabled.)
PageSize	The size of the database server cache pages. This can be set using the <code>-gp</code> option, otherwise, it is the maximum database page size of the databases specified on the command line.
PeakCacheSize	The largest value the cache has reached in the current session, in kilobytes.
Platform	The operating system on which the software is running.
PlatformVer	The operating system on which the software is running, including build numbers, service packs, etc.

Property	Description
ProcessCPU	CPU usage statistics for the server process. Values are in microseconds.
ProcessCPUSystem	Process CPU system usage.
ProcessCPUUser	Process CPU user usage.
ProductName	The name of the software.
ProductVersion	The version of the software being run.
QuittingTime	Shutdown time for the server. If none is specified, the value is none .
Req	The number of times the server has been entered to allow it to handle a new request or continue processing an existing request.
RequestLogFile	The name of the request-level logging filename. An empty string is returned if there is no request level logging. 🔗 For information, see "sa_server_option system procedure" on page 716 of the book <i>ASA SQL Reference Manual</i> .
RequestLogging	ALL, SQL, or NONE. 🔗 For information, see "sa_server_option system procedure" on page 716 of the book <i>ASA SQL Reference Manual</i> .
RequestQueueWait	The number of times the server has had to wait for room in the request queue.
SendFail	The number of times that the underlying communications protocols have failed to send a packet.
ServerIdleWaits	The number of times the server has gone idle waiting for I/O completion or a new request.
StartTime	The date/time that the server started
Tempdir	The directory in which temporary files are stored by the server.
TimeZoneAdjustment	The number of minutes that must be added to the Coordinated Universal Time (UTC) to display time local to the server.
TotalBuffers	The total number of network buffers.
UnschReq	The number of requests that are currently queued up waiting for an available server thread.

Database-level properties

The following table lists properties available for each database on the server.

Examples

- ❖ **To retrieve the value of a database property:**
 - ◆ Use the *db_property* system function. For example, the following statement returns the page size of the current database:

```
select db_property ( 'PageSize' )
```

- ❖ **To retrieve the values of all database properties:**
 - ◆ Use the *sa_db_properties* system procedure:

```
call sa_db_properties
```

Descriptions

Property	Description
Alias	The database name.
BlankPadding	The status of the blank padding feature. Returns ON if the database has blank padding enabled. Otherwise, it returns OFF.
BlobArenas	The status of the BlobArenas feature. Returns ON if the database stores extension (blob) pages separately from table pages for the database. Otherwise, returns OFF.
CacheHits	The number of database page lookups satisfied by finding the page in the cache.
CacheRead	The number of database pages that have been looked up in the cache.
CacheReadIndInt	The number of index internal-node pages that have been read from the cache.
CacheReadIndLeaf	The number of index leaf pages that have been read from the cache.
CacheReadTable	The number of table pages that have been read from the cache.
Capabilities	The capability bits enabled for the database. This property is primarily for use by technical support.
CaseSensitive	The status of the case sensitivity feature. Returns ON if the database is case sensitive. Otherwise, it returns OFF
CharSet	The character set of the database.
CheckpointUrgency	The time that has elapsed since the last checkpoint as a percentage of the checkpoint time setting of the

Property	Description
	database.
Chkpt	The number of checkpoints that have been performed.
ChkptFlush	The number of ranges of adjacent pages written out during a checkpoint.
ChkptPage	The number of transaction log checkpoints.
CommitFile	The number of times the server has forced a flush of the disk cache. On Windows NT/2000/XP and NetWare platforms, the disk cache does not need to be flushed if unbuffered (direct) I/O is used.
CompressedBTrees	Returns ON if Compressed B-tree indexes are supported. Otherwise, returns OFF.
Compression	The compression status of the database. Returns either ON (meaning the database is compressed) or OFF.
ConnCount	The number of connections to the database.
CurrentRedoPos	The current offset in the transaction log file where the next database operation is to be logged.
CurrIO	The current number of file I/Os that were issued by the server but have not yet completed.
CurrRead	The current number of file reads that were issued by the server but have not yet completed.
CurrWrite	The current number of file writes that were issued by the server but have not yet completed.
DBFileFragments	The number of database file fragments. This property is supported on Windows NT/2000/XP.
DiskRead	The number of pages that have been read from disk.
DiskReadIndInt	The number of index internal-node pages that have been read from disk.
DiskReadIndLeaf	The number of index leaf pages that have been read from disk.
DiskReadTable	The number of table pages that have been read from disk.
DiskWrite	The number of modified pages that have been written to disk.
Encryption	The type of encryption applied to the database. Returns None, Simple, AES, or MDSR.
ExtendDB	The number of pages by which the database file has been extended.

Property	Description
ExtendTempWrite	The number of pages by which temporary files have been extended.
File	The file name of the database root file, including path.
FileVersion	The version of the database file. This does not correspond to a software release version.
FreePageBitMaps	Returns ON if free database pages are managed via bitmaps. Otherwise, returns OFF.
FullCompare	The number of comparisons that have been performed beyond the hash value in an index.
GlobalDBId	The value of the GLOBAL_DATABASE_ID option used to generate unique primary key values in a replication environment.
Histograms	Returns ON if optimizer statistics are maintained as histograms. Otherwise, returns OFF.
IdleCheck	The number of times that the server's idle thread has become active to do idle writes, idle checkpoints, and so on.
IdleChkpt	The number of checkpoints completed by the server's idle thread. An idle checkpoint occurs whenever the idle thread writes out the last dirty page in the cache.
IdleChkTime	The number of 100ths of a second spent checkpointing during idle I/O.
IdleWrite	The number of disk writes that have been issued by the server's idle thread.
IndAdd	The number of entries that have been added to indexes.
IndLookup	The number of entries that have been looked up in indexes.
IOToRecover	The estimated number of I/O operations required to recover the database.
IQStore	Reserved.
JavaHeapSize	Heap size per Java VM.
JavaNSSize	Java VM Namespace size.
JDKVersion	The Java runtime library version used by this database.
Language	The locale language of the database.
LargeProcedureIDs	Returns ON if 32-bit stored procedure IDs are supported for the database. Otherwise, returns OFF.

Property	Description
LockTablePages	The number of pages used to store lock information.
LogFileFragments	The number of log file fragments. This property is supported on Windows NT/2000/XP.
LogFreeCommit	The number of Redo Free Commits. A "Redo Free Commit" occurs when a commit of the transaction log is requested but the log has already been written (so the commit was done for "free").
LogName	The file name of the transaction log, including path.
LogWrite	The number of pages that have been written to the transaction log.
LTMGeneration	The generation number of the LTM, or Replication Agent. This property is primarily for use by technical support.
LTMTrunc	The minimal confirmed log offset for the Replication Agent.
MapPages	The number of map pages used for accessing the lock table, frequency table, and table layout.
MaxIO	The maximum value that CurrIO has reached.
MaxRead	The maximum value that CurrRead has reached.
MaxWrite	The maximum value that CurrWrite has reached.
MultiByteCharSet	Returns ON if the database uses a multi-byte character set. Otherwise, returns OFF.
Name	The database name (identical to alias).
PageRelocations	The number of relocatable heap pages that have been read from the temporary file.
PageSize	The page size of the database, in bytes.
PreserveSource	Returns ON if the database preserves the source for procedures and views. Otherwise, returns OFF.
ProcedurePages	The number of relocatable heap pages that have been used for procedures.
QueryCachePages	The number of pages used to cache execution plans.
QueryLowMemoryStrategy	The number of times the server changed its execution plan during execution as a result of low memory conditions. The strategy can change because less memory is available than the optimizer estimated, or because the execution plan required more memory than the optimizer estimated.
ReadOnly	Returns ON if the database is being run in read-only

Property	Description
	mode. Otherwise, returns OFF.
RecoveryUrgency	An estimate of the amount of time required to recover the database.
RelocatableHeapPages	The number of pages used for relocatable heaps (cursors, statements, procedures, triggers, views, etc.).
RemoteTrunc	The minimal confirmed log offset for the SQL Remote Message Agent.
RollbackLogPages	The number of pages in the rollback log.
SeparateCheckpointLog	Returns ON if the checkpoint log for the database is maintained at the end of the SYSTEM dbspace. Otherwise, returns OFF.
SeparateForeignKeys	Returns ON if primary and foreign keys are stored separately. Otherwise, returns OFF.
SyncTrunc	The minimal confirmed log offset for the MobiLink client <i>dbmlsync</i> executable.
TableBitMaps	Returns ON if the database supports table bitmaps. Otherwise, returns OFF.
TempFileName	The named of the database temporary file.
TempTablePages	The number of pages in the temporary file used for temporary tables.
TransactionsSpanLogs	Returns ON if transactions can span multiple log files. Otherwise, returns OFF.
TriggerPages	The number of relocatable heap pages used for triggers.
VariableHashSize	Returns ON if the hash size can be specified for B-tree indexes. Otherwise, returns OFF.
ViewPages	The number of relocatable heap pages used for views.

C H A P T E R 1 8

Physical Limitations

About this chapter This chapter describes the limitations on size and number of objects in Adaptive Server Anywhere databases.

Contents	<table><tr><th>Topic</th><th>Page</th></tr><tr><td>Size and number limitations</td><td>636</td></tr></table>	Topic	Page	Size and number limitations	636
Topic	Page				
Size and number limitations	636				

Size and number limitations

The following table lists the physical limitations on size and number of objects in Adaptive Server Anywhere database. The memory, CPU, and disk drive of the computer are more limiting factors in most cases.

Item	Limitation
Database size	13 files per database. For each file, the largest file allowed by operating system and file system
Field size	2 Gb
File size (FAT 12)	16 Mb
File size (FAT 16)	2 Gb
File size (FAT 32)	4 Gb
File size (NTFS, HP-UX 11.0 and later, Solaris 2.6 and later)	<ul style="list-style-type: none">◆ 256 Gb for 1 kb pages◆ 512 Gb for 2 kb pages◆ 1 Tb for 4 kb pages
File size (NetWare)	4 Gb
File size (all other platforms and file systems)	2 Gb
Maximum cache size	<ul style="list-style-type: none">◆ Windows 95/98 1.5 Gb◆ Windows NT Advanced Server 1.5 Gb◆ Windows NT 2.5 Gb◆ Windows 2000 Professional (non-AWE cache) 1.6 Gb◆ Windows 2000 Server (non-AWE cache) 2.3 Gb◆ Windows 2000 Advanced Server (non-AWE cache) 2.3 Gb◆ Windows 2000 Datacenter Server (non-AWE cache) 2.3 Gb◆ Windows XP Home Edition (non-AWE cache) 1.6 Gb◆ Windows XP Professional (non-AWE cache) 1.6 Gb◆ Windows .NET Web Server*

Item	Limitation
	(non-AWE cache) 1.6 Gb
	◆ Windows .NET Standard Server* (non-AWE cache) 1.6 Gb
	◆ Windows .NET Enterprise Server* (non-AWE cache) 2.3 Gb
	◆ Windows .NET Datacenter Server* (non-AWE cache) 2.3 Gb
	◆ Windows CE 32Mb
	◆ UNIX 4 Gb
	◆ NetWare 2 Gb
Maximum index entry size	No limit
Number of databases per server	255
Number of columns per table	$((\text{page size})/4)^2$. Note: An excessive number of columns, although allowed, will affect performance.
Number of indexes	The maximum number of indexes per table is based on the database page size: ◆ 4 kb or less 2048 indexes ◆ 8 kb 1024 indexes ◆ 16 kb 512 indexes ◆ 32 kb 256 indexes
Number of rows per database	Limited by file size
Number of rows per table	Limited by file size
Number of tables per database	$2^{32} - 1 = 4\,294\,967\,295$
Number of temporary tables per connection	$2^{32} - 1 = 4\,294\,967\,295$
Number of tables referenced per transaction	No limit
Number of stored procedures per database	$2^{32} - 1 = 4\,294\,967\,295$
Number of events per database	$2^{31} - 1 = 2\,147\,483\,647$
Number of triggers per database	$2^{32} - 1 = 4\,294\,967\,295$
Row size	Limited by file size
Table size	Maximum file size. User-created

Size and number limitations

Item	Limitation
	indexes for the table can be stored separately from the table

* At the time of printing, Windows .NET Web Server, Windows .NET Standard Server, Windows .NET Enterprise Server, and Windows .NET Datacenter Server had not been released.

Index

%

% comment indicator, 590

% operator
 modulo function, 590

&

&
 UNIX command line, 21

?

-? server option
 database server, 127

@

@environment-variable server option
 database server, 125

@filename option
 database server, 10

@filename server option
 database server, 126

A

-a option
 database server, 157
 Log Translation [dbtran] utility, 488, 490

ping [dbping] utility, 494
service creation [dbsvc] utility, 499, 500

-A option
 Log Transfer Manger [dbltm] utility, 481, 482

-ac option
 unload [dbunload] utility, 514, 516

access
 security features, 388

ActiveReq property
 description, 626

Adaptive Server Anywhere
 connecting using 8.0 sample data source, 48
 setting up as an Open Server, 108

Adaptive Server Anywhere as an Open Server, 105

Address Windowing Extensions
 setting cache size, 131

ADO
 connecting, 69

AES encryption algorithm
 about, 399
 initialization [dbinit] utility, 468
 unload [dbunload] utility, 518

Alias property
 description, 630

ALL
 permissions, 359

ALLOW_NULLS_BY_DEFAULT option
 compatibility option, 544, 545
 description, 550
 obtaining value, 619
 Open Client, 116

alphabetic characters
 defined, 286

- ALTER permissions, 359
- ALTER PROCEDURE statement
 - about, 424
 - effects, 425
- ALTER TABLE statement
 - REPLICATE ON, 413
- an option
 - unload [dbunload] utility, 514, 517
- ANSI
 - about code pages, 255
 - choosing code pages, 274
 - COMMIT behavior, 561, 565
 - conformance, 599
 - cursors, 551
 - delete permissions, 551
 - update permissions, 551
 - variable behavior, 550
- ANSI_BLANKS option
 - compatibility option, 544, 545
 - description, 550
 - obtaining value, 619
 - Open Client, 116
- ANSI_CLOSE_CURSORS_ON_ROLLBACK
 - option
 - compatibility option, 545
 - description, 551
 - obtaining value, 619
- ANSI_INTEGER_OVERFLOW option
 - compatibility option, 545
 - description, 551
 - obtaining value, 619
 - Open Client, 116
- ANSI_PERMISSIONS option
 - compatibility option, 545
 - description, 551
 - obtaining value, 619
- ANSI_UPDATE_CONSTRAINTS option
 - compatibility option, 545
 - description, 552
 - obtaining value, 619
- ANSINULL option
 - compatibility option, 544, 545
 - description, 553
 - obtaining value, 619
 - Open Client, 116
- APC_pw parameter, 417
 - LTM configuration file, 484
- APC_user parameter, 417, 425
 - LTM configuration file, 484
- APCs
 - function APCs, 424
 - Replication Server, 407
- AppInfo [APP] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 165
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- AppInfo property
 - description, 619
- ar option
 - unload [dbunload] utility, 497, 514, 517
- as option
 - service creation [dbsvc] utility, 499, 501
- asademo.db file
 - about, xvi
- asany.ini
 - location, 207
- ASANY8 environment variable
 - description, 209
- ASANYSH8 environment variable
 - description, 209
- ASAProv
 - OLE DB provider, 68
- asasrv.ini
 - troubleshooting server startup, 35
- asasrv.ini file
 - server information, 81
- ASCII
 - character set, 254, 255
 - collating sequences, 471
 - file format, 570, 588
- ASTMP environment variable
 - description, 209
 - UNIX, 212

- asynchronous procedures
 - about, 425
 - Replication Server, 407
 - user ID for, 484
 - auditing
 - about, 393
 - comments, 395
 - controlling, 553
 - example, 395
 - log translation [dbtran] utility, 396
 - recovering uncommitted operations, 344
 - retrieving audit information, 394
 - security features, 388, 393
 - transaction log [dblog] utility, 396
 - turning on, 393
 - write file [dbwrite] utility, 396
 - AUDITING option
 - description, 553
 - Log Translation [dbtran] utility, 491
 - setting, 541
 - AUTO_COMMIT option
 - description, 554
 - Interactive SQL settings, 548
 - AUTO_REFETCH option
 - description, 554
 - Interactive SQL settings, 548
 - autocommit
 - ODBC configuration, 57
 - autoexec.ncf
 - automatic loading, 6
 - AUTOMATIC_TIMESTAMP option
 - compatibility option, 545
 - description, 554
 - obtaining value, 619
 - Open Client, 116
 - automation
 - administration tasks, 231
 - AutoStart [ASTART] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 166
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
 - AutoStop [ASTOP] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 167
 - introduction to connection parameters, 39
 - ODBC configuration, 59
 - tips, 72
 - using default connection parameters, 51
 - Avail IO statistic, 617
 - availability
 - database servers, 21
 - high, 340
 - AvailIO property
 - description, 626
 - AWE cache
 - setting cache size, 131
- ## B
- b option
 - data source [dbdsn] utility, 451, 454
 - database server, 127
 - initialization [dbinit] utility, 466, 467
 - background
 - running database servers, 21
 - BACKGROUND_PRIORITY option
 - description, 555
 - obtaining value, 619
 - setting, 541
 - backup and data recovery, 299
 - backup database wizard
 - using, 339
 - backup directory
 - backup [dbbackup] utility, 440
 - backup plans
 - about, 315
 - backup utility
 - about, 438
 - options, 440
 - syntax, 439
 - backups
 - about, 300
 - backup utility, 438

- concepts, 305
 - database only, 440
 - databases not involved in replication, 311
 - dbltm and, 312
 - dbmlsync and, 312
 - dbremote and, 312
 - designing procedures, 308
 - external, 303
 - for remote databases, 314
 - full, 329
 - internal, 303
 - internals of, 321
 - live, 340, 440
 - MobiLink and, 312
 - MobiLink consolidated databases, 311
 - offline, 303
 - online, 303
 - options, 440
 - planning, 309
 - rename and start new transaction log, 440
 - replication, 429, 430
 - Replication Agent and, 312
 - restrictions during, 322
 - scheduling, 309
 - SQL Remote and, 312
 - SQL statement, 303
 - strategies, 309
 - Sybase Central, 303
 - types of, 303
 - unfinished, 337
 - validating, 315
- batch files
- starting servers, 503
- batch mode
- for LTM, 426
- batch_ltl_cmds parameter
- LTM configuration file, 484
- batch_ltl_mem parameter
- LTM configuration file, 484
- batch_ltl_sz parameter
- LTM configuration file, 484
- BELL option
- description, 555
 - Interactive SQL settings, 548
- bindery
- NetWare, 96
- blank padding
- initialization [dbinit] utility, 467
- BlankPadding property
- description, 630
- blanks
- ANSI behavior, 550
- BLOB_THRESHOLD option
- description, 555
 - replication option, 547
- BlobArenas property
- description, 630
- BlockedOn property
- description, 619
- BLOCKING option
- description, 556
 - obtaining value, 619
 - setting, 541
- BLOCKING_TIMEOUT option
- description, 556
 - setting, 541
- Broadcast [BCAST] communication parameter
- description, 190
- BroadcastListener [BLISTENER] communication parameter
- description, 190
- buffer size
- ODBC configuration, 61
- buffering
- replication commands, 426
- BufferMisses property
- description, 626
- BuildChange property
- description, 626
- BuildClient property
- description, 626
- BuildReproducible property
- description, 626
- bulk loading
- option, 13
- bulk operations
- b server option, 127

BytesReceived property
description, 619, 626

BytesReceivedUncomp property
description, 619, 626

BytesSent property
description, 619, 626

BytesSentUncomp property
description, 619, 626

C

–c option

backup [dbbackup] utility, 439, 440
collation [dbcollat] utility, 443
console [dbconsole] utility, 450
data source [dbdsn] utility, 451, 454
database server, 127
debugger [dbprdbg] utility, 456, 457
histogram [dbhist] utility, 461, 462
information [dbinfo] utility, 463
initialization [dbinit] utility, 466, 467
Interactive SQL [dbisql] utility, 472, 474
Log Translation [dbtran] utility, 488, 490
ping [dbping] utility, 494
stop [dbstop] utility, 505, 506
unload [dbunload] utility, 514, 518
upgrade [dbupgrad] utility, 522, 523
validation [dbvalid] utility, 527, 528
write file [dbwrite] utility, 531, 532

–C option

Log Transfer Manger [dbltn] utility, 481, 482

C2 property
description, 626

C2 security
guidelines, 388

–ca option

database server, 129

cache

database server options, 124
maximum size, 636
size, 124
size option, 12

cache buffers
performance, 127, 129, 130, 131

Cache Hits/sec statistic, 611

Cache Reads Index Interior/sec statistic, 611

Cache Reads Index Leaf/sec statistic, 611

Cache Reads Table/sec statistic, 611

Cache Reads Total Pages/sec statistic, 611

cache size

encrypted database issues, 402
limit, 636
limiting, 130, 131
setting, 127
setting for Windows .NET Server, 131
setting for Windows 2000, 131
setting for Windows XP, 131
static, 129

Cache Size Current statistic, 611

Cache Size Maximum statistic, 611

Cache Size Minimum statistic, 611

Cache Size Peak statistic, 611

cache statistics, 611

CacheHits property
description, 619, 630

CacheHitsEng property
description, 626

CacheRead property
description, 619, 630

CacheReadIndInt property
description, 619, 630

CacheReadIndLeaf property
description, 619, 630

CacheReadTable property
description, 619, 630

CacheReplacements property
description, 626

caching
server names, 80

cannot find database server, 76

Capabilities property
description, 630

capability ID, 630

- case sensitivity, 466
 - collations, 285
 - command line, 9
 - connection parameters, 71, 164
 - database name, 11
 - initialization [dbinit] utility, 467
 - international aspects, 261
 - server name, 11
- CaseSensitive property
 - description, 630
- CD-ROM
 - databases, 224
 - deployment, 149
- Certicom
 - encrypting client/server communications, 135
 - security, 135
- ch option
 - database server, 130
- CHAINED option
 - compatibility option, 544, 545
 - description, 557
 - obtaining value, 619
 - Open Client, 116
- chained transaction mode, 557
- change log file settings wizard
 - changing the location of a transaction log, 345
 - starting a transaction log mirror for an existing database, 347
 - using, 507
- changing
 - collations, 294
- CHAR data type
 - and host variables, 550
- character LTM, 427
- character set conversion
 - enabling, 131
- character set translation
 - about, 291
 - disabling, 292
 - enabling, 131, 291
 - error messages, 278
- character sets
 - about, 250
 - application, 264
 - avoiding translation, 280
 - choosing, 287
 - connection parameter, 167
 - definition, 253
 - determining, 264
 - encoding, 250
 - fixed width, 259
 - LTM, 427, 428
 - multibyte, 259, 277
 - Replication Server, 421
 - server, 264
 - single-byte, 254
 - translation, 291, 292
 - Unicode, 277
 - UNIX default, 264
 - variable width, 259
 - Windows, 255
 - Windows default, 264
- characters
 - alphabetic, 286
 - digits, 286
 - white space, 286
- CharSet [CS] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 167
 - introduction to connection parameters, 39
 - overriding client character set with, 295
 - tips, 72
 - using default connection parameters, 51
- CharSet property
 - description, 619, 626, 630
- Checkpoint Flushes/sec statistic, 611
- Checkpoint Log statistic, 611
- checkpoint logs
 - about, 322
- checkpoint statistics, 611
- Checkpoint Urgency statistic, 611
- CHECKPOINT_TIME option
 - description, 557
 - obtaining value, 619
 - setting, 541
 - using, 325

- checkpoints
 - backups, 305
 - CHECKPOINT_TIME option, 557
 - interval between, 138
 - log, 322
 - scheduling, 326
 - urgency, 325
- Checkpoints/sec statistic, 611
- CheckpointUrgency property
 - description, 630
- Chkpt property
 - description, 630
- ChkptFlush property
 - description, 630
- ChkptPage property
 - description, 630
- choosing drivers, 43
- CIS_option property
 - description, 619
- CIS_ROWSET_SIZE option
 - description, 557
 - setting, 541
- Cis_rowset_size property
 - description, 619
- cl option
 - data source [dbdsn] utility, 454
- cl option
 - data source [dbdsn] utility, 451
 - database server, 130
- client liveness timeout
 - data source [dbdsn] utility, 455
- client side
 - setting encryption parameters, 177, 179
- ClientLibrary property
 - description, 619
- ClientPort [CPORT] communication parameter
 - description, 191
- clients
 - identifying, 165
- CLOSE_ON_ENDTRANS option
 - compatibility option, 545
 - description, 558
 - obtaining value, 619
- code pages
 - ANSI, 255
 - definition, 253
 - Interactive SQL [dbisql] utility, 474
 - list, 471
 - OEM, 255
 - overview, 254
 - Windows, 255
- codepage option
 - Interactive SQL [dbisql] utility, 472, 474
- collation file
 - editing, 282
- collation sequence label
 - collation [dbcollat] utility, 445
- collation sequences
 - collation utility, 442
 - dbcollat command line syntax, 443
 - initialization [dbinit] utility, 470
 - list, 471
- collation utility
 - about, 442
 - custom collations, 292, 293
 - options, 443
 - syntax, 443
- collations
 - about, 250
 - changing, 294
 - choosing, 287
 - creating, 292
 - custom, 282, 292, 293
 - default, 226
 - definition, 253
 - extracting using Sybase Central, 442
 - file format, 282
 - internals, 282
 - ISO_1, 274
 - list of, 269
 - LTM, 427, 428
 - multibyte, 277
 - OEM, 275
 - options, 443
 - recommended for UNIX, 271
 - recommended for Windows, 271

- Replication Server, 421
 - sorting characters with, 260
 - supplied, 269
- column names
 - international aspects, 261
- column permissions, 359
- columns
 - limitations, 636
 - permissions, 359
- Comm Buffer Misses statistic, 612
- Comm Bytes Received Uncompressed/sec statistic, 612
- Comm Bytes Sent Uncompressed/sec statistic, 612
- Comm Bytes Sent/sec statistic, 612
- Comm Free Buffers statistic, 612
- Comm Multi-packets Received/sec statistic, 612
- Comm Multi-packets Sent/sec statistic, 612
- Comm Packets Received Uncompressed/sec statistic, 612
- Comm Packets Received/sec statistic, 612
- Comm Packets Sent Uncompressed/sec statistic, 612
- Comm Packets Sent/sec statistic, 612
- Comm Send Fails/sec statistic, 612
- Comm TotalBuffers statistic, 612
- comma delimited files, 570, 588
- command delimiter
 - Interactive SQL [dbisql] utility, 474
- command echo, 567
- command file
 - option, 126
- command files
 - adding simple encryption with dbfhide, 446
- command line
 - case sensitivity, 9
 - database server, 120
 - in configuration file, 126
 - options, 10, 126
 - server, 126
 - starting the server, 8
- command line utilities
 - backup [dbbackup] syntax, 439
 - collation [dbcollat] syntax, 443
 - compression [dbshrink] syntax, 448
 - console [dbconsole] syntax, 450
 - data source [dbdsn] syntax, 451
 - debugger [dbprdbg] syntax, 456
 - erase [dberase] syntax, 459
 - file hiding [dbfhide] syntax, 446
 - histogram [dbhist] syntax, 461
 - information [dbinfo] syntax, 463
 - initialization [dbinit] syntax, 466
 - Interactive SQL [dbisql] syntax, 472
 - language [dblang] syntax, 477
 - license [dblic] syntax, 479
 - log translation [dbtran] syntax, 488
 - log transfer manager [dbltm] syntax, 481
 - ping [dbping] syntax, 494
 - rebuild syntax, 497
 - server location [dblocate] syntax, 498
 - service creation [dbsvc] syntax, 499
 - spawn [dbspawn] syntax, 503
 - stop [dbstop] syntax, 505
 - transaction log [dblog] syntax, 507
 - uncompression [dbexpand] syntax, 511
 - unload [dbunload] syntax, 514
 - upgrade [dbupgrad] syntax, 522
 - validation [dbvalid] syntax, 527
 - write file [dbwrite] syntax, 531
- COMMAND_DELIMITER option
 - description, 558
 - Interactive SQL settings, 548
- CommBufferSize [CBSIZE] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 168
 - introduction to connection parameters, 39
 - TCP/IP, 93
 - tips, 72
 - using default connection parameters, 51
- comments, 395
 - comment delimiters, 590
- Commit property
 - description, 619

-
- COMMIT statement
 - and AUTO_COMMIT option, 554
 - LTM, 418
 - COMMIT_ON_EXIT option
 - description, 559
 - Interactive SQL settings, 548
 - CommitFile property
 - description, 630
 - commits
 - COOPERATIVE_COMMIT_TIMEOUT option, 561
 - COOPERATIVE_COMMITS option, 561
 - DELAYED_COMMIT_TIMEOUT option, 564
 - DELAYED_COMMITS option, 565
 - CommLink property
 - description, 619
 - CommLinks [LINKS] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 169
 - introduction to connection parameters, 39
 - options, 16
 - parentheses, 279
 - tips, 72
 - using default connection parameters, 51
 - CommProtocol property
 - description, 619
 - communication compression, 170
 - compression threshold, 172
 - communication parameters
 - about, 163
 - Broadcast [BCAST], 190
 - BroadcastListener [BLISTENER], 190
 - ClientPort [CPORT], 191
 - DLL, 192
 - DoBroadcast [DOBROAD], 192
 - ExtendedName [ENAME], 193
 - HOST [IP], 194
 - list of, 189
 - LocalOnly [LOCAL], 195
 - MyIP [ME], 195
 - PreFetchOnOpen, 196
 - ReceiveBufferSize [RCVBUFSZ], 197
 - RegisterBindery [REGBIN], 197
 - SearchBindery [BINSEARCH], 197
 - SendBufferSize [SNDBUFSZ], 197
 - ServerPort [PORT], 198
 - TDS, 198
 - TIMEOUT [TO], 199
 - VerifyServerName [VERIFY], 199
 - communications
 - about, 91
 - database server, 153
 - debugging, 156
 - encrypting, 135, 137, 159, 177, 179
 - parameters, 189
 - supported, 92
 - troubleshooting, 101
 - communications statistics, 612
 - company names
 - license [dblic] utility, 480
 - CompanyName property
 - description, 626
 - comparing
 - TIMESTAMP, 564
 - comparisons
 - DATE data type, 604
 - compatibility
 - ANSI, 544
 - SQL, 544
 - Transact-SQL, 544
 - compatibility options
 - ALLOW_NULLS_BY_DEFAULT, 550
 - ANSI_BLANKS, 550
 - ANSI_CLOSE_CURSORS_ON_ROLLBACK, 551
 - ANSI_INTEGER_OVERFLOW, 551
 - ANSI_PERMISSIONS, 551
 - ANSI_UPDATE_CONSTRAINTS, 552
 - ANSINULL, 553
 - AUTOMATIC_TIMESTAMP, 554
 - CHAINED, 557
 - classification, 539
 - CLOSE_ON_ENDTRANS, 558
 - CONTINUE_AFTER_RAISERROR, 560
 - CONVERSION_ERROR, 560
 - DATE_FORMAT, 562
 - DATE_ORDER, 564
 - DIVIDE_BY_ZERO_ERROR, 566
 - ESCAPE_CHARACTER, 567
 - FIRE_TRIGGERS, 568
 - FLOAT_AS_DOUBLE, 569

- initial settings, 539
- ISOLATION_LEVEL, 571
- list of, 544, 545
- NON_KEYWORDS, 584
- ON_TSQL_ERROR, 587
- PERCENT_AS_COMMENT, 590
- QUERY_PLAN_ON_OPEN, 594
- QUOTED_IDENTIFIER, 594
- RI_TRIGGER_TIME, 597
- SQL_FLAGGER_ERROR_LEVEL, 599
- SQL_FLAGGER_WARNING_LEVEL, 599
- STRING_RTRUNCATION, 600
- TIME_FORMAT, 601
- TIMESTAMP_FORMAT, 602
- TSQL_HEX_CONSTANT, 606
- TSQL_VARIABLES, 606
- Compress [COMP] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 170
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- compress database wizard
 - using, 448
- compressed databases, 512
 - compressing with dbshrink, 448
 - creating from Sybase Central, 448, 456
 - uncompressing, 511
- CompressedBTrees property
 - description, 630
- compressing
 - database files with dbshrink, 448
 - packets, 147
- compression
 - compression [dbshrink] utility, 448
 - options, 449
 - performance, 98
- COMPRESSION option
 - description, 559
 - replication option, 547
- Compression property
 - description, 619, 630
- compression utility
 - about, 448
 - options, 449
 - syntax, 448
- CompressionThreshold [COMPPTH] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 172
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- configuration file
 - creating LTM, 425
 - database server, 10
 - for the LTM, 483
 - format for LTM, 425
 - Log Transfer Manger [dbltm] utility, 482
 - LTM, 417, 425
 - LTM command line, 482
 - option, 126
- configuration files
 - adding simple encryption with dbfhide, 446
 - hiding, 446
- configuring
 - Adaptive Server Anywhere for Replication Server, 419, 421
 - interfaces file, 110
 - LTM, 425
 - ODBC data sources, 56
 - SQL.ini, 110
 - using dbconsole, 450
- ConnCount property
 - description, 630
- Connect dialog
 - accessing, 42
 - overview, 44
- connected users, 368
 - managing, 368
- connecting
 - ADO, 69
 - character sets, 279
 - Connect dialog overview, 44
 - database connection scenarios, 45
 - firewalls, 93, 190, 191, 194, 198
 - from Interactive SQL, 42

- from Sybase Central, 42
 - integrated logins, 390
 - OLE DB, 68
 - permission, 357
 - RAS, 94
 - starting a local server, 46
 - to a local database, 46
 - using a data source, 48
 - utility database, 227
 - Windows CE, 63, 65, 66
- connecting to a database, 37
- connection name
- ODBC configuration, 62
- connection parameters
- about, 163
 - and connection strings, 40
 - AppInfo [APP], 165
 - AutoStart [ASTART], 166
 - AutoStop [ASTOP], 167
 - case sensitivity, 71, 164
 - CharSet [CS], 167
 - collation [dbcollat] utility, 443
 - CommBufferSize [CBSIZE], 168
 - CommLinks [LINKS], 169
 - Compress [COMP], 170
 - Compression Threshold [COMPTH], 172
 - conflicts, 72
 - ConnectionName [CON], 172
 - console [dbconsole] utility, 450
 - data source [dbdsn] utility, 454
 - data sources, 53
 - DatabaseFile [DBF], 173
 - DatabaseName [DBN], 174
 - DatabaseSwitches [DBS], 174
 - DataSourceName [DSN], 175
 - dbbackup command-line utility, 440
 - debugger [dbprdbg] utility, 457
 - DisableMultiRowFetch [DMRF], 175
 - embedded databases, 72
 - EncryptedPassword [ENP], 177
 - Encryption [ENC], 177
 - Encryption Key [DBKEY], 179
 - EngineName [ENG], 176
 - establishing connections, 38
 - FileDataSourceName [FILEDSN], 180
 - ForceStart [FORCESTART], 180
 - Idle [IDLE], 181
 - information [dbinfo] utility, 463
 - Integrated [INT], 181
 - Interactive SQL [dbisql] utility, 474
 - introduction, 39
 - LazyClose [LCLOSE], 182
 - list of, 70, 72
 - LivenessTimeout [LTO], 182
 - location of, 75
 - Log Translation [dbtran] utility, 490
 - Logfile [LOG], 183
 - ODBC, 70
 - overview, 164
 - Password [PWD], 184
 - ping [dbping] utility, 494
 - PrefetchBuffer [PBUF], 185
 - PrefetchRows [PROWS], 186
 - priority, 71
 - ServerName [ENG], 187
 - setting, 40
 - StartLine [START], 187
 - stop [dbstop] utility, 506
 - Unconditional [UNC], 187
 - unload [dbunload] utility, 516
 - upgrade [dbupgrad] utility, 523
 - Userid [UID], 188
 - using default parameters, 51
 - validation [dbvalid] utility, 528
- connection properties
- Allow_nulls_by_default, 619
 - Ansi_blanks, 619
 - Ansi_close_cursors_on_rollback, 619
 - Ansi_integer_overflow, 619
 - Ansi_permissions, 619
 - Ansi_update_constraints, 619
 - AnsiNull, 619
 - AppInfo, 619
 - Automatic_timestamp, 619
 - Background_priority, 619
 - BlockedOn, 619
 - Blocking, 619
 - BytesReceived, 619
 - BytesReceivedUncomp, 619
 - BytesSent, 619
 - BytesSentUncomp, 619
 - CacheHits, 619
 - CacheRead, 619
 - CacheReadIndInt, 619
 - CacheReadIndLeaf, 619
 - CacheReadTable, 619
 - Chained, 619
 - CharSet, 619
 - Checkpoint_time, 619

CIS_option, 619
Cis_rowset_size, 619
ClientLibrary, 619
Close_on_endtrans, 619
Commit, 619
CommLink, 619
CommProtocol, 619
Compression, 619
Continue_after_raisererror, 619
Conversion_error, 619
Cooperative_commit_timeout, 619
Cooperative_commits, 619
CurrTaskSwitch, 619
Cursor, 619
CursorOpen, 619
Database_authentication, 619
Date_format, 619
Date_order, 619
DBNumber, 619
Default_Timestamp_Increment, 619
Delayed_commit_timeout, 619
Delayed_commits, 619
DiskRead, 619
DiskReadIndInt, 619
DiskReadIndLeaf, 619
DiskReadTable, 619
DiskWrite, 619
Divide_by_zero_error, 619
Encryption, 619
Escape_character, 619
EventName, 619
Extended_join_syntax, 619
Fire_triggers, 619
Float_as_double, 619
FullCompare, 619
IdleTimeout, 619
IndAdd, 619
IndLookup, 619
Isolation_level, 619
Java_heap_size, 619
Java_input_output, 619
Java_namespace_size, 619
Java_page_buffer_size, 619
JavaHeapSize, 619
Language, 619
LastIdle, 619
LastReqTime, 619
LastStatement, 619
list of, 619
LivenessTimeout, 619
Lock_rejected_rows, 619
LockName, 619
LogFreeCommit, 619
Login_mode, 619
Login_procedure, 619
LogWrite, 619
Max_cursor_count, 619
Max_plans_cached, 619
Max_statement_count, 619
Min_password_length, 619
Min_table_size_for_histogram, 619
Name, 619
Nearest_century, 619
NodeAddress, 619
Non-keywords, 619
Number, 619
On_TSQL_error, 619
Optimization_goal, 619
Optimization_level, 619
PacketSize, 619
PacketsReceived, 619
PacketsReceivedUncomp, 619
PacketsSent, 619
PacketsSentUncomp, 619
Percent_as_comment, 619
Precision, 619
Prefetch, 619
Prepares, 619
PrepStmt, 619
Query_plan_on_open, 619
QueryCachedPages, 619
QueryLowMemoryStrategy, 619
Quoted_identifier, 619
Recovery_time, 619
Replicate_all, 619
reporting, 495
ReqType, 619
RI_trigger_time, 619
Rlbk, 619
RollbackLogPages, 619
Row_counts, 619
Scale, 619
SQL_Flagger_error_level, 619
SQL_Flagger_warning_level, 619
String_rtruncation, 619
Suppress_TDS_debugging, 619
TaskSwitch, 619
TDS_empty_string_is_null, 619
TempTablePages, 619
Time_format, 619
Timestamp_format, 619
TimeZoneAdjustment, 619

- Truncate_timestamp_values, 619
- Truncate_with_auto_commit, 619
- TSQL_Hex_constant, 619
- TSQL_variables, 619
- UncommitOp, 619
- User_estimates, 619
- Userid, 619
- UtilCmdsPermitted, 619
- Wait_for_commit, 619
- connection scenarios, 45
- connection strings
 - about, 40
 - character sets, 279
 - introduction, 39
 - representing, 40
- CONNECTION_PROPERTY function
 - obtaining option settings, 537
- ConnectionName [CON] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 172
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- connections
 - about, 38
 - connection property syntax, 618
 - creating, 415, 416
 - default parameters, 51
 - definition, 38
 - details, 73
 - disallowing with RAISERROR, 578
 - dropping, 150, 151
 - embedded database, 47
 - from utilities, 52
 - Interactive SQL, 81
 - limit, 141
 - list of properties, 619
 - liveness, 151
 - local database, 45
 - network, 49
 - overview, 38
 - performance, 80
 - problems, 73
 - programming interfaces, 38
 - setting a maximum number, 578
 - simple, 45
 - syntax of server-level property, 625
 - troubleshooting, 73, 82, 494, 498
- connectivity
 - jConnect, 43
 - JDBC-ODBC bridge, 43
- ConnsDisabled property
 - description, 626
- console utility
 - about, 450
 - options, 450
 - syntax, 450
- Context Switches statistic, 617
- CONTINUE_AFTER_RAISERROR option
 - compatibility option, 544, 545
 - description, 560
 - obtaining value, 619
 - Open Client, 116
- conventions
 - documentation, xiii
- CONVERSION_ERROR option
 - compatibility option, 545
 - description, 560
 - obtaining value, 619
- COOPERATIVE_COMMIT_TIMEOUT option
 - description, 561
 - obtaining value, 619
 - setting, 541
- COOPERATIVE_COMMITS option
 - description, 561
 - obtaining value, 619
 - setting, 541
- CPU
 - number used, 12, 143
- create backup images wizard
 - using, 333, 438
- create connection statement, 416
 - Replication Server, 415
- create custom collation wizard
 - using, 442
- CREATE DATABASE statement
 - permissions, 12, 229
 - utility database, 226

- ul style="list-style-type: none;">
- create database wizard
 - list of collating sequences, 470
 - using, 465
- CREATE DBSPACE statement
 - using, 220
- create replication definition statement
 - qualifying table owner for Replication Server, 416
 - Replication Server, 416
- CREATE SUBSCRIPTION statement
 - Replication Server, 418
- create write file wizard
 - using, 530
- creating
 - databases, 465
 - databases with dbinit, 466
 - dbspaces, 220
 - groups, 369
 - ODBC data sources, 53, 56
 - ODBC data sources with dbdsn, 451
 - replication definitions for Replication Server, 416
 - Replication Server connections, 415, 416
 - subscriptions for Replication Server, 418
 - users, 357
- creating databases
 - options, 467
 - security, 397
 - Sybase Central, 465
 - Windows CE, 273
- cs option
 - database server, 131
- ct option
 - database server, 131
- CT-Library, 106
- CURRENT USER
 - environment settings, 213
- CurrentCacheSize property
 - description, 626
- CurrentRedoPos property
 - description, 630
- CurrIO property
 - description, 630
- CurrRead property
 - description, 630
- CurrTaskSwitch property
 - description, 619
- CurrWrite property
 - description, 630
- Cursor property
 - description, 619
- CursorOpen property
 - description, 619
- cursors
 - and transactions, 558
 - ANSI_CLOSE_CURSORS_ON_ROLLBACK
 - option, 551
 - closing, 558
 - connection limit, 384, 581
 - database options, 537
 - describing behavior, 58
 - ODBC configuration, 58
 - opening, 594
 - redescribing, 58
- Cursors Open statistic, 616
- Cursors statistic, 616
- custom collations
 - about, 282
 - creating, 282
 - creating databases, 293
 - using with collation [dbcollat] utility, 444
- cw option
 - database server, 131
- cxmp directory
 - about, 205

D

- d option
 - backup [dbbackup] utility, 439, 440
 - collation [dbcollat] utility, 443, 444
 - data source [dbdsn] utility, 451, 454
 - database server, 134
 - Interactive SQL [dbisql] utility, 472, 474
 - language [dblang] utility, 477, 478
 - Log Translation [dbtran] utility, 488, 490
 - ping [dbping] utility, 494

- service creation [dbsvc] utility, 499, 501
 - stop [dbstop] utility, 505, 506
 - unload [dbunload] utility, 514, 518
 - write file [dbwrite] utility, 531, 532
- d1 option
 - Interactive SQL [dbisql] utility, 474
- daemon
 - database server as, 21
 - Log Transfer Manger [dbltm] utility, 483
 - LTM, 483
 - Replication Agent, 483
 - running the server as, 153
- data source description
 - ODBC, 56
- data source name
 - ODBC, 56
- data source utility
 - about, 451
 - options, 453, 454
 - syntax, 451
- data sources
 - about, 41, 53
 - configuring, 56
 - connection parameters, 39
 - creating, 53, 56
 - creating ODBC with dbdsn, 451
 - data source [dbdsn] utility, 454
 - embedded SQL, 53
 - example, 48
 - file, 62
 - Interactive SQL [dbisql] utility, 474
 - ODBC, 53
 - UNIX, 63
 - using with jConnect, 43
 - Windows CE, 63
- data type conversion
 - errors, 560
- data types
 - supported, 423
- database access
 - controlling, 390
- database administration utilities
 - about, 435
- database administrator
 - defined, 352
- database connections
 - ping [dbping] utility, 494
- database file encryption
 - initialization [dbinit] utility, 468
- database files
 - backups, 305
 - compressing with dbshrink, 448
 - encrypting, 399
 - erasing, 458
 - limit, 220
 - location, 8
 - maximum size, 636
 - media failure, 340
 - NDS, 124
 - ODBC configuration, 59
 - paths, 124
 - security, 399, 403
 - UNC filenames, 124
 - uncompressing, 511
- database information
 - getting with dbinfo, 463
- database name
 - case sensitivity, 11
 - ODBC configuration, 59
 - option, 10
 - setting, 160
- database object debugger wizard
 - using, 456
- database options, 535
 - about, 536
 - AUDITING, 553
 - BACKGROUND_PRIORITY, 555
 - BLOCKING, 556
 - BLOCKING_TIMEOUT, 556
 - CHECKPOINT_TIME, 557
 - CIS_ROWSET_SIZE, 557
 - classification, 539
 - COOPERATIVE_COMMIT_TIMEOUT, 561
 - COOPERATIVE_COMMITS, 561
 - DEFAULT_TIMESTAMP_INCREMENT, 564
 - DELAYED_COMMIT_TIMEOUT, 564
 - DELAYED_COMMITS, 565
 - deleting settings, 539
 - ESCAPE_CHARACTER, 567
 - finding values, 537
 - FIRST_DAY_OF_WEEK, 568
 - GLOBAL_DATABASE_ID, 569
 - initial settings, 539

- JAVA_HEAP_SIZE, 576
- JAVA_INPUT_OUTPUT, 577
- JAVA_NAMESPACE_SIZE, 577
- list of compatibility options, 544, 545
- list of general database options, 541
- list of Interactive SQL options, 548
- list of replication options, 547
- LOGIN_MODE, 577
- LOGIN_PROCEDURE, 578
- MAX_CURSOR_COUNT, 581
- MAX_HASH_SIZE, 581
- MAX_PLANS_CACHED, 581
- MAX_STATEMENT_COUNT, 583
- MAX_WORK_TABLE_HASH_SIZE, 582
- MIN_PASSWORD_LENGTH, 583
- MIN_TABLE_SIZE_FOR_HISTOGRAM, 583
- NEAREST_CENTURY option, 584
- ON_CHARSET_CONVERSION_FAILURE, 585
- Open Client, 116
- OPTIMIZATION_GOAL, 587
- options, 10
- PINNED_CURSOR_PERCENT_OF_CACHE, 591
- PRECISION, 591
- PREFETCH, 592
- PRESERVE_SOURCE_FORMAT, 593
- PREVENT_ARTICLE_PKEY_UPDATE, 593
- RECOVERY_TIME, 595
- RETURN_DATE_TIME_AS_STRING, 596
- RETURN_JAVA_AS_STRING, 597
- ROW_COUNTS, 597
- SCALE, 598
- scope and duration, 537
- setting, 536
- setting in Interactive SQL, 548
- SORT_COLLATION, 598
- startup settings, 116
- SUPPRESS_TDS_DEBUGGING, 601
- TDS_EMPTY_STRING_IS_NULL, 601
- TIME_ZONE_ADJUSTMENT, 602
- TRUNCATE_DATE_VALUES, 604
- TRUNCATE_TIMESTAMP_VALUES, 604
- TRUNCATE_WITH_AUTO_COMMIT, 605
- USER_ESTIMATES, 606
- WAIT_FOR_COMMIT, 608
- database pages
 - displaying size of, 463
- database performance and connection properties, 609
- database properties
 - Alias, 630
 - BlankPadding, 630
 - BlobArenas, 630
 - CacheHits, 630
 - CacheRead, 630
 - CacheReadIndInt, 630
 - CacheReadIndLeaf, 630
 - CacheReadTable, 630
 - Capabilities, 630
 - CaseSensitive, 630
 - CharSet, 630
 - CheckpointUrgency, 630
 - Chkpt, 630
 - ChkptFlush, 630
 - ChkptPage, 630
 - CommitFile, 630
 - CompressedBTrees, 630
 - Compression, 630
 - ConnCount, 630
 - CurrentRedoPos, 630
 - CurrIO, 630
 - CurrRead, 630
 - CurrWrite, 630
 - DBFileFragments, 630
 - DiskRead, 630
 - DiskReadIndInt, 630
 - DiskReadIndLeaf, 630
 - DiskReadTable, 630
 - DiskWrite, 630
 - Encryption, 630
 - ExtendDB, 630
 - ExtendTempWrite, 630
 - File, 630
 - FileVersion, 630
 - FreePageBitMaps, 630
 - FullCompare, 630
 - GlobalDBId, 630
 - Histograms, 630
 - IdleCheck, 630
 - IdleChkpt, 630
 - IdleChkTime, 630
 - IdleWrite, 630
 - IndAdd, 630
 - IndLookup, 630
 - IOToRecover, 630
 - IQStore, 630
 - JavaHeapSize, 630
 - JavaNSSize, 630
 - JDKVersion, 630
 - Language, 630

- LargeProcedureIDs, 630
- list of, 630
- LockTablePages, 630
- LogFileFragments, 630
- LogFreeCommit, 630
- LogName, 630
- LogWrite, 630
- LTMGeneration, 630
- LTMTrunc, 630
- MapPages, 630
- MaxIO, 630
- MaxRead, 630
- MaxWrite, 630
- MultiByteCharSet, 630
- Name, 630
- PageRelocations, 630
- PageSize, 630
- PreserveSource, 630
- ProcedurePages, 630
- QueryCachedPages, 630
- QueryLowMemoryStrategy, 630
- ReadOnly, 630
- RecoveryUrgency, 630
- RelocatableHeapPages, 630
- RemoteTrunc, 630
- reporting, 495
- RollbackLogPages, 630
- SeparateCheckpointLog, 630
- SeparateForeignKeys, 630
- SyncTrunc, 630
- syntax, 618, 630
- TableBitMaps, 630
- TempFileName, 630
- TempTablePages, 630
- TransactionsSpanLogs, 630
- TriggerPages, 630
- VariableHashSize, 630
- ViewPages, 630
- database servers
 - automatic starting, 6
 - command line, 120
 - locating, 76, 82
 - name, 145
 - name caching, 80
 - name option, 10
 - NetWare, 6, 124
 - options, 10
 - preventing from starting, 166
 - running as a daemon, 21
 - running in the background, 21
 - security, 397
 - shutting down, 152
 - starting, 5, 8
 - starting on Windows CE, 9
 - stopping, 18, 187
 - UNIX, 6
 - Windows operating systems, 5
 - Windows services, 21
- database size
 - limit, 636
- database utilities
 - and database connections, 52
 - backup [dbbackup], 438
 - collation [dbcollat], 442
 - compression [dbshrink], 448
 - console [dbconsole], 450
 - data source [dbdsn], 451
 - debugger [dbprdbg], 456
 - erase [dberase], 458
 - file hiding [dbfhide], 446
 - histogram [dbhist], 461
 - information [dbinfo], 463
 - initialization [dbinit], 465
 - Interactive SQL [dbisql], 472
 - language [dblang], 477
 - license [dblic], 479
 - Log Transfer Manager [dbltm], 481
 - log translation [dbtran], 488
 - ping [dbping], 494
 - rebuild [rebuild], 497
 - server location [dblocate], 498
 - service creation [dbsvc], 499
 - spawn [dbspawn], 503
 - stop [dbstop], 505
 - transaction log [dblog], 507
 - uncompression [dbexpand], 511
 - unload [dbunload], 513
 - upgrade [dbupgrad], 521
 - validation [dbvalid], 526
 - write file [dbwrite], 530
- DATABASE_AUTHENTICATION property
 - description, 619
- DatabaseFile [DBF] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 173
 - embedded databases, 47
 - introduction to connection parameters, 39

- tips, 72
 - using default connection parameters, 51
- DatabaseName [DBN] connection parameter
- about, 70
 - about connection parameters, 164
 - description, 174
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- databases, 451
- accessing the Connect dialog, 42
 - allocating space, 222
 - altering dbspaces, 222
 - automatic stopping, 138
 - backing up from Sybase Central, 438
 - backup, 300
 - character set, 278
 - compressed, 512
 - compressing from Sybase Central, 448, 456
 - connecting from Interactive SQL, 42
 - connecting from Sybase Central, 42
 - connecting to, 38, 73
 - connecting to a local database, 46
 - connection scenarios, 45
 - creating, 465
 - creating dbspaces, 220
 - creating with dbinit, 466
 - custom collations, 293
 - deleting dbspaces, 222
 - erasing, 458
 - information, 463
 - initializing, 465
 - Java-enabling, 524
 - large databases, 220
 - maximum size, 636
 - multiple-file, 220
 - new, 29
 - page usage, 463
 - permissions, 351
 - permissions to load, 141
 - permissions to start, 139
 - permissions to stop, 140
 - permissions to unload, 141
 - properties of, 630
 - read-only, 13, 149, 161, 224
 - rebuilding, 497
 - recovery, 300
 - replicating entire, 430
 - starting, 20
 - stopping, 20, 505
 - unloading, 513
 - upgrading, 521
 - utilities, 438, 442, 446, 448, 450, 456, 458, 461, 463, 465, 472, 477, 479, 481, 488, 494, 497, 498, 499, 503, 505, 507, 511, 513, 521, 526, 530
 - utility, 226
 - validating, 329, 526
 - validating from Sybase Central, 526
 - write files, 530
- DatabaseSwitches [DBS] connection parameter
- about, 70
 - about connection parameters, 164
 - description, 174
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- datasource option
- Interactive SQL [dbisql] utility, 472, 474
- DataSourceName [DSN] connection parameter
- about, 70
 - about connection parameters, 164
 - description, 175
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
 - Windows CE, 63
- DATE data types
- time values in, 604
- DATE_FORMAT option
- compatibility option, 544
 - description, 562
 - obtaining value, 619
 - setting, 541
- DATE_ORDER option
- compatibility option, 544
 - description, 564
 - obtaining value, 619
 - setting, 541
- dates
- NEAREST_CENTURY option, 584
 - time values in, 604
- DBA authority
- about, 352
 - defined, 352
 - granting, 358

- not inheritable, 369
- security tips, 403
- dBASE file format, 570
- dBASE II file format, 570, 588
- dBASE III file format, 570, 588
- dbbackup utility
 - about, 438
 - full backup, 329
 - live backup, 340
 - options, 440
 - syntax, 439
- dbcc function, 509
- dbcollat utility
 - about, 442
 - custom collations, 292, 293
 - options, 443
 - syntax, 443
- dbconsole utility
 - about, 450
 - options, 450
 - syntax, 450
- dbdsn utility
 - about, 451
 - options, 453
 - syntax, 451
 - using, 56
- dbeng8
 - command line, 120
 - personal database server, 4
 - syntax, 120
- dberase utility
 - about, 458
 - options, 459
 - syntax, 459
- dbexpand utility
 - about, 511
 - options, 512
 - syntax, 511
- dbfhide utility
 - about, 446
 - syntax, 446
- DBFileFragments property
 - description, 630
- dbhist utility
 - about, 461
 - options, 462
 - syntax, 461
- dbinfo utility
 - about, 463
 - options, 463
 - syntax, 463
- dbinit utility
 - about, 465
 - options, 467
 - syntax, 466
- dbisql utility
 - about, 472
 - options, 474
 - syntax, 472
- dblang utility
 - about, 477
 - options, 478
 - syntax, 477
- dblgen8.res
 - locating, 206
- DB-Library, 106
- dblic utility
 - about, 479
 - options, 480
 - syntax, 479
- dblocate utility
 - about, 498
 - options, 498
 - syntax, 498
- dblog utility
 - about, 507
 - auditing, 396
 - command line, 508
 - options, 509
 - syntax, 507
 - transaction log mirrors, 347
- dbltm utility
 - options, 482
 - syntax, 481
- dbmlsrv
 - licensing, 479

- DBNumber property
 - description, 619
- dbo user
 - system objects, 513
- dbping utility
 - about, 494
 - options, 494
 - syntax, 494
 - using, 82
- dbshrink utility
 - about, 448
 - options, 449
 - syntax, 448
- dbspace creation wizard
 - using, 221
- dbspaces
 - altering, 222
 - creating, 220
 - deleting, 222
 - limit, 636
- dbspawn utility
 - about, 503
 - options, 504
 - syntax, 503
- dbsrv
 - licensing, 479
- dbsrv8
 - command line, 120
 - network database server, 4
 - syntax, 120
- dbstop utility
 - about, 505
 - options, 505
 - permissions, 12, 140
 - syntax, 505
 - using, 18
- dbsvc utility
 - about, 499
 - options, 500
 - syntax, 499
- dbtran command-line utility
 - command line, 490
- dbtran utility
 - about, 488
 - auditing, 394, 396
 - options, 490
 - syntax, 488
 - transaction logs, 344
 - uncommitted changes, 344
- dbunload utility
 - about, 513
 - options, 516
 - syntax, 514
- dbupgrad utility
 - about, 521
 - options, 523
 - syntax, 522
- dbvalid utility
 - about, 526
 - options, 528
 - syntax, 527
 - using, 329
- dbwrite utility
 - about, 530
 - auditing, 396
 - options, 532
 - syntax, 531
- debugger [dbprdbg] utility
 - about, 456
 - options, 457
 - syntax, 456
- debugging
 - debugger [dbprdbg] utility, 456
 - event handlers, 240, 246
 - SQL scripts, 472
- decimal precision
 - database option, 591
- default character set
 - about, 264
 - UNIX, 264
 - Windows, 264
- DEFAULT_TIMESTAMP_INCREMENT option
 - description, 564
 - obtaining value, 619
 - setting, 541

- DefaultCollation property
 - about, 226
 - description, 626
- defaults
 - connection parameters, 51
- DELAYED_COMMIT_TIMEOUT option
 - description, 564
 - obtaining value, 619
 - setting, 541
- DELAYED_COMMITS option
 - description, 565
 - obtaining value, 619
 - setting, 541
- DELETE permissions, 359
- DELETE statement, 422
- DELETE_OLD_LOGS option
 - description, 566
 - replication option, 547
 - resetting truncation offset, 509
 - transaction log options, 509
 - using, 430
- deletes
 - ANSI behavior, 551
 - Transact-SQL permissions, 551
- deleting
 - dbspaces, 222
 - erase utility, 458
 - groups, 375
 - integrated logins, 86
 - services, 501
 - users, 366
- deleting databases
 - security, 397
- Delphi
 - ODBC configuration for, 57
- dependencies
 - of services, 32, 33
 - setting, 501
- DESCRIBE_JAVA_FORMAT option
 - description, 566
 - Interactive SQL settings, 548
- describing
 - cursor behavior, 58
- dial-up networking
 - connections, 94
- digit characters
 - defined, 286
- digits
 - maximum number, 591
- directory structure, 204
- DisableMultiRowFetch [DMRF] connection
 - parameter
 - about, 70
 - about connection parameters, 164
 - description, 175
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- disconnecting users
 - using the Console utility, 450
- disk access
 - options, 134
- Disk Active I/O statistics, 613
- disk cache
 - operating system, 152
- disk controllers
 - transaction log management, 318
- disk crashes
 - about, 304
- disk full
 - error writing to transaction log, 318
- disk I/O statistics, 613
- Disk Maximum I/O statistics, 613
- disk mirroring
 - transaction logs, 318
- disk read statistics, 614
- Disk Reads Active statistics, 614
- Disk Reads Index interior/sec statistics, 614
- Disk Reads Index leaf/sec statistics, 614
- Disk Reads Maximum Active statistics, 614
- Disk Reads Table/sec statistics, 614
- Disk Reads Total Pages/sec statistics, 614

- disk space
 - event example, 237
- disk write statistics, 614
- Disk Writes Active statistics, 614
- Disk Writes Commit Files/sec statistics, 614
- Disk Writes Database Extends/sec statistics, 614
- Disk Writes Maximum Active statistics, 614
- Disk Writes Pages/sec statistics, 614
- Disk Writes Temp Extends/sec statistics, 614
- Disk Writes Transaction Log Group Commits statistics, 614
- Disk Writes Transaction Log/sec statistics, 614
- DiskRead property
 - description, 619, 630
- DiskReadIndInt property
 - description, 619, 630
- DiskReadIndLeaf property
 - description, 619, 630
- DiskReadTable property
 - description, 619, 630
- disks
 - fragmentation and performance, 222
 - recovery from failure of, 340
- DiskWrite property
 - description, 619, 630
- distributed transactions
 - enlistment, 151
- DIVIDE_BY_ZERO option
 - compatibility option, 545
- DIVIDE_BY_ZERO_ERROR option
 - description, 566
 - obtaining value, 619
- dl option
 - Interactive SQL [dbisql] utility, 472
 - Log Transfer Manger [dbltm] utility, 481, 483
- DLL
 - location, 206
- DLL communication parameter
 - description, 192

- DoBroadcast [DOBROAD] communication
 - parameter description, 192
- documentation
 - conventions, xiii
 - SQL Anywhere Studio, x
- Driver Not Capable error
 - ODBC configuration, 57
- drivers
 - choosing a driver for a connection, 43
- dropping
 - groups, 375
 - users, 366
- DSEdit utility
 - about, 108
 - entries, 112
 - not included with SQL Anywhere Studio, 108
 - starting, 111
 - using, 110, 410
- DSN connection parameter
 - about, 53
 - Windows CE, 63

E

- e option
 - collation [dbcollat] utility, 443, 444
 - initialization [dbinit] utility, 466, 468
 - unload [dbunload] utility, 514, 518
- ea option
 - initialization [dbinit] utility, 466, 468
 - unload [dbunload] utility, 514, 518
- EBCDIC
 - collating sequence, 471
- ec option
 - data source [dbdsn] utility, 454
 - database server, 135
- ec server option
 - data source [dbdsn] utility, 451
- ECHO option
 - description, 567
 - Interactive SQL settings, 548

- ek option
 - compression [dbshrink] utility, 448, 449
 - database server, 159
 - erase [dberase] utility, 459
 - initialization [dbinit] utility, 466, 468
 - Log Transfer Manger [dbltm] utility, 481, 482
 - Log Translation [dbtran] utility, 488, 491
 - transaction log [dblog] utility, 507, 509
 - uncompression [dbexpand] utility, 511, 512
 - unload [dbunload] utility, 514, 518
 - write file [dbwrite] utility, 531, 532
- embedded databases
 - connecting, 47
 - connection parameters, 72
 - Java, 48
 - starting, 47
- embedded SQL
 - connections, 38
 - interface library, 74
- encoding
 - character sets, 253
 - definition, 253
 - multibyte character sets, 285
- encrypted passwords
 - ODBC configuration, 59
- EncryptedPassword [ENP] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 177
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- encryption
 - AES algorithm, 399
 - Certicom, 135
 - client/server communications, 135, 137, 159, 177
 - communications, 198, 398
 - creating databases with dbinit, 466
 - database files, 399
 - file hiding [dbfhide] utility, 446
 - MDSR algorithm, 399
 - network packets, 62
 - passwords, 177, 391
 - performance of encrypted databases, 402
 - Rijndael algorithm, 399
 - simple, 399
 - strong, 135, 137, 159, 177, 179, 399
- Encryption [ENC] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 177
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- encryption algorithms
 - AES, 399
 - initialization [dbinit] utility, 468
 - MDSR, 399
 - Rijndael, 399
 - unload [dbunload] utility, 518
- Encryption Key [DBKEY] connection parameter
 - description, 179
- encryption keys
 - compression [dbshrink] utility, 449
 - erase [dberase] utility, 460
 - initialization [dbinit] utility, 468
 - Log Transfer Manger [dbltm] utility, 482
 - Log Translation [dbtran] utility, 491
 - ODBC configuration, 59
 - transaction log [dblog] utility, 509
 - uncompression [dbexpand] utility, 512
 - unload [dbunload] utility, 518
 - write file [dbwrite] utility, 532
- Encryption property
 - connection property description, 619
 - database property description, 630
- EngineName [ENG] connection parameter
 - about, 70
 - about connection parameters, 164
 - character sets, 71
 - description, 176
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- enlistment
 - distributed transactions, 151
- environment variables, 52
 - about, 208
 - ASANY8, 209
 - ASANYSH8, 209
 - ASTMP, 209
 - ERRORLEVEL, 475
 - LD_LIBRARY_PATH, 210
 - PATH, 210

- setting, 208
 - setting TEMP directory on CE, 214
 - SQLCONNECT, 52, 210
 - SQLLOCALE, 211
 - SQLPATH, 211
 - SQLREMOTE, 211
 - SYBASE, 211
 - TEMP, 212
 - TEMPDIR, 212
 - TMP, 212
- environment-variable server option
- database server, 125
- ep option
- compression [dbshrink] utility, 448, 449
 - database server, 137
 - erase [dberase] utility, 459, 460
 - initialization [dbinit] utility, 466, 468
 - Log Transfer Manger [dbltm] utility, 481, 482
 - Log Translation [dbtran] utility, 488, 491
 - transaction log [dblog] utility, 507, 509
 - uncompression [dbexpand] utility, 511, 512
 - unload [dbunload] utility, 514, 518
 - write file [dbwrite] utility, 531, 532
- erase database wizard
- using, 458
- erase utility
- about, 458
 - options, 459
 - syntax, 459, 461
- erasing databases, 458
- Sybase Central, 458
- error handling
- Interactive SQL, 586
 - Transact-SQL procedures, 587
- error messages
- character set translation, 278
- ERRORLEVEL environment variable
- Interactive SQL return code, 475
- errors
- divide by zero, 566
 - in Interactive SQL, 586
 - Transact-SQL procedures, 587
- escape characters
- unload [dbunload] utility, 519
- ESCAPE_CHARACTER option
- compatibility option, 544, 545
 - description, 567
 - obtaining value, 619
- estimates
- recovery time, 595
 - row count, 597
 - USER_ESTIMATES option, 606
- ethernet, 102
- euro symbol
- 1252LATIN1 collation, 274, 275
- event creation wizard
- using, 244
- event handlers
- debugging, 240, 246
 - defined, 232
 - internals, 243
- event handling
- about, 231
- event log
- suppressing messages, 125
- event types
- about, 242
- EventName property
- description, 619
- events
- about, 236
 - defined, 232
 - handling, 231, 232
 - internals, 242
 - limitations, 636
- examples
- directories, 205
- Excel file format, 570, 588
- executable name
- license [dblic] utility, 480
- execution threads
- number, 141
- expanding
- database files, 511
- explicit selectivity estimates
- USER_ESTIMATES option, 606

exporting data
output format, 588

ExtendDB property
description, 630

extended characters
about, 255
collation [dbcollat] utility, 445

EXTENDED_JOIN_SYNTAX option
description, 567
obtaining value, 619
setting, 541

ExtendedName [ENAME] communication
parameter
description, 193

ExtendTempWrite property
description, 630

external functions
stack size, 140

F

–f option
database server, 158
Log Translation [dbtran] utility, 488, 491
spawn [dbspawn] utility, 503, 504
validation [dbvalid] utility, 527, 528
write file [dbwrite] utility, 531, 532

failure
recovery from, 300

–fd option
validation [dbvalid] utility, 527, 528

feedback
documentation, xvii
providing, xvii

fetch operation
suppressing warnings, 57

–fi option
validation [dbvalid] utility, 527, 528

file data sources
creating, 62

file hiding utility
about, 446
syntax, 446

file locations
NetWare, 205
Windows CE, 205

File property
description, 630

FileDataSourceName [FILEDSN] connection
parameter
about, 70
about connection parameters, 164
description, 180
introduction to connection parameters, 39
referencing file data sources, 53
tips, 72
using default connection parameters, 51
Windows CE, 63

filename server option
database server, 126

files
location, 206

FileVersion property
description, 630

FIPS
SQL_FLAGGER_ERROR option, 599

FIRE_TRIGGERS option
compatibility option, 545
description, 568
obtaining value, 619

firewalls
connecting across, 93, 190, 191, 194, 198

first row optimization option
OPTIMIZATION_GOAL, 587

FIRST_DAY_OF_WEEK option
description, 568
setting, 541

FIXED file format, 570, 588

fixed width character sets
about, 259

FLOAT_AS_DOUBLE option
compatibility option, 544, 545
description, 569

- obtaining value, 619
 - Open Client, 116
 - follow bytes
 - about, 260, 279
 - ForceStart [FORCESTART] connection parameter
 - description, 180
 - format
 - input file, 570
 - forward log, 305
 - FoxPro file format, 570, 588
 - fragmentation and performance, 222
 - frame type, 102
 - FreeBuffers property
 - description, 626
 - FreePageBitMaps property
 - description, 630
 - full backups
 - performing, 329
 - FullCompare property
 - description, 619, 630
 - functions
 - replicating, 424
 - support for APC format, 424
 - fx option
 - validation [dbvalid] utility, 527, 528
- ## G
- g option
 - data source [dbdsn] utility, 451, 454
 - Log Translation [dbtran] utility, 488, 491
 - service creation [dbsvc] utility, 499, 501
 - transaction log [dblog] utility, 507, 509
 - ga option
 - database server, 138
 - gb option
 - database server, 138
 - gc option
 - database server, 138
 - gd option
 - database server, 139
 - ge option
 - database server, 140
 - gf option
 - database server, 140
 - gk option
 - database server, 140
 - gl option
 - database server, 141
 - GLOBAL_DATABASE_ID option
 - description, 569
 - setting, 541
 - GlobalDBId property
 - description, 630
 - gm option
 - database server, 141
 - gn option
 - database server, 141
 - gp option
 - database server, 142
 - gr option
 - database server, 143
 - GRANT statement
 - creating groups, 369
 - DBA authority, 358
 - group membership, 370
 - new users, 357
 - passwords, 358
 - permissions, 359
 - procedures, 363
 - RESOURCE authority, 358
 - table permissions, 359
 - WITH GRANT OPTION, 363
 - without password, 373
 - granting
 - remote permissions, 365
 - group creation wizard
 - using, 370
 - group dependencies
 - setting, 501

GROUP permissions
not inheritable, 369

groups

- creating, 369
- deleting, 375
- deleting integrated logins, 86
- granting integrated logins, 84
- leaving, 371
- managing, 369
- membership, 370
- of services, 32
- permissions, 354, 372
- PUBLIC, 374
- remote permissions, 365
- revoking membership, 371
- setting dependencies, 501
- setting options, 356
- SYS, 374
- without passwords, 373

–gss option
database server, 143

–gt option
database server, 143

–gu option
database server, 143

–gx option
database server, 144

H

handling events, 231, 232

hardware mirroring
transaction logs, 318

hash size
MAX_HASH_SIZE option, 581
MAX_WORK_TABLE_HASH_SIZE option,
582

hexadecimal constants
data type, 606

high availability
database servers, 21
live backups, 340

histogram utility
about, 461
options, 462

histograms
minimum table size, 583
viewing with dbhist, 461

Histograms property
description, 630

Host [IP] communication parameter
description, 194

–host option
Interactive SQL [dbisql] utility, 472, 474

HTML file format, 588

I

–i option
initialization [dbinit] utility, 468
service creation [dbsvc] utility, 501
upgrade [dbupgrad] utility, 524
validation [dbvalid] utility, 529

–I option
initialization [dbinit] utility, 466
Log Transfer Manger [dbltn] utility, 481, 482
service creation [dbsvc] utility, 499
upgrade [dbupgrad] utility, 522
validation [dbvalid] utility, 527

I/O
idle, 326
operations and the database server, 134

IANA
port number, 198

icons
for running services, 29
used in manuals, xiv

identification
client applications, 165

identifiers
case insensitivity, 261
international aspects, 261

Idle [IDLE] connection parameter
about, 70
about connection parameters, 164

- description, 181
- introduction to connection parameters, 39
- tips, 72
- using default connection parameters, 51
- Idle Actives/sec statistic, 611
- Idle Checkpoint Time statistic, 611
- Idle Checkpoints/sec statistic, 611
- idle I/O
 - task, 326
- idle server
 - event example, 238
- idle timeout
 - ODBC configuration, 61
- Idle Writes/sec statistic, 611
- IdleCheck property
 - description, 630
- IdleChkpt property
 - description, 630
- IdleChkTime property
 - description, 630
- IdleTime event
 - polling, 242
- IdleTimeout property
 - description, 619, 626
- IdleWrite property
 - description, 630
- ii option
 - unload [dbunload] utility, 514, 519
- il option
 - transaction log [dblog] utility, 507, 509
- IN keyword
 - CREATE TABLE statement, 220
- incremental backups
 - dbbackup utility, 438
- IndAdd property
 - description, 619, 630
- Index Adds/sec statistics, 615
- Index Full Compares/sec statistics, 615
- Index Lookups/sec statistics, 615
- indexes
 - hash size, 581, 582
 - limitations, 636
 - statistics, 615
- IndLookup property
 - description, 619, 630
- information utility
 - about, 463
 - options, 463
 - syntax, 463
- INI files
 - about, 213
- initialization utility
 - about, 465
 - options, 467
 - syntax, 466
- initializing databases with dbinit, 466
- INPUT_FORMAT option
 - description, 570
 - Interactive SQL settings, 548
- INSERT permissions, 359
- INSERT statement, 422
 - truncation of strings, 600
- installation
 - directory, 204
 - NetWare, 205
 - registry settings, 214
 - Windows CE, 205
- integer overflow
 - ANSI behavior, 551
- Integrated [INT] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 181
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- integrated login creation wizard
 - using, 85
- integrated logins
 - about, 83
 - creating, 84
 - default user, 90

- deleting, 86
- network aspects, 89
- ODBC configuration, 58
- operating systems, 83
- security features, 87, 89, 390
- using, 84, 86

Interactive SQL

- AUTO_COMMIT option, 554
- command line, 473
- Java output, 566
- list of options, 548, 554, 555, 558, 559, 566, 567, 570, 572, 573, 574, 575, 576, 585, 586, 588, 590, 600, 605
- return codes, 475
- setting options, 548
- starting from Sybase Central, 472

Interactive SQL options

- AUTO_COMMIT, 554
- AUTO_REFETCH, 554
- BELL, 555
- classification, 539
- COMMAND_DELIMITER, 558
- COMMIT_ON_EXIT, 559
- DESCRIBE_JAVA_FORMAT, 566
- ECHO, 567
- initial settings, 539
- INPUT_FORMAT, 570
- ISQL_COMMAND_TIMING, 572
- ISQL_ESCAPE_CHARACTER, 572
- ISQL_FIELD_SEPARATOR, 573
- ISQL_LOG, 574
- ISQL_PLAN, 574
- ISQL_PRINT_RESULT_SET, 575
- ISQL_QUOTE, 576
- list of, 548
- NULLS, 585
- ON_ERROR, 586
- OUTPUT_FORMAT, 588
- OUTPUT_LENGTH, 590
- OUTPUT_NULLS, 590
- setting, 548
- STATISTICS, 600
- TRUNCATION_LENGTH, 605

Interactive SQL utility

- about, 472
- options, 474
- syntax, 472

- interface libraries
 - connections, 38
 - locating, 74

interfaces file

- configuring, 110
- Log Transfer Manger [dbltm] utility, 482
- Open Servers, 410

internals

- backups, 321
- event handlers, 243
- event handling, 242
- events, 242
- schedules, 242

- international languages and character sets, 249

IOTorecover property

- description, 630

IP address

- about, 112
- ping, 102

IQStore property

- description, 630

-ir option

- Log Translation [dbtran] utility, 488, 492
- transaction log [dblog] utility, 507, 509

-is option

- Log Translation [dbtran] utility, 488, 491

IsIQ property

- description, 626

IsJavaAvailable property

- description, 626

IsNetworkServer property

- description, 626

ISO_1 collation

- about, 274

isolation levels

- ODBC configuration, 56

ISOLATION_LEVEL option

- compatibility option, 544
- description, 571
- obtaining value, 619
- Open Client, 116
- setting, 541

ISQL_COMMAND_TIMING option
description, 572
Interactive SQL settings, 548

ISQL_ESCAPE_CHARACTER option
description, 572
Interactive SQL settings, 548

ISQL_FIELD_SEPARATOR option
description, 573
Interactive SQL settings, 548

ISQL_LOG option
description, 574
Interactive SQL settings, 548

ISQL_PLAN option
description, 574
Interactive SQL settings, 548

ISQL_PRINT_RESULT_SET option
description, 575

ISQL_QUOTE option
description, 576
Interactive SQL settings, 548

IsRuntimeServer property
description, 626

-it option
Log Translation [dbtran] utility, 488, 492

-ix option
unload [dbunload] utility, 514, 519

J

-j option
Log Translation [dbtran] utility, 488, 492
unload [dbunload] utility, 514, 519
upgrade [dbupgrad] utility, 522, 524

-ja option
initialization [dbinit] utility, 466, 469
upgrade [dbupgrad] utility, 522, 524

Java
and memory use, 576, 577
connection parameters, 72
enabling a database, 524
file access, 577

Java classes
upgrade [dbupgrad] utility, 524

Java debugger
about, 456
syntax, 456

Java in the database
upgrading, 523

Java VM statistics, 615

java.io package
support, 577

JAVA_HEAP_SIZE option
description, 576
obtaining value, 619
setting, 541

JAVA_INPUT_OUTPUT option
description, 577
obtaining value, 619

JAVA_NAMESPACE_SIZE option
description, 577
obtaining value, 619
setting, 541

Java_page_buffer_size property
description, 619

JavaGlobFix property
description, 626

JavaHeapSize property
description, 619, 630

JavaNssize property
description, 630

jConnect
initialization [dbinit] utility, 468
Interactive SQL [dbisql] utility, 475
TDS, 106
upgrade [dbupgrad] utility, 524

jConnect driver, 43

-jConnect option
Interactive SQL [dbisql] utility, 472, 475

JDBC
compatibility options, 544

JDBC connectivity, 43

JDBC/ODBC bridge driver
Interactive SQL [dbisql] utility, 475

JDBC-ODBC bridge driver, 43

- jdk option
 - initialization [dbinit] utility, 466, 469
 - upgrade [dbupgrad] utility, 522, 524

- JDKVersion property
 - description, 630

- jr option
 - upgrade [dbupgrad] utility, 522, 524

- JVM Global Fixed Size statistics, 615

- JVM Heap Size statistics, 615

- JVM Namespace Size statistics, 615

- jxmp directory
 - about, 205

K

- k option
 - initialization [dbinit] utility, 466, 470

- keeping your data secure, 387

- keyboard mapping
 - about, 253

- keywords
 - NON_KEYWORDS option, 584
 - turning off, 584

L

- l option
 - backup [dbbackup] utility, 439, 440
 - data source [dbdsn] utility, 451, 454
 - initialization [dbinit] utility, 466, 470
 - license [dblic] utility, 479, 480
 - Log Translation [dbtran] utility, 490
 - ping [dbping] utility, 494, 495
 - service creation [dbsvc] utility, 499, 501

- LANalyzer, 102

- language
 - locale, 263
 - registry settings, 214, 477

- language DLL
 - locating, 206
 - registry settings, 477

- Language property
 - description, 619, 626, 630

- language resource library
 - messages file, 253
 - registry settings, 477

- language support
 - about, 250
 - collations, 287
 - multibyte character sets, 277
 - overview, 250

- language utility
 - about, 477
 - options, 478
 - syntax, 477

- LargeProcedureIDs property
 - description, 630

- LastIdle property
 - description, 619

- LastReqTime property
 - description, 619

- LastStatement property
 - description, 619

- LazyClose [LCLOSE] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 182
 - introduction to connection parameters, 39
 - using default connection parameters, 51

- LD_LIBRARY_PATH environment variable
 - description, 210

- leaving
 - groups, 371

- LegalCopyright property
 - description, 626

- LegalTrademarks property
 - description, 626

- libctl.cfg file
 - DSEdit, 111

- libraries
 - ping [dbping] utility, 495

- ul style="list-style-type: none;">
- license type
 - license [dblic] utility, 480
 - specifying with Log Transfer Manger [dbltm] utility, 482
 - specifying with Log Translation [dbtran] utility, 490
- license utility
 - about, 479
 - options, 480
 - syntax, 479
- LicenseCount property
 - description, 626
- LicensedCompany property
 - description, 626
- LicensedUser property
 - description, 626
- licenses
 - adding with dblic command-line utility, 479
- LicenseType property
 - description, 626
- licensing
 - database servers, 479
 - synchronization servers, 479
- limit
 - connections, 141
- limitations
 - Adaptive Server Anywhere, 636
 - cache size, 636
 - columns, 636
 - databases, 636
 - events, 636
 - indexes, 636
 - tables, 636
 - temporary tables, 636
- live backups, 340
 - backup [dbbackup] utility, 440
 - regular backups and, 319
- liveness
 - connections, 151
 - ODBC configuration, 60
- LivenessTimeout [LTO] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 182
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- LivenessTimeout property
 - description, 619, 626
- LOAD TABLE statement
 - security, 397
- loading data
 - security, 397
- local machine
 - environment settings, 213
- locales
 - about, 262
 - character sets, 278
 - language, 263
 - setting, 289
- localhost machine name, 112
- LocalOnly [LOCAL] communication parameter
 - description, 195
- LocalSystem account
 - about, 23
 - options, 28
- LOCATION registry entry, 207
- Lock_rejected_rows property
 - description, 619
- LockedHeapPages property
 - description, 626
- locking conflicts
 - BLOCKING option, 556
 - BLOCKING_TIMEOUT option, 556
- LockName property
 - description, 619
- LockTablePages property
 - description, 630
- log files, 508, 509
 - changing name from Sybase Central, 507
 - checkpoint log, 322
 - rollback, 324
 - transaction, 305
- Log Transfer Manager, 107

log transfer manager utility
options, 482

Log Transfer Manager utility
syntax, 481

log translation utility, 344
about, 488
auditing, 394, 396
options, 490
syntax, 488

Logfile [LOG] connection parameter
about, 70
about connection parameters, 164
description, 183
introduction to connection parameters, 39
tips, 72
using default connection parameters, 51

LogFileFragments property
description, 630

LogFreeCommit property
description, 619, 630

logging off
keeping server running, 153

LOGIN_MODE option
description, 577
integrated logins, 84
obtaining value, 619
setting, 541

LOGIN_PROCEDURE option
description, 578
disallowing connections with RAISERROR, 578
obtaining value, 619
setting, 541

logins
integrated, 83, 84

LogName property
description, 630

LogWrite property
description, 619, 630

LOTUS file format, 570, 588

lower code page
about, 255

LTM
and interfaces file, 482
character set configuration, 428
character sets, 427
collations, 427, 428
configuration, 410, 425
configuration file, 417, 483
Open Client/Open Server character sets, 427
Open Client/Open Server collations, 427
starting, 417
supported operations, 422
transaction log management, 430
transaction log options, 508, 509

LTM configuration file
about, 425
character sets, 428
creating, 425
format, 425

LTM utility
syntax, 481

LTM_admin_pw parameter, 417
LTM configuration file, 484

LTM_admin_user parameter, 417
LTM configuration file, 484

LTM_charset parameter, 417
LTM configuration file, 428, 484

LTM_language parameter
LTM configuration file, 428

LTM_sortorder parameter
LTM configuration file, 428

LTMGeneration property
description, 630

LTMTrunc property
description, 630

M

–m option
database server, 144, 159
initialization [dbinit] utility, 466, 470
Log Translation [dbtran] utility, 488, 492
ping [dbping] utility, 494, 495
transaction log [dblog] utility, 507, 509
write file [dbwrite] utility, 531, 533

- M option
 - Log Transfer Manger [dbltn] utility, 481, 482
- MachineName property
 - description, 626
- MAGIC
 - USER_ESTIMATES option, 606
- Main Heap Bytes statistic, 617
- MainHeapBytes property
 - description, 626
- MainHeapPages property
 - description, 626
- maintenance user
 - primary sites, 412
 - replicate sites, 414
 - user ID, 419
- managing connected users, 368
- MapPages property
 - description, 630
- MAX_CURSOR_COUNT option
 - description, 581
 - obtaining value, 619
 - setting, 541
- MAX_HASH_SIZE option
 - description, 581
 - setting, 541
- MAX_PLANS_CACHED option
 - description, 581
 - obtaining values, 619
 - setting, 541
- MAX_STATEMENT_COUNT option
 - description, 583
 - obtaining value, 619
 - setting, 541
- MAX_WORK_TABLE_HASH_SIZE option
 - description, 582
 - setting, 541
- MaxCacheSize property
 - description, 626
- maximum size
 - database files, 636
 - databases, 636
- MaxIO property
 - description, 630
- MaxMessage property
 - description, 626
- MaxRead property
 - description, 630
- MaxWrite property
 - description, 630
- MDSR encryption algorithm
 - about, 399
 - initialization [dbinit] utility, 468
 - unload [dbunload] utility, 518
- media failures
 - about, 304
 - protection, 317
 - recovery, 306
 - recovery of, 340
- Mem Pages
 - Map Pages statistics, 615
- Mem Pages Lock Table statistics, 615
- Mem Pages Locked Heap statistics, 615
- Mem Pages Main Heap statistics, 615
- Mem Pages Procedure Definitions statistics, 615
- Mem Pages Relocatable statistics, 615
- Mem Pages Relocations/sec statistics, 615
- Mem Pages Rollback Log statistics, 615
- Mem Pages Trigger Definitions statistics, 615
- Mem Pages View Definitions statistics, 615
- membership
 - revoking group membership, 371
- memory
 - cache size, 127, 129, 130, 131
 - connection limit, 384
 - Java and memory use, 576
- memory pages statistics, 615
- memory use
 - and Java, 577
- Message Agent
 - transaction log management, 314

- Message property
 - description, 626
 - messages
 - language resource library, 253
 - MessageText property
 - description, 626
 - MessageTime property
 - description, 626
 - MessageWindowSize property
 - description, 626
 - Microsoft Access
 - ODBC configuration for, 57
 - TIMESTAMP comparison, 564
 - Microsoft Visual Basic
 - ODBC configuration for, 57
 - MIN_PASSWORD_LENGTH option
 - description, 583
 - obtaining value, 619
 - setting, 541
 - MIN_TABLE_SIZE_FOR_HISTOGRAM option
 - description, 583
 - obtaining value, 619
 - setting, 541
 - MinCacheSize property
 - description, 626
 - minimal column definitions, 423
 - mirror
 - transaction log, 305, 306, 346, 347
 - miscellaneous statistics, 617
 - MobiLink synchronization
 - backups, 312
 - purposes, 406
 - modulo operator
 - PERCENT_AS_COMMENT option, 590
 - monitoring
 - performance with histograms, 461
 - who is logged on, 450
 - monitoring performance, 610
 - MSDASQL OLE DB provider
 - about, 68
 - multibyte character sets
 - about, 259
 - using, 277
 - multibyte characters
 - encodings, 285
 - properties, 286
 - MultiByteCharSet property
 - description, 630
 - MultiPacketsReceived property
 - description, 626
 - MultiPacketsSent property
 - description, 626
 - multiple databases
 - DSEdit entries, 112
 - multiple record fetching
 - data source [dbdsn] utility, 455
 - ODBC configuration, 62
 - multi-processing
 - support, 14
 - multi-processor machines
 - option, 12
 - multi-processor support, 14
 - multi-tasking, 14
 - MyIP [ME] communication parameter
 - description, 195
- ## N
- n option
 - backup [dbbackup] utility, 439, 440
 - database server, 145
 - histogram [dbhist] utility, 461, 462
 - initialization [dbinit] utility, 466, 470
 - Log Translation [dbtran] utility, 488, 492
 - setting database name, 160
 - transaction log [dblog] utility, 507, 509
 - unload [dbunload] utility, 514, 519
 - name
 - database, 160
 - Name property
 - description, 619, 626, 630

- Named Pipes
 - about, 97
 - server configuration, 153
 - starting protocol, 16
 - supported protocols, 92
 - national language support
 - about, 250
 - collations, 287
 - multibyte character sets, 277
 - NDIS
 - drivers, 101
 - NDS
 - filenames, 125
 - NEAREST_CENTURY option
 - compatibility option, 545
 - description, 584
 - obtaining value, 619
 - negative permissions, 391
 - net.cfg file, 102
 - NetWare
 - bindery, 96
 - cache buffers, 127
 - database servers, 6, 124
 - enabling character set conversion, 131
 - file locations, 205
 - installation, 205
 - network adapter settings, 102
 - server performance, 127
 - subdirectories, 205
 - threading behavior, 14
 - network adapters
 - drivers, 101
 - network communications, 34, 35
 - command line options, 189
 - network connections
 - option, 16
 - network drives
 - database files, 8
 - network protocols
 - about, 91
 - network connections, 49
 - ODBC configuration, 60
 - supported, 92
 - troubleshooting, 101
 - network server
 - about, 4
 - connecting to, 49
 - encryption, 398
 - new databases, 29
 - newsgroups
 - technical support, xvii
 - NodeAddress property
 - description, 619
 - nogui option
 - Interactive SQL [dbisql] utility, 472, 475
 - NON_KEYWORDS option
 - compatibility option, 545
 - description, 584
 - obtaining value, 619
 - Novell client software, 101
 - NULL
 - ANSI behavior, 553
 - defining for output, 590
 - NULLS option, 585
 - Transact-SQL behavior, 553
 - NULLS option
 - description, 585
 - Interactive SQL settings, 548
 - Number property
 - description, 619
 - numeric precision
 - database option, 591
 - NumProcessorsAvail property
 - description, 626
 - NumProcessorsMax property
 - description, 626
- ## O
- o option
 - backup [dbbackup] utility, 439, 440
 - collation [dbcollat] utility, 443, 444
 - compression [dbshrink] utility, 448, 449
 - data source [dbdsn] utility, 451
 - database server, 146
 - erase [dberase] utility, 459, 460
 - information [dbinfo] utility, 463, 464

- initialization [dbinit] utility, 466, 470
- license [dblic] utility, 479, 480
- Log Transfer Manger [dbltm] utility, 481, 483
- Log Translation [dbtran] utility, 488, 492
- ping [dbping] utility, 494, 496
- server location [dblocate] utility, 498
- stop [dbstop] utility, 505, 506
- transaction log [dblog] utility, 507, 510
- uncompression [dbexpand] utility, 511, 512
- unload [dbunload] utility, 514, 519
- upgrade [dbupgrad] utility, 522, 525
- validation [dbvalid] utility, 527, 529
- write file [dbwrite] utility, 531, 533
- o server option
 - data source [dbdsn] utility, 455
- objects
 - qualified names, 376
- ODBC
 - Administrator, 53
 - connection parameters, 70, 164
 - connections, 38
 - data sources, 53
 - driver location, 74
 - initialization file for UNIX, 63
 - troubleshooting, 494
 - UNIX support, 63
 - Windows CE, 63
- ODBC Administrator
 - using, 53
- ODBC connectivity, 43
- ODBC data sources
 - about, 41
 - configuring, 56
 - creating, 53, 56
 - creating with dbdsn, 451
 - UNIX, 63
 - using with jConnect, 43
- ODBC option
 - Interactive SQL [dbisql] utility, 472, 475
- ODI
 - drivers, 101
- OEM code pages
 - about, 255
 - choosing, 274
- offline backups, 303
 - defined, 308
- OLE DB
 - ASAProv provider, 68
 - connecting, 68
 - providers, 68
- OmniConnect support, 107
- ON EXCEPTION RESUME clause
 - stored procedures, 587
- ON_CHARSET_CONVERSION_FAILURE option, 585
 - setting, 541
- ON_ERROR option
 - description, 586
 - Interactive SQL settings, 548
- ON_TSQL_ERROR option
 - compatibility option, 544, 545
 - description, 587
 - obtaining value, 619
- onerror option
 - Interactive SQL [dbisql] utility, 472, 475
- online backups, 303
 - defined, 308
- Open Client
 - compatibility options, 544
 - configuring, 110
 - interface, 106
 - options, 116
- Open Server
 - adding, 110, 410
 - address, 112
 - architecture, 106
 - configuring servers for JDBC, 115
 - connecting to, 411
 - deleting server entries, 115
 - renaming server entries, 115
 - starting, 109
 - system requirements, 108
- OPEN statement
 - QUERY_PLAN_ON_OPEN option, 594
- opening cursors, 594

OPTIMIZATION_GOAL option

description, 587

setting, 541

Optimization_goal property

description, 619

Optimization_level property

description, 619

options

about, 536

ALLOW_NULLS_BY_DEFAULT, 550

ANSI_BLANKS, 550

ANSI_CLOSE_CURSORS_ON_ROLLBACK,
551

ANSI_INTEGER_OVERFLOW, 551

ANSI_PERMISSIONS, 551

ANSI_UPDATE_CONSTRAINTS, 552

ANSINULL, 553

AUDITING, 553

AUTO_COMMIT, 554

AUTO_REFETCH, 554

AUTOMATIC_TIMESTAMP, 554

BACKGROUND_PRIORITY, 555

backup utility [dbbackup], 440

BELL, 555

BLOB_THRESHOLD, 555

BLOCKING, 556

BLOCKING_TIMEOUT, 556

CHAINED, 557

CHECKPOINT_TIME, 557

CIS_ROWSET_SIZE, 557

classification, 539

CLOSE_ON_ENDTRANS, 558

collation utility [dbcollat], 443

COMMAND_DELIMITER, 558

COMMIT_ON_EXIT, 559

compatibility, 545

COMPRESSION, 559

compression utility [dbshrink], 449

console utility [dbconsole], 450

CONTINUE_AFTER_RAISERROR, 560

CONVERSION_ERROR, 560

COOPERATIVE_COMMIT_TIMEOUT, 561

COOPERATIVE_COMMITS, 561

data source utility [dbdsn], 453

database server, 120

DATE_FORMAT, 562

DATE_ORDER, 564

debugger [dbprdbg] utility, 457

DEFAULT_TIMESTAMP_INCREMENT, 564

DELAYED_COMMIT_TIMEOUT, 564

DELAYED_COMMITS, 565

DELETE_OLD_LOGS, 566

deleting settings, 539

DESCRIBE_JAVA_FORMAT, 566

DIVIDE_BY_ZERO_ERROR, 566

ECHO, 567

erase utility [dberase], 459

ESCAPE_CHARACTER, 567

finding values, 537

FIRE_TRIGGERS, 568

FIRST_DAY_OF_WEEK, 568

FLOAT_AS_DOUBLE, 569

GLOBAL_DATABASE_ID, 569

histogram utility [dbhist], 462

information utility [dbinfo], 463

initial settings, 539

initialization utility [dbinit], 467

INPUT_FORMAT, 570

Interactive SQL, 548

Interactive SQL utility [dbisql], 474

ISOLATION_LEVEL, 571

ISQL_COMMAND_TIMING, 572

ISQL_ESCAPE_CHARACTER, 572

ISQL_FIELD_SEPARATOR, 573

ISQL_LOG, 574

ISQL_PLAN, 574

ISQL_PRINT_RESULT_SET, 575

ISQL_QUOTE, 576

JAVA_HEAP_SIZE, 576

JAVA_INPUT_OUTPUT, 577

JAVA_NAMESPACE_SIZE, 577

language utility [dblang], 478

license utility [dblic], 480

list of all, 550

list of compatibility options, 544, 545

list of database options, 541

list of Interactive SQL options, 548

list of replication options, 547

log transfer manager utility [dbltn], 482

log translation utility [dbtran], 490

LOGIN_MODE, 577

LOGIN_PROCEDURE, 578

MAX_CURSOR_COUNT, 581

MAX_HASH_SIZE, 581

MAX_PLANS_CACHED, 581

MAX_STATEMENT_COUNT, 583

MAX_WORK_TABLE_HASH_SIZE, 582

MIN_PASSWORD_LENGTH, 583

MIN_TABLE_SIZE_FOR_HISTOGRAM, 583

NEAREST_CENTURY option, 584

NON_KEYWORDS, 584
 NULLS, 585
 ON_CHARSET_CONVERSION_FAILURE, 585
 ON_ERROR, 586
 ON_TSQL_ERROR, 587
 Open Client, 116
 OPTIMIZATION_GOAL, 587
 OUTPUT_FORMAT, 588
 OUTPUT_LENGTH, 590
 OUTPUT_NULLS, 590
 PERCENT_AS_COMMENT, 590
 ping utility [dbping], 494
 PINNED_CURSOR_PERCENT_OF_CACHE, 591
 PRECISION, 591
 PREFETCH, 592
 PRESERVE_SOURCE_FORMAT, 593
 PREVENT_ARTICLE_PKEY_UPDATE, 593
 PUBLIC options, 538
 QUALIFY_OWNERS, 593
 QUERY_PLAN_ON_OPEN, 594
 QUOTE_ALL_IDENTIFIERS, 594
 QUOTED_IDENTIFIER, 594
 RECOVERY_TIME, 595
 REPLICATE_ALL, 595
 REPLICATION_ERROR, 596
 RETURN_DATE_TIME_AS_STRING, 596
 RETURN_JAVA_AS_STRING, 597
 RI_TRIGGER_TIME, 597
 ROW_COUNTS, 597
 SCALE, 598
 scope and duration of database options, 537
 server location utility [dblocate], 498
 service creation utility [dbsvc], 500
 setting database options, 536
 setting in Interactive SQL, 548
 setting temporary, 537
 setting user and group options, 356
 SORT_COLLATION, 598
 spawn utility [dbspawn], 504
 SQL_FLAGGER_ERROR_LEVEL, 599
 SQL_FLAGGER_WARNING_LEVEL, 599
 startup settings, 116
 STATISTICS, 600
 stop utility [dbstop], 505
 STRING_RTRUNCATION, 600
 SUBSCRIBE_BY_REMOTE, 600
 SUPPRESS_TDS_DEBUGGING, 601
 TDS_EMPTY_STRING_IS_NULL, 601
 TIME_FORMAT, 601

TIME_ZONE_ADJUSTMENT, 602
 TIMESTAMP_FORMAT, 602
 transaction log utility [dblog], 509
 TRUNCATE_DATE_VALUES, 604
 TRUNCATE_TIMESTAMP_VALUES, 604
 TRUNCATE_WITH_AUTO_COMMIT, 605
 TRUNCATION_LENGTH, 605
 TSQL_HEX_CONSTANT, 606
 TSQL_VARIABLES, 606
 uncompression utility [dbexpand], 512
 unload utility [dbunload], 516
 upgrade utility [dbupgrad], 523
 USER_ESTIMATES, 606
 validation utility [dbvalid], 528
 VERIFY_ALL_COLUMNS, 607
 VERIFY_THRESHOLD, 608
 WAIT_FOR_COMMIT, 608
 write file utility [dbwrite], 532

-os option
 database server, 146
 Log Transfer Manger [dbltn] utility, 481, 483

-ot option
 Log Transfer Manger [dbltn] utility, 481, 483

OUTPUT_FORMAT option
 description, 588
 Interactive SQL settings, 548

OUTPUT_LENGTH option
 description, 590
 Interactive SQL settings, 548

OUTPUT_NULLS option
 description, 590
 Interactive SQL settings, 548

OVERRIDE-MAGIC
 USER_ESTIMATES option, 606

owners
 about, 353

P

-p option
 data source [dbdsn] utility, 451, 455
 database server, 147
 initialization [dbinit] utility, 466, 470
 license [dblic] utility, 479, 480
 service creation [dbsvc] utility, 499, 501

- spawn [dbspawn] utility, 503, 504
- unload [dbunload] utility, 514, 519
- packet encryption
 - data source [dbdsn] utility, 454
- packet size
 - data source [dbdsn] utility, 455
 - limiting, 147, 148
- PacketSize property
 - description, 619
- PacketsReceived property
 - description, 619, 626
- PacketsReceivedUncomp property
 - description, 619, 626
- PacketsSent property
 - description, 619, 626
- PacketsSentUncomp property
 - description, 619, 626
- page size
 - choosing, 470
 - databases, 466
 - maximum allowed, 142
 - option, 13
- page usage
 - information [dbinfo] utility, 464
- PageRelocations property
 - description, 630
- pages
 - displaying usage in database files, 463
- PageSize property
 - description, 626, 630
- parallel execution
 - processors, 143
 - threads, 144
- Password [PWD] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 184
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- passwords
 - changing, 358
 - default, 352
 - encrypting, 177
 - length, 403
 - LTM configuration file, 484
 - minimum length, 583
 - ODBC configuration, 58
 - security features, 391
 - security tips, 403
 - utility database, 228
- PATH environment variable
 - description, 210
- pc option
 - database server, 147
 - ping [dbping] utility, 494, 495
- pd option
 - ping [dbping] utility, 494, 495
- PeakCacheSize property
 - description, 626
- percent sign
 - modulo or comment, 590
- PERCENT_AS_COMMENT option
 - compatibility option, 545
 - description, 590
 - obtaining value, 619
- performance
 - cache size, 127, 129, 130, 131
 - compression, 98
 - disk fragmentation, 222
 - encrypted databases, 402
 - improving, 327
 - LTM, 426
 - monitoring, 610
 - options, 12
 - PowerBuilder DataWindow, 587
 - prefetch, 592
 - primary keys, 319
 - result sets, 587
 - server options, 6
 - setting priorities, 555
 - TCP/IP, 93
 - transaction log, 305, 319
 - transaction log mirrors, 306
 - TRUNCATE TABLE statement, 605

Performance Monitor
Windows, 610

permissions

- about, 391
- conflicts, 383
- connect, 357
- DBA authority, 352
- file administration statements, 229
- granting passwords, 357
- granting remote, 365
- group, 369
- group membership, 370
- groups, 354, 372
- individual, 357
- inheriting, 363, 369
- integrated login permissions, 84
- listing, 385
- loading data, 141
- managing, 351
- negative, 391
- options, 12
- overview, 352
- passwords, 358
- procedures, 363
- RESOURCE authority, 353, 358
- revoking, 366
- revoking remote, 365
- scheme, 390
- security features, 390
- setting table permissions, 359
- tables, 354, 359
- the right to grant, 363
- triggers, 353, 365
- unloading data, 141
- views, 354, 361, 378
- WITH GRANT OPTION, 363

personal server
about, 4

physical layer
troubleshooting, 102

physical limitations
Adaptive Server Anywhere, 636

ping utility
about, 494
options, 494
syntax, 494
TCP/IP, 102

testing networks, 34
testing Open Client, 114

pinging
servers, 494

PINNED_CURSOR_PERCENT_OF_CACHE
database option, 591

planning
backups, 309, 315

plans
caching of, 581

Platform property
description, 626

PlatformVer property
description, 626

plug-ins
registry settings, 214

polling frequency, 30

port number
database server, 198
TCP/IP, 109, 155, 198

-port option
Interactive SQL [dbisql] utility, 472, 475

PowerBuilder DataWindow
query performance, 587

PRECISION option
description, 591
setting, 541

PRECISION property
description, 619

PREFETCH option
description, 592
DisableMultiRowFetch, 175
setting, 541

PREFETCH property
description, 619

PrefetchBuffer [PBUF] connection parameter
about, 70
about connection parameters, 164
description, 185
introduction to connection parameters, 39
tips, 72
using default connection parameters, 51

PreFetchOnOpen communication parameter
description, 196

PrefetchRows [PROWS] connection parameter
about, 70
about connection parameters, 164
description, 186
introduction to connection parameters, 39
tips, 72
using default connection parameters, 51

prepared statements
connection limits, 384, 583

Prepares property
description, 619

PrepStmt property
description, 619

PRESERVE_SOURCE_FORMAT option
description, 593
setting, 541

PreserveSource property
description, 630

PREVENT_ARTICLE_PKEY_UPDATE option
description, 593
setting, 541

primary keys
transaction log, 319

primary sites
adding Replication Server information, 412
creating, 409
Replication Server, 407, 408
uses of LTM at, 408

priority
process, 138

ProcedurePages property
description, 630

procedures
and Adaptive Server Anywhere LTM, 422
permissions, 363
permissions for creating, 353
replicating, 424
security, 378

ProcessCPU property
description, 626

ProcessCPUSystem property
description, 626

ProcessCPUUser property
description, 626

processors
multiple, 12, 143
number used, 12, 143

ProductName property
description, 626

ProductVersion property
description, 626

programming interfaces
connections, 38

properties
list of connection properties, 619
list of database properties, 630
list of server properties, 626
syntax, 618
syntax of connection-level properties, 618
syntax of database-level properties, 630
syntax of server-level properties, 625

property functions
syntax, 618

protocol stacks
TCP/IP, 192

protocols
about, 91
disabling shared memory, 150
option, 16
supported, 16, 92
troubleshooting, 101

providers
ASAProv, 68
MSDASQL, 68
OLE DB, 68

–ps option
ping [dbping] utility, 494, 496

–pt option
database server, 148

PUBLIC group, 374

PUBLIC options
about, 538
DBA authority required, 538

Q

–q option

- backup [dbbackup] utility, 439, 440
- collation [dbcollat] utility, 443, 444
- compression [dbshrink] utility, 448, 449
- data source [dbdsn] utility, 451, 454, 455
- database server, 148
- erase [dberase] utility, 459, 460
- information [dbinfo] utility, 463, 464
- initialization [dbinit] utility, 466, 470
- Interactive SQL [dbisql] utility, 472, 475
- language [dblang] utility, 477, 478
- license [dblic] utility, 479, 480
- Log Transfer Manger [dbltn] utility, 481, 483
- Log Translation [dbtran] utility, 488, 492
- ping [dbping] utility, 494, 496
- server location [dblocate] utility, 498
- service creation [dbsvc] utility, 499, 501
- spawn [dbspawn] utility, 503, 504
- stop [dbstop] utility, 505, 506
- transaction log [dblog] utility, 507, 510
- uncompression [dbexpand] utility, 511, 512
- unload [dbunload] utility, 514, 520
- upgrade [dbupgrad] utility, 522, 525
- validation [dbvalid] utility, 527, 529
- write file [dbwrite] utility, 531, 533

–qp option

- database server, 148

–qq option

- data source [dbdsn] utility, 454

qualified names

- database objects, 376
- tables, 373

QUALIFY_OWNERS option

- description, 593
- replication option, 547

qualify_table_owner parameter

- LTM configuration file, 484

Query Low Memory Strategy statistic, 617

QUERY_PLAN_ON_OPEN option

- compatibility option, 545
- description, 594
- obtaining value, 619

QueryCachePages property

- description, 619, 630

QueryLowMemoryStrategy property

- description, 619, 630

questions

- character sets, 251

quiet mode

- backup [dbbackup] utility, 440
- collation [dbcollat] utility, 444
- compression [dbshrink] utility, 449
- data source [dbdsn] utility, 455
- erase [dberase] utility, 460
- information [dbinfo] utility, 464
- initialization [dbinit] utility, 470
- Interactive SQL [dbisql] utility, 475
- language [dblang] utility, 478
- license [dblic] utility, 480
- Log Transfer Manger [dbltn] utility, 483
- Log Translation [dbtran] utility, 492
- ping [dbping] utility, 496
- server location [dblocate] utility, 498
- spawn [dbspawn] utility, 504
- stop [dbstop] utility, 506
- transaction log [dblog] utility, 510
- uncompression [dbexpand] utility, 512
- unload [dbunload] utility, 520
- upgrade [dbupgrad] utility, 525
- validation [dbvalid] utility, 529
- write file [dbwrite] utility, 533

QuittingTime property

- description, 626

QUOTE_ALL_IDENTIFIERS option

- description, 594
- replication option, 547, 594

QUOTED_IDENTIFIER option

- compatibility option, 544, 545
- description, 594
- obtaining value, 619
- Open Client, 116

R

–r option

- backup [dbbackup] utility, 439, 440
- data source [dbdsn] utility, 451, 455
- database, 161
- database server, 149
- Log Translation [dbtran] utility, 488, 492, 493

- transaction log [dblog] utility, 507, 510
- unload [dbunload] utility, 514, 520
- RAISERROR**
 - behavior of, 560
- RAISERROR** statement
 - CONTINUE_AFTER_RAISERROR option, 560
 - ON_TSQL_ERROR option, 587
- RAS**
 - dial-up networking, 94
- read-only
 - databases, 13, 149, 161
- read-only media
 - modifying databases, 224
- ReadOnly property
 - description, 630
- rebuild command line utility
 - syntax, 497
- rebuild utility
 - about, 497
- rebuilding databases, 497
- ReceiveBufferSize [RCVBUFSZ] communication
 - parameter
 - description, 197
- recommended collations
 - UNIX, 271
 - Windows, 271
- recovery
 - about, 300
 - distributed transactions, 151
 - maximum time, 143
 - media failure, 306, 340
 - option, 13
 - rapid, 340
 - server options, 157
 - system failure, 303
 - transaction log, 305
 - transaction log mirrors, 306
 - uncommitted changes, 344
 - urgency, 325
- Recovery I/O Estimate statistic, 611
- recovery mode
 - Log Transfer Manger [dbltm] utility, 482
- recovery statistics, 611
- Recovery Urgency statistic, 611
- RECOVERY_TIME option
 - description, 595
 - obtaining value, 619
 - setting, 541
 - using, 325
- RecoveryUrgency property
 - description, 630
- redescribing cursors, 58
- REFERENCES permissions, 359
- referential integrity
 - triggers, 597
- RegisterBindery [REGBIN] communication
 - parameter
 - description, 197
- registry
 - about, 213
 - environment variables, 208
 - language settings, 214, 477
 - location setting, 214
 - modifying, 208
 - setting SQLREMOTE environment variable, 211
 - Sybase Central, 214
 - TEMP environment variable, 212
 - tools location setting, 214
 - Windows CE, 214
 - Windows services, 213
- RelocatableHeapPages property
 - description, 630
- remote data access
 - CIS_ROWSET_SIZE option, 557
- remote permissions, 365
- RemoteTrunc property
 - description, 630
- removing
 - users from groups, 371
- rep_func parameter
 - LTM configuration file, 484
- replicate minimal columns
 - support for, 423
- REPLICATE ON clause, 413

- replicate sites
 - adding Replication Server information, 414
 - creating, 409
 - Replication Server, 407
 - uses of LTM at, 408
- REPLICATE_ALL option
 - description, 595
 - obtaining value, 619
 - replication option, 547
- replication
 - backup procedures, 314, 429, 430
 - benefits, 406
 - buffering, 426
 - creating a replication definition for Replication Server, 416
 - dbcc, 508, 509
 - defining, 422
 - entire databases, 430
 - introduction, 406
 - list of options, 547
 - Log Transfer Manager, 481
 - of procedures, 424
 - of trigger actions, 567, 568
 - procedures, 424
 - replication server, 508, 509
 - Replication Server, 407
 - setting flag, 413
 - stored procedures, 425
 - transaction log management, 314, 429, 430
- Replication Agent
 - backups, 312
 - dbltm command-line syntax, 481
- replication options
 - BLOB_THRESHOLD, 555
 - classification, 539
 - COMPRESSION, 559
 - DELETE_OLD_LOGS, 566
 - initial settings, 539
 - list of, 547
 - QUALIFY_OWNERS, 593
 - QUOTE_ALL_IDENTIFIERS, 594
 - REPLICATE_ALL, 595
 - REPLICATION_ERROR, 596
 - SUBSCRIBE_BY_REMOTE, 600
 - VERIFY_ALL_COLUMNS, 607
 - VERIFY_THRESHOLD, 608
- Replication Server
 - Adaptive Server Anywhere character sets, 421
 - Adaptive Server Anywhere collations, 421
 - backup procedures, 429
 - creating a connection, 415, 416
 - creating a replication definition, 416
 - creating a subscription, 418
 - Log Transfer Manager, 481
 - preparing Adaptive Server Anywhere databases, 414
 - primary sites, 408, 415
 - purposes, 406
 - replicate sites, 407, 416
 - replicating an entire database, 430
 - replicating procedures, 424, 425
 - rssetup.sql script, 412
 - starting an Adaptive Server Anywhere server, 410
 - support, 107
 - supported versions, 408
 - transaction log management, 429
- Replication Server Adaptive Server Anywhere configuration, 419, 421
- REPLICATION_ERROR option
 - description, 596
 - replication option, 547
- Req property
 - description, 626
- ReqType property
 - description, 619
- Request Queue Waits/sec statistic, 617
- RequestLogFile property
 - description, 626
- RequestLogging property
 - description, 626
- RequestQueueWait property
 - description, 626
- requests
 - about, 14
- Requests Active statistic, 616
- Requests statistic, 616
- Requests Unscheduled statistic, 616

RESOURCE authority

- about, 353
- granting, 358
- not inheritable, 369

resource governor

- cursors, 581
- statements, 583

restore database wizard

- using, 343

return codes

- Interactive SQL, 475

RETURN_DATE_TIME_AS_STRING option

- description, 596
- setting, 541

RETURN_JAVA_AS_STRING option

- description, 597
- setting, 541

revoking

- group membership, 371
- permissions, 366
- remote permissions, 365

–rg option

- service creation [dbsvc] utility, 499, 501

RI_TRIGGER_TIME option

- compatibility option, 545
- description, 597
- obtaining value, 619

Rijndael encryption algorithm

- about, 399
- initialization [dbinit] utility, 468
- unload [dbunload] utility, 518

Rlbc property

- description, 619

rollback logs

- about, 324

ROLLBACK statement

- cursors, 551
- log, 324

RollbackLogPages property

- description, 619, 630

rounding

- SCALE option, 598

routers

- broadcasting over, 192

row counts, 597**ROW_COUNTS option**

- description, 597
- obtaining value, 619
- setting, 541

–rs option

- service creation [dbsvc] utility, 499, 501

RS parameter, 417

- LTM configuration file, 484

RS_pw parameter, 417

- LTM configuration file, 484

RS_source_db parameter, 417

- LTM configuration file, 484

RS_source_ds parameter, 417

- LTM configuration file, 484

RS_user parameter, 417

- LTM configuration file, 484

rssetup.sql script, 412, 414

- about, 419
- preparing to run, 419
- running, 421

–rsu option

- Log Translation [dbtran] utility, 488, 492

running the database server, 3

S

–s option

- database server, 149
- Log Transfer Manger [dbltm] utility, 481, 483
- Log Translation [dbtran] utility, 488, 492
- service creation [dbsvc] utility, 499, 501
- write file [dbwrite] utility, 531, 533

–S option

- Log Transfer Manger [dbltm] utility, 481, 482

sa_conn_properties

- using, 537

sample database

- about asademo.db, xvi

- saving
 - statements, 574
- sb option
 - database server, 150
- sc option
 - database server, 150
- SCALE option
 - description, 598
 - obtaining value, 619
 - setting, 541
- scan_retry parameter, 417
 - LTM configuration file, 484
- schedules
 - about, 234
 - defined, 232
 - internals, 242
- scheduling
 - about, 231, 232, 234
 - backups, 309, 315
 - server shutdown, 152
- schemas
 - unloading definition, 519
- scripts directory
 - about, 205
- search conditions
 - USER_ESTIMATES option, 606
- SearchBindery [BINSEARCH] communication
 - parameter
 - description, 197
- security
 - about, 351
 - adding simple encryption to command files, 446
 - AES encryption, 399
 - auditing, 393, 394, 553
 - C2, 150, 169
 - C2 guidelines, 388
 - creating databases, 397
 - database server, 397, 403
 - deleting databases, 397
 - encrypting communications, 135, 137, 159, 177, 179
 - encrypting database files, 399
 - encryption, 398
 - event example, 238
 - file access, 141
 - file hiding [dbfhide] utility, 446
 - integrated logins, 87, 89, 390
 - loading data, 397
 - MDSR encryption, 399
 - minimum password length, 583
 - overview, 388
 - passwords, 391
 - procedures, 363, 378
 - Rijndael encryption, 399
 - server command line, 388
 - services, 29
 - system functions, 403
 - temporary file, 209
 - tips, 403
 - unloading data, 397
 - utility database, 228
 - views, 378
- SecurityManager
 - JAVA_INPUT_OUTPUT option, 577
- SELECT permissions, 359
- selectivity estimates
 - USER_ESTIMATES option, 606
- SendBufferSize [SNDBUFSZ] communication
 - parameter
 - description, 197
- SendFail property
 - description, 626
- SeparateCheckpointLog property
 - description, 630
- SeparateForeignKeys property
 - description, 630
- server address
 - DSEdit, 112
- Server Idle Waits/sec statistic, 611
- server information
 - asasrv.ini, 81
- server location utility, 498
 - options, 498
 - syntax, 498
- server messages
 - displaying, 148
 - output to file, 146
 - suppressing performance messages, 148

- server name, 145
 - case sensitivity, 11
 - data source [dbdsn] utility, 455
 - ODBC configuration, 59
 - stop [dbstop] syntax, 505

- server options
 - database, 159
 - options, 10
 - recovery, 157

- server properties
 - ActiveReq, 626
 - AvailIO, 626
 - BufferMisses, 626
 - BuildChange, 626
 - BuildClient, 626
 - BuildReproducible, 626
 - BytesReceived, 626
 - BytesReceivedUncomp, 626
 - BytesSent, 626
 - BytesSentUncomp, 626
 - C2, 626
 - CacheHitsEng, 626
 - CacheReplacements, 626
 - CharSet, 626
 - CompanyName, 626
 - ConnsDisabled, 626
 - CurrentCacheSize, 626
 - DefaultCollation, 626
 - FreeBuffers, 626
 - IdleTimeout, 626
 - IsIQ, 626
 - IsJavaAvailable, 626
 - IsNetworkServer, 626
 - IsRuntimeServer, 626
 - JavaGlobFix, 626
 - Language, 626
 - LegalCopyright, 626
 - LegalTrademarks, 626
 - LicenseCount, 626
 - LicensedCompany, 626
 - LicensedUser, 626
 - LicenseType, 626
 - list of, 626
 - LivenessTimeout, 626
 - LockedHeapPages, 626
 - MachineName, 626
 - MainHeapBytes, 626
 - MainHeapPages, 626
 - MaxCacheSize, 626
 - MaxMessage, 626

- MessageText, 626
- MessageTime, 626
- MessageWindowSize, 626
- MinCacheSize, 626
- MultiPacketsReceived, 626
- MultiPacketsSent, 626
- Name, 626
- NumProcessorsAvail, 626
- NumProcessorsMax, 626
- PacketsReceived, 626
- PacketsReceivedUncomp, 626
- PacketsSent, 626
- PacketsSentUncomp, 626
- PageSize, 626
- PeakCacheSize, 626
- Platform, 626
- PlatformVer, 626
- ProcessCPU, 626
- ProcessCPUSystem, 626
- ProcessCPUUser, 626
- ProductName, 626
- ProductVersion, 626
- QuittingTime, 626
- reporting, 496
- Req, 626
- RequestLogFile, 626
- RequestLogging, 626
- RequestQueueWait, 626
- SendFail, 626
- ServerIdleWaits, 626
- StartTime, 626
- Tempdir, 626
- TimeZoneAdjustment, 626
- TotalBuffers, 626
- UnschReq, 626

- server side
 - setting encryption parameters, 135, 137, 159

- ServerIdleWaits property
 - description, 626

- ServerName [ENG] connection parameter
 - character sets, 71
 - description, 187

- ServerPort (PORT) communication parameter, 109

- ServerPort [PORT] communication parameter
 - description, 198

- servers
 - locating, 498
 - managing, 499

- properties of, 625
- starting from batch files, 503
- service creation utility
 - about, 499
 - options, 500
 - syntax, 499
- service creation wizard
 - using, 24
- services
 - about, 24, 29, 31
 - account, 28
 - adding, 24, 502
 - adding new databases, 29
 - configuring, 26
 - database server for Windows, 21
 - dependencies, 32, 33, 501
 - eligible programs, 23
 - event log, 125
 - executable file, 29
 - failure to start, 27
 - groups, 32
 - icon on the desktop, 29
 - listing, 501
 - managing, 24
 - multiple, 31
 - options, 27
 - parameters, 26
 - pausing, 31
 - registry settings, 213
 - removing, 24, 26
 - security, 29
 - service creation [dbsvc] utility, 499
 - service manager, 31
 - setting dependencies, 501
 - setting group dependencies, 501
 - setting type, 501
 - starting, 31
 - starting order, 33
 - startup options, 27
 - stopping, 31
 - Windows, 23, 155, 499
 - Windows NT Control Panel, 31
- SET OPTION statement
 - Interactive SQL syntax, 548
 - using, 536
- SET TEMPORARY OPTION statement
 - Interactive SQL syntax, 548
 - using, 536
- setting
 - database options, 541
 - temporary options, 537
- setup script
 - about, 419
 - preparing to run, 419
 - running, 421
- seven-bit characters
 - about, 255
- shared memory
 - about, 16
 - and CommLinks connection parameter, 169
 - server configuration, 153
- shutdown
 - specifying time, 152
- simple encryption, 399
- single-byte character sets
 - about, 254
- SMP
 - number of processors, 12, 143
- software
 - dblog, 508
 - dbshrink, 449
 - dbtran, 490
 - Interactive SQL, 473
 - licensing database servers, 479
 - rebuilding, 497
 - version, 153
- software license
 - licensing servers, 479
- software version
 - database server, 153
- sort order
 - collations, 250
 - definition, 253
- SORT_COLLATION option
 - description, 598
 - setting, 541
- sorting
 - collation file, 283
- sp_setreplicate procedure
 - about, 413, 424

- sp_setrepproc procedure
 - about, 424
- spawn utility
 - about, 503
 - options, 504
 - syntax, 503
- specifying
 - drivers, 43
- SPX
 - about, 96
 - Host communication parameter, 194
 - starting, 16
 - supported protocols, 92
- SQL Anywhere Studio
 - documentation, x
- SQL compatibility
 - options, 544
- SQL file format, 588
- SQL Remote
 - backup procedures, 314
 - backups, 312
 - options, 547
 - purposes, 406
 - transaction log management, 314
- SQL standards
 - compatibility options, 545
 - UPDATE statement, 492
- SQL statements
 - capturing recently-prepared, 156
- SQL.ini file, 410
 - configuring, 110
- SQL/92 compliance
 - SQL_FLAGGER_ERROR option, 599
 - updates, 552
- SQL_database parameter, 417
 - LTM configuration file, 484
- SQL_FLAGGER_ERROR_LEVEL option
 - compatibility option, 545
 - description, 599
 - obtaining value, 619
- SQL_FLAGGER_WARNING_LEVEL option
 - compatibility option, 545
 - description, 599
 - obtaining value, 619
- SQL_pw parameter, 417
 - LTM configuration file, 484
- SQL_server parameter, 417
 - LTM configuration file, 484
- SQL_user parameter, 417
 - LTM configuration file, 484
- SQLCONNECT environment variable
 - connections, 52
 - description, 210
- SQLDA
 - and ANSI_BLANKS, 550
- SQLLOCALE environment variable
 - about, 268, 278
 - description, 211
 - setting, 289
- SQLPATH environment variable
 - description, 211
- SQLREMOTE environment variable
 - description, 211
- sr option
 - Log Translation [dbtran] utility, 488, 493
- stack overflow
 - errors, 143
- stack size
 - external functions, 140
 - maximum, 143
- start line
 - ODBC configuration, 59
- StartLine [START] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 187
 - embedded databases, 48
 - introduction to connection parameters, 39
 - options, 10
 - tips, 72
 - using default connection parameters, 51
- StartTime property
 - description, 626

Statement Prepares statistic, 616

statements

- ALTER TABLE, 413
- COMMIT, 418
- CREATE DATABASE, 226, 229
- DELETE, 422
- INSERT, 422
- UPDATE, 422

Statements statistic, 616

statistics

- Avail IO, 617
- cache, 611
- Cache Hits/sec, 611
- Cache Reads Index Interior/sec, 611
- Cache Reads Index Leaf/sec, 611
- Cache Reads Table/sec, 611
- Cache Reads Total Pages/sec, 611
- Cache Size Current, 611
- Cache Size Maximum, 611
- Cache Size Minimum, 611
- Cache Size Peak, 611
- checkpoint and recovery, 611
- Checkpoint Flushes/sec, 611
- Checkpoint Log, 611
- Checkpoint Urgency, 611
- Checkpoints/sec, 611
- Comm Buffer Misses, 612
- Comm Bytes Received Uncompressed/sec, 612
- Comm Bytes Sent Uncompressed/sec, 612
- Comm Bytes Sent/sec, 612
- Comm Free Buffers, 612
- Comm Multi-packets Received/sec, 612
- Comm Multi-packets Sent/sec, 612
- Comm Packets Received Uncompressed/sec, 612
- Comm Packets Received/sec, 612
- Comm Packets Sent Uncompressed/sec, 612
- Comm Packets Sent/sec, 612
- Comm Send Fails/sec, 612
- Comm TotalBuffers, 612
- communications, 612
- Context Switches, 617
- Cursors, 616
- Cursors Open, 616
- Disk Active I/O, 613
- disk I/O, 613
- Disk Maximum I/O, 613
- disk read, 614
- Disk Reads Active, 614
- Disk Reads Index interior/sec, 614

- Disk Reads Index leaf/sec, 614
- Disk Reads Maximum Active, 614
- Disk Reads Table/sec, 614
- Disk Reads Total Pages/sec, 614
- disk write, 614
- Disk Write Active, 614
- Disk Write Commit Files/sec, 614
- Disk Write Database Extends/sec, 614
- Disk Write Maximum Active, 614
- Disk Write Pages/sec, 614
- Disk Write Temp Extends/sec, 614
- Disk Write Transaction Log Group Commits, 614
- Disk Write Transaction Log/sec, 614
- Idle Actives/sec, 611
- Idle Checkpoint Time, 611
- Idle Checkpoints/sec, 611
- Idle Writes/sec, 611
- index, 615
- Index Adds/sec, 615
- Index Full Compares/sec, 615
- Index Lookups/sec, 615
- Java VM, 615
- JVM Global Fixed Size, 615
- JVM Heap Size, 615
- JVM Namespace Size, 615
- Main Heap Bytes, 617
- Mem Pages Lock Table, 615
- Mem Pages Locked Heap, 615
- Mem Pages Main Heap, 615
- Mem Pages Map Pages, 615
- Mem Pages Procedure Definitions, 615
- Mem Pages Relocatable, 615
- Mem Pages Relocations/sec, 615
- Mem Pages Rollback Log, 615
- Mem Pages Trigger Definitions, 615
- Mem Pages View Definitions, 615
- memory pages, 615
- miscellaneous, 617
- monitoring performance, 610
- Query Low Memory Strategy, 617
- Recovery I/O Estimate, 611
- Recovery Urgency, 611
- Request Queue Waits/sec, 617
- Requests, 616
- Requests Active, 616
- Requests Unscheduled, 616
- Server Idle Waits/sec, 611
- Statement Prepares, 616
- Statements, 616
- Temporary Table Pages, 617

- Transaction Commits, 616
- Transaction Rollbacks, 616
- STATISTICS option
 - description, 600
 - Interactive SQL settings, 548
- stop utility
 - about, 505
 - options, 505
 - permissions, 140
 - syntax, 505
- stopping
 - databases, 505
- stored procedures
 - security features, 388
- STRING_RTRUNCATION option
 - compatibility option, 545
 - description, 600
 - obtaining value, 619
- strings
 - and host variables, 550
- strong encryption
 - AES algorithm, 399
 - client/server communications, 135, 137, 159, 177, 179
 - database files, 399
 - initialization [dbinit] utility, 468
 - Rijndael, 399
 - unload [dbunload] utility, 518
 - using MDSR algorithm, 399
- subdirectories
 - NetWare, 205
 - Windows CE, 205
- SUBSCRIBE_BY_REMOTE option
 - description, 600
 - replication option, 547
- subscriptions
 - creating for Replication Server, 418
- support
 - newsgroups, xvii
- SUPPRESS_TDS_DEBUGGING option
 - description, 601
 - obtaining value, 619
 - setting, 541
- switches
 - backup utility [dbbackup], 440
 - collation utility [dbcollat], 443
 - compression utility [dbshrink], 449
 - console utility [dbconsole], 450
 - data source utility [dbdsn], 453
 - debugger [dbprdbg] utility, 457
 - erase utility [dberase], 459
 - histogram utility [dbhist], 462
 - information utility [dbinfo], 463
 - initialization utility [dbinit], 467
 - Interactive SQL utility [dbisql], 474
 - language utility [dblang], 478
 - license utility [dblic], 480
 - log transfer manager utility [dbltm], 482
 - log translation utility [dbtran], 490
 - ping utility [dbping], 494
 - server location utility [dblocate], 498
 - service creation utility [dbsvc], 500
 - spawn utility [dbspawn], 504
 - stop utility [dbstop], 505
 - transaction log utility [dblog], 509
 - uncompression utility [dbexpand], 512
 - unload utility [dbunload], 516
 - upgrade utility [dbupgrad], 523
 - validation utility [dbvalid], 528
 - write file utility [dbwrite], 532
- Sybase Central
 - backing up databases, 438
 - changing log file names, 507
 - compressing databases, 448, 456
 - creating custom collations, 442
 - creating databases, 465
 - creating write files, 530
 - erasing databases, 458
 - managing services, 24
 - registry settings, 214
 - starting Interactive SQL from, 472
 - syntax highlighting, 590
 - uncompressing database files, 511
 - unloading databases, 514
 - upgrading databases, 521
 - validating databases, 526
- SYBASE environment variable
 - description, 211
 - DSEdit, 111
- sybinit utility
 - about, 108
- sybping, 411

- SyncTrunc property
 - description, 630
 - syntax
 - connection properties, 618
 - database properties, 618
 - database-level properties, 630
 - server-level properties, 625
 - syntax errors
 - joins, 567
 - syntax highlighting
 - comments, 590
 - SYS group, 374
 - SYSCOLAUTH view
 - permissions, 386
 - SYSCOLLATION table
 - collation files, 282
 - SYSCOLPERM table
 - permissions, 385
 - SYSDDUMMY table
 - permissions, 385
 - SYSGROUP table
 - permissions, 385
 - SYSGROUPS view
 - permissions, 386
 - SYSINFO table
 - collation files, 282
 - syslog
 - user ID for, 149
 - SYSOPTION system table
 - storing Interactive SQL options, 548
 - SYSPROCAUTH view
 - permissions, 386
 - SYSROCPERM table
 - permissions, 385
 - SYSTABAUTH view
 - permissions, 386
 - SYSTABLEPERM table
 - permissions, 385
 - system events
 - about, 236
 - defined, 232
 - system failures
 - about, 304
 - recovery, 303
 - system objects
 - and the dbo user, 513
 - system tables
 - character sets, 282
 - national languages, 282
 - permissions, 385
 - PRESERVE_SOURCE_FORMAT, 593
 - PREVENT_ARTICLE_PKEY_UPDATE, 593
 - source column, 593
 - SYSCOLLATION, 282
 - SYSINFO, 282
 - users and groups, 385
 - system views
 - permissions, 385
 - SYSUSERAUTH view
 - permissions, 386
 - SYSUSERLIST view
 - permissions, 386
 - SYSUSERPERM table
 - permissions, 385
 - SYSUSERPERMS view
 - permissions, 386
- ## T
- t option
 - backup [dbbackup] utility, 439, 441
 - histogram [dbhist] utility, 461, 462
 - initialization [dbinit] utility, 466, 471
 - Log Translation [dbtran] utility, 488
 - service creation [dbsvc] utility, 499, 501
 - transaction log [dblog] utility, 507, 510
 - unload [dbunload] utility, 514, 520
 - validation [dbvalid] utility, 527, 529
 - write file [dbwrite] utility, 531, 533
 - table names
 - international aspects, 261
 - table permissions
 - setting, 359

- table size
 - limit, 636
 - number of rows, 636
- TableBitMaps property
 - description, 630
- tables
 - group owners, 373
 - limitations, 636
 - owner, 353
 - permissions, 353, 354
 - qualified names, 373, 376
 - replicating, 413, 422
- tabular data streams, 106
- tasks
 - about, 14
 - backup, 329
 - events, 244
 - recovery, 329
 - schedules, 244
- TaskSwitch property
 - description, 619
- TCP/IP
 - about, 93
 - addresses, 112
 - connecting across firewalls, 93, 190, 191, 194, 198
 - encrypting client/server communications, 95
 - Host communication parameter, 194
 - Open Server, 109
 - performance, 93
 - port number, 109, 155, 198
 - protocol stacks, 192
 - server configuration, 153, 189
 - starting, 16
 - supported protocols, 92
 - troubleshooting, 101
 - Windows 95/98/Me, 93
 - Windows NT/2000/XP, 93
- TDS communication parameter
 - description, 198
- TDS communication protocol
 - about, 106
- TDS_EMPTY_STRING_IS_NULL option
 - description, 601
 - obtaining value, 619
 - setting, 541
- technical support
 - newsgroups, xvii
- telnet
 - testing networks, 34
- TEMP environment variable
 - CE, 214
 - description, 212
 - disk space, 34
- Tempdir property
 - description, 626
- TempFileName property
 - description, 630
- temporary files
 - location, 209, 212
 - location on CE, 214
 - security, 209
- temporary options
 - setting, 537
- Temporary Table Pages statistic, 617
- temporary tables
 - limitations, 636
- TempTablePages property
 - description, 619, 630
- threading behavior
 - NetWare, 14
 - UNIX, 14
 - Windows 95/98/Me, 14
- threads, 14
 - about, 14
 - controlling behavior, 15
 - execution, 141
 - multiple processors, 143
 - number of, 144
 - Windows NT/2000/XP, 15
- ti option
 - database server, 150
- TIME_FORMAT option
 - compatibility option, 544
 - description, 601
 - obtaining value, 619
 - setting, 541

-
- TIME_ZONE_ADJUSTMENT option
 - description, 602
 - setting, 541
 - Timeout [TO] communication parameter
 - description, 199
 - timeouts
 - troubleshooting, 103
 - TIMESTAMP
 - DEFAULT_TIMESTAMP_INCREMENT option, 564
 - TIMESTAMP columns, 554
 - TIMESTAMP data type
 - comparing, 564
 - TIMESTAMP_FORMAT option
 - compatibility option, 544
 - description, 602
 - obtaining value, 619
 - setting, 541
 - TimeZoneAdjustment property
 - description, 619, 626
 - tl option
 - data source [dbdsn] utility, 451, 455
 - database server, 151
 - tmf option
 - database server, 151
 - TMP environment variable
 - CE, 214
 - description, 212
 - TMPDIR environment variable
 - CE, 214
 - description, 212
 - tmt option
 - database server, 151
 - TotalBuffers property
 - description, 626
 - tq time option
 - database server, 152
 - Transaction Commits statistic, 616
 - transaction log
 - about, 305
 - allocating space, 222
 - backup [dbbackup] utility, 440
 - changing location, 345
 - delete after checkpoint, 144
 - erasing, 458
 - initialization [dbinit] utility, 470, 471
 - limiting size, 237
 - location, 8
 - Log Transfer Manager, 408
 - Log Translation [dbtran] utility, 488, 492
 - management, 422
 - media failure, 341
 - mirror, 306, 346, 347, 429
 - option, 13
 - placing, 317
 - primary keys, 319
 - size, 319
 - transaction log [dblog] utility, 508, 509
 - translation utilities, 490
 - uncommitted changes, 344
 - validating, 327
 - write file [dbwrite] utility, 533
 - transaction log mirror
 - initialization [dbinit] utility, 470
 - purpose, 318
 - starting, 346, 347
 - write file [dbwrite] utility, 533
 - transaction log utility
 - about, 507
 - auditing, 396
 - options, 509
 - syntax, 507
 - transaction modes
 - chained/unchained, 557
 - Transaction Rollbacks statistic, 616
 - transactions
 - closing cursors, 558
 - distributed, 151
 - TransactionsSpanLogs property
 - description, 630
 - Transact-SQL
 - ALLOW_NULLS_BY_DEFAULT option, 550
 - AUTOMATIC_TIMESTAMP option, 554
 - column NULLs compatibility, 594
 - compatibility options, 545
 - delete permissions, 551
 - NULL behavior, 553
 - QUOTED_IDENTIFIER option, 594
 - update permissions, 551

translate log file wizard
using, 344, 488

trigger condition
defined, 236

TriggerPages property
description, 630

triggers
and referential integrity actions, 597
and replication, 567, 568
disabling, 140
permissions, 353, 365
permissions for creating, 353

troubleshooting
backups, 337
common problems, 102
connections, 494, 498
database connections, 73
database server, 156
database server request-level logging, 157
encrypted database performance, 402
identifying client applications, 165
ODBC, 494
protocols, 101
server address, 114
server startup, 34, 35
timeouts, 103
Windows CE connections, 66
wiring problems, 102

TRUNCATE TABLE statement
autocommit behavior, 605

TRUNCATE_DATE_VALUES option
description, 604
setting, 541

TRUNCATE_TIMESTAMP_VALUES option
description, 604
obtaining value, 619
setting, 541

TRUNCATE_WITH_AUTO_COMMIT option
description, 605
obtaining value, 619
setting, 541

truncating
character strings, 600

TRUNCATION_LENGTH option
description, 605
Interactive SQL settings, 548

TSQL_HEX_CONSTANT option
compatibility option, 544, 545
description, 606
obtaining value, 619
Open Client, 116

TSQL_VARIABLE option
compatibility option, 544

TSQL_VARIABLES option
compatibility option, 545
description, 606
obtaining value, 619
Open Client, 116

U

–u option
database server, 152
histogram [dbhist] utility, 461, 462
information [dbinfo] utility, 463, 464
license [dblic] utility, 479, 480
Log Translation [dbtran] utility, 488, 493
unload [dbunload] utility, 514, 520

–ud option
database server, 153
Log Transfer Manger [dbltm] utility, 481, 483

unchained transaction mode, 557

UncommitOp property
description, 619

uncompress database wizard
using, 511

uncompressing
database files, 511
database files from Sybase Central, 511

uncompression utility
about, 511
options, 512
syntax, 511

Unconditional [UNC] connection parameter
about, 70
about connection parameters, 164
description, 187

- introduction to connection parameters, 39
- tips, 72
- using default connection parameters, 51
- Unicode character sets
 - about, 277
- UNIX
 - choosing a collation, 271
 - database servers, 6
 - default character set, 264
 - environment variables, 208
 - ODBC support, 63
 - threading behavior, 14
- unload database wizard
 - using, 514
- UNLOAD statement
 - security, 397
- UNLOAD TABLE statement
 - security, 397
- unload utility
 - about, 513
 - options, 516
 - syntax, 514
- unloading
 - data, 397
 - databases, 513
- unloading data
 - security, 397
- unloading databases
 - from Sybase Central, 514
- UnschReq property
 - description, 626
- UPDATE permissions, 359
- UPDATE statement, 422
 - truncation of strings, 600
- updates
 - ANSI behavior, 551, 552
 - SQL/92behavior, 552
 - Transact-SQL permissions, 551
- upgrade database wizard
 - using, 521
- upgrade utility
 - about, 521
 - options, 523
 - syntax, 522
- upgrading
 - databases, 521
 - Java in the database, 523
- upgrading databases
 - from Sybase Central, 521
- upper code page
 - about, 255
- usage information
 - displaying, 127
- user creation wizard
 - using, 357
- user estimates
 - overriding, 606
- user id
 - PUBLIC options, 538
- user IDs
 - default, 352
 - listing, 385
 - managing, 351
 - ODBC configuration, 58
 - security features, 388
 - security tip, 403
- user names
 - license [dblic] utility, 480
- user option
 - debugger [dbprdbg] utility, 456, 457
- USER_ESTIMATES option
 - description, 606
 - obtaining value, 619
 - setting, 541
- Userid [UID] connection parameter
 - about, 70
 - about connection parameters, 164
 - description, 188
 - introduction to connection parameters, 39
 - tips, 72
 - using default connection parameters, 51
- Userid property
 - description, 619

- users
 - connected users, 368
 - creating, 357
 - deleting, 366
 - deleting integrated logins, 86
 - granting integrated logins, 84
 - remote permissions, 365
 - removing from groups, 371
 - setting options, 356
 - user-supplied selectivity estimates
 - USER_ESTIMATES option, 606
 - ut option
 - database server, 153
 - UtilCmdsPermitted property
 - description, 619
 - utilities. *See* See database utilities
 - backup [dbbackup], 438
 - backup [dbbackup] syntax, 439
 - collation [dbcollat], 292, 293
 - collation [dbcollat] syntax, 443
 - compression [dbshrink], 448
 - compression [dbshrink] syntax, 448
 - console [dbconsole], 450
 - console [dbconsole] syntax, 450
 - data source [dbdsn], 451
 - data source [dbdsn] syntax, 451
 - debugger [dbprdbg], 456
 - debugger [dbprdbg] syntax, 456
 - DSEdit, 110
 - erase [dberase], 458
 - erase [dberase] syntax, 459
 - file hiding [dbfhide], 446
 - file hiding [dbfhide] syntax, 446
 - histogram [dbhist], 461
 - histogram [dbhist] syntax, 461
 - information [dbinfo], 463
 - information [dbinfo] syntax, 463
 - initialization [dbinit], 465
 - initialization [dbinit] syntax, 466
 - Interactive SQL [dbisql], 472
 - Interactive SQL [dbisql] syntax, 472
 - introduction, 437
 - language [dblang], 477
 - language [dblang] syntax, 477
 - license [dblic], 479
 - license [dblic] syntax, 479
 - Log Transfer Manager [dbltm], 481
 - Log Transfer Manager [dbltm] syntax, 481
 - log translation [dbtran], 488
 - log translation [dbtran] auditing, 394, 396
 - log translation [dbtran] syntax, 488
 - ping [dbping], 494
 - ping [dbping] syntax, 494
 - rebuild [rebuild], 497
 - rebuild [rebuild] syntax, 497
 - server location [dblocate], 498
 - server location [dblocate] syntax, 498
 - service creation [dbsvc], 499
 - service creation [dbsvc] syntax, 499
 - spawn [dbspawn], 503
 - spawn [dbspawn] syntax, 503
 - stop [dbstop], 505
 - stop [dbstop] permissions, 140
 - stop [dbstop] syntax, 505
 - transaction log [dblog], 507
 - transaction log [dblog] auditing, 396
 - transaction log [dblog] syntax, 507
 - uncompression [dbexpand], 511
 - uncompression [dbexpand] syntax, 511
 - unload [dbunload], 513
 - unload [dbunload] syntax, 514
 - upgrade [dbupgrad], 521
 - upgrade [dbupgrad] syntax, 522
 - validation [dbvalid], 526
 - validation [dbvalid] syntax, 527
 - write file [dbwrite], 530
 - write file [dbwrite] auditing, 396
 - write file [dbwrite] syntax, 531
 - utility commands
 - permissions, 143
 - utility database
 - about, 226
 - allowed SQL statements, 226
 - connecting to, 227
 - passwords, 228
 - security and passwords, 228
- ## V
- v option
 - data source [dbdsn] utility, 451, 454
 - database server, 153
 - Log Transfer Manager [dbltm] utility, 481, 483
 - unload [dbunload] utility, 514, 520
 - validate database wizard
 - using, 331, 526

- validating
 - backups, 315
 - databases, 315, 526
- validating databases, 327
 - from Sybase Central, 526
- validation utility
 - about, 526
 - options, 528
 - syntax, 527
- variable width character sets
 - about, 259
- VariableHashSize property
 - description, 630
- verbose mode
 - Log Transfer Manger [dbltm] utility, 483
 - unload [dbunload] utility, 520
- VERIFY_ALL_COLUMNS option
 - description, 607
 - replication option, 547
- VERIFY_THRESHOLD option
 - description, 608
 - replication option, 547
- VerifyServerName [VERIFY] communication
 - parameter
 - description, 199
- version
 - database server, 153
- version mismatch
 - file locations, 206
- ViewPages property
 - description, 630
- views
 - owner, 353
 - permissions, 353, 354, 361
 - security, 378
 - security features, 388
- Visual Basic
 - ODBC configuration for, 57

W

- w option
 - backup [dbbackup] utility, 439, 441
 - data source [dbdsn] utility, 451, 454
 - service creation [dbsvc] utility, 499, 502
- WAIT_FOR_COMMIT option
 - description, 608
 - obtaining value, 619
 - setting, 541
- white space characters
 - defined, 286
- Windows
 - choosing a collation, 271
 - default character set, 264
 - services, 23
- Windows .NET Server
 - setting cache size, 131
- Windows 2000
 - setting cache size, 131
- Windows 95/98/Me
 - threading behavior, 14
- Windows CE
 - character set, 273
 - connecting from the desktop, 65, 66
 - creating databases, 273
 - file locations, 205
 - installation, 205
 - ODBC, 63
 - setting TEMP file, 214
 - starting the database server, 9
 - subdirectories, 205
- Windows NT/2000/XP
 - and TCP/IP, 93
 - event log, 125
 - Performance Monitor, 610
 - threading behavior, 15
- Windows operating systems
 - database servers, 5
- Windows Performance Monitor, 610
- Windows services
 - registry settings, 213
- Windows XP
 - setting cache size, 131

wiring

- troubleshooting, 102

WITH GRANT OPTION clause, 363

wizards

- backup database, 339
- change log file, 345, 347
- change log file settings, 507
- compress database, 448
- create backup images, 333, 438
- create custom collation, 442
- create database, 465
- create write file, 530
- database object debugger, 456
- dbspace creation, 221
- erase database, 458
- event creation, 244
- group creation, 370
- integrated login creation, 85
- restore database, 343
- service creation, 24
- translate log file, 344, 488
- uncompress database, 511
- unload database, 514
- upgrade database, 521
- user creation, 357
- validate database, 331, 526

write file utility

- about, 530
- auditing, 396
- options, 532
- syntax, 531

write files

- about, 224
- backup [dbbackup] utility, 441
- creating from Sybase Central, 530
- erasing, 458
- managing, 530
- managing using dbwrite, 531
- write file [dbwrite] utility, 532

X

–x option

- backup [dbbackup] utility, 439, 441
- collation [dbcollat] utility, 443, 445
- data source [dbdsn] utility, 451, 455
- database server, 153

- Interactive SQL [dbisql] utility, 472, 475

- Log Translation [dbtran] utility, 488, 493

- stop [dbstop] utility, 505, 506

- transaction log [dblog] utility, 507, 510

–xi option

- unload [dbunload] utility, 514, 520

XML file format, 588

–xo option

- backup [dbbackup] utility, 439, 441

xp_cmdshell system procedure

- security features, 403

xp_sendmail system procedure

- security features, 403

xp_startmail system procedure

- security features, 403

xp_startsmtp system procedure

- security features, 403

xp_stopmail system procedure

- security features, 403

xp_stopsmtmp system procedure

- security features, 403

–xx option

- unload [dbunload] utility, 514, 520

Y

–y option

- backup [dbbackup] utility, 439, 441

- collation [dbcollat] utility, 443, 445

- compression [dbshrink] utility, 448, 449

- data source [dbdsn] utility, 451, 454

- database server, 155

- erase [dberase] utility, 459, 460

- Log Translation [dbtran] utility, 488, 493

- service creation [dbsvc] utility, 499, 502

- stop [dbstop] utility, 505, 506

- uncompression [dbexpand] utility, 511, 512

- unload [dbunload] utility, 514, 520

- write file [dbwrite] utility, 531, 533

year 2000

- NEAREST_CENTURY option, 584

Z**-z option**

collation [dbcollat] utility, 443, 445
data source [dbdsn] utility, 451, 455
database server, 35, 156
initialization [dbinit] utility, 466, 471
Log Translation [dbtran] utility, 493
ping [dbping] utility, 494, 496
transaction log [dblog] utility, 507, 510

-zl option

database server, 156

-zo option

database server, 156

-zr option

database server, 157

-zs option

database server, 157

