SYBASE®
An SAP® Company

Configuration Guide

Adaptive Server® Enterprise

15.7

[Windows]

DOCUMENT ID: DC38421-01-1570-01

LAST REVISED: September 2011

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Contents

CHAPTER 1

Introduction

Adaptive Server® Enterprise for Windows is a full-featured Adaptive Server that runs on Windows operating systems in the Windows environment.

The instructions in this book assume that Adaptive Server is installed and running. For information about installing and starting Adaptive Server, see the *Installation Guide* for your platform.

This chapter provides an overview of how to configure Adaptive Server and the steps you need to take to customize it for your use.

About Adaptive Server

Adaptive Server performs data management and transaction functions, independent of client applications and user interface functions.

Adaptive Server also:

- Manages multiple databases and multiple users
- Keeps track of the data's location on disks

- Maintains the mapping of logical data description to physical data storage
- Maintains data and procedure caches in memory

Adaptive Server uses these auxiliary programs to perform dedicated tasks:

- Backup Server manages database load, dump, backup, and restoration activities.
- XP Server stores the extended stored procedures (ESPs) that allow Adaptive Server to run operating-system level commands.

System-specific issues

Adaptive Server runs on a variety of hardware and operating system platforms. System-specific issues do not affect the basic functionality of Adaptive Server, but there are differences among platform implementations. These differences may include:

- Adaptive Server configuration
- Changes to the operating system that enable or enhance Adaptive Server performance
- Adaptive Server features that are available only on Windows
- The structure of entries in the *sql.ini* file
- Options for selecting database devices
- Operating system commands or utilities that simplify or automate routine system administration tasks
- Operating system utilities for monitoring Adaptive Server performance

System-specific issues are described in this document. For more information about system-specific issues, see the *Installation Guide* and release bulletin for your platform.

Definition of terms

The following terms are used in this book:

- *Server* – provides a service in client/server computing. Examples include Adaptive Server, Backup Server, and XP Server.
- *Client* – requests a service in client/server computing. Sybase Central™, PowerDesigner®, SQL Modeler™, and end-user applications are clients.

User roles

The Adaptive Server installation and setup process defines various user roles. Different user roles have different responsibilities and privileges. These user roles clarify the way Adaptive Server is to be integrated into your system:

- Operating System Administrator – the individual who maintains the operating system. This individual has administrator privileges.
- System Administrator – the individual in charge of Adaptive Server system administration, creating user accounts, assigning permissions on databases, and creating new databases. At installation time, the System Administrator's login name is "sa". The "sa" login is specific to Adaptive Server and is used to log in to Adaptive Server using the isql command.

Environment variables

It is crucial to the operation of Sybase® products that the system environment variables are set correctly. The installer will set the environment variables automatically at the system level.

Note As part of the installation, the installer sets up these environment variables in the system.

As part of the installation, the installer sets up these environment variables in the system:

- DSLISTEN – defines the name Adaptive Server uses to listen for client connections if no name is provided during the Adaptive Server start-up. If DSLISTEN is not set, and no name is given during the Adaptive Server start-up, the Adaptive Server name defaults to the server name given at installation.

- DSQUERY – defines the Adaptive Server name that client programs try to connect to if no Adaptive Server is specified with a command line option. If DSQUERY is not set, and you do not supply the Adaptive Server name with a command-line option, clients attempt to connect to the server name given at installation.
- SYBASE – defines the path of the Sybase installation directory. The installation program sets up the variable SYBASE to point to the release directory specified during installation.
- SYBASE_ASE – defines the subdirectory directory of the Adaptive Server components.
- SYBASE_OCS – defines the subdirectory to which the Open Client™ is set.
- SYBASE_SYSAM – points to the license-management software directory.
- SYBASE_TS_MODE – On Windows, Adaptive Server uses SYBASE_TS_MODE to determine if the shared memory should use a GLOBAL name space or a session-specific LOCAL name space. Sybase recommends that Adaptive Server uses a GLOBAL name space to which it can attach diagnostic tools for servers you start as a service, or when connecting to the server through terminal services.

  The default mode in versions of Adaptive Server earlier than 15.7 was LOCAL, which imposed diagnostic limitations. In Adaptive Server 15.7 and later, the default is GLOBAL.

  Setting SYBASE_TS_MODE to LOCAL starts Adaptive Server in a pre-15.7 default mode. There is no advantage in using LOCAL name space and Sybase does not recommend doing this because it resticts shared memory access for diagnostic tools.
- PATH – specifies which directory paths to search for executables and dynamic link libraries (DLLs). The Sybase executables are in the *%SYBASE%|product_name\bin* directory. The installation program appends these paths to the current PATH environment variable.
- TEMP – defines the location used by the installation program to write files temporarily during the installation process. The installation process frees the disk space after installation is completed.
- INCLUDE – specifies which directory to set to or append for Open Client.
- LIB – is appended with *lib* directory for Open Client.

Adaptive Server devices and system databases

Devices are files or portions of a disk that are used to store databases and database objects. You can initialize devices using raw disk partitions or operating system files.

Adaptive Server requires the following devices:

- master – to store system databases.
- sybsystemdb – to store information about transaction in process.
- sysprocsdev – to store system procedures.

The master, sybsystemdb, and sysprocsdev devices are created when you create a new Adaptive Server.

The master device

The master device contains the following databases:

- master – controls the operation of Adaptive Server as a whole and stores information about all users, user databases, devices, objects, and system table entries. The master database is contained entirely on the master device and cannot be expanded onto any other device.
- model – provides a template for new user databases. The model database contains required system tables, which are copied into a new user database with the create database command.
- tempdb – the work area for Adaptive Server. Each time Adaptive Server is started the tempdb database is cleared and rebuilt from the model database.
- The sample databases are stored on the master device at installation, but should be moved to a user-defined device after installation. For more information, see “The sample databases” on page 7.

Note For recovery purposes, Sybase recommends that you do not create other system or user databases or user objects on the master device.

The sybsystemdb device

The sybsystemdb device stores the sybsystemdb database, which stores information about transactions in progress, and which is also used during recovery.

For instructions about creating the sybsystemdb device and database for Data Transfer Management (two-phase commit), see "Upgrading Sybase Servers" in the Adaptive Server *Installation Guide* for your platform.

The sysprocsdev device

The sybprocsdev devices stores the sybsystemprocs database, which contains most of the Sybase-supplied system procedures. System procedures are a collection of SQL statements and flow-of-control statements that perform system tasks, for example, sp_configure.

The system procedures that are needed during recovery situations are stored in the master database.

Note sysprocsdev is the default system name for this device. However, it is frequently referred to as the sybsystemprocs device, since it stores the sybsystemprocs database.

Optional devices and databases

The devices and databases described in the following sections are optional.

Pluggable component interface (PCI) database

The pluggable component interface (PCI) allows you to add libraries that provide different functionalities to the Adaptive Server. Java support (pluggable component adaptor/Java virtual machine) is included as a pluggable component with Adaptive Server 15.0.3.

The sybpcidb database stores necessary configuration information for the PCI and the pluggable component adaptor/Java virtual machine (PCA/JVM) plug-in.

To enable PCI in Adaptive Server by using sybatch, add PCI/Java-related properties to the resource files used by these utilities. Enter these values:

```
sqlsrv.do_configure_pci: yes
sqlsrv.sybpcidb_device_physical_name:\device_path
```

```
sqlsrv.sybpcidb_device_size: USE_DEFAULT
sqlsrv.sybpcidb_database_size: USE_DEFAULT
```

The sample databases

The sample databases are:

- pubs2 and pubs3 are provided as learning tools for Adaptive Server. pubs2 is used for most of the examples in the Adaptive Server documentation; other examples use the pubs3 database. Both are available in U.S. English versions of Adaptive Server.
- interpubs database contains French and German data.
- jpubs contains Japanese data.

For information about installing the sample databases, see Chapter 5, "Post-Installation Tasks" in the Adaptive Server *Installation Guide* for your platform.

For information about the contents of the sample databases, see the *Transact-SQL Users Guide*.

The sybsecurity device and database

The sybsecurity device is created as part of the auditing installation process. The sybsecurity device stores the sybsecurity database and the auditing system procedures with which you can configure auditing for your system.

The auditing system records system security information in an Adaptive Server audit trail. You can use this audit trail to monitor the use of Adaptive Server or system resources.

For instructions on configuring Adaptive Server for auditing, see Chapter 11, "Adding Optional Functionality to Adaptive Server." For information about installing and using the auditing system, see Chapter 12, "Auditing," in the *System Administration Guide: Volume 1*.

dbccdb database

The database consistency checker (dbcc) provides commands for checking the logical and physical consistency of a database. The dbccdb database stores the results of dbcc when dbcc checkstorage or dbcc checkverify is used.

dbcc checkstorage records configuration information for the **target database**, operation activity, and the results of the operation in the dbccdb database. Stored in the database are dbcc stored procedures for creating and maintaining dbccdb and for generating reports on the results of dbcc checkstorage operations.

For information on installing and using dbccdb, see "Checking Database Consistency," in the *System Administration Guide: Volume 2*.

Client/server communication (the interfaces file)

Adaptive Server communicates with other Adaptive Servers, Open Server applications (such as Backup Server), and client software on your network. Clients can talk to one or more servers, and servers can communicate with other servers by remote procedure calls.

For Sybase products to interact with one another, each product needs to know where the others reside on the network. Names and addresses of every known server are listed in a directory services file. This information can be stored in a directory services file two different ways:

- In an interfaces file called *sql.ini* on Windows platforms, located in the *%SYBASE%\ini* installation directory, or
- In an LDAP server

After your Adaptive Server or client software is installed, it can connect with any server on the network that is listed in the interfaces file

When you are using a client program, and you want to connect with a particular server, the client program looks up the server name in the interfaces file and connects to that server, as shown in Figure 1-1. You can supply the name of the server by using the DSQUERY environment variable.

On TCP/IP networks, the port number gives clients a way to identify the Adaptive Server to which they want to connect. It also tells Adaptive Server where to listen for incoming connection attempts from clients. Adaptive Server uses a single port for these two services (referred to as *listener servicequery service*).

On SPX networks, the socket number gives clients and servers a way to identify each other.

Figure 1-1: Communicating with a server using the sql.ini file

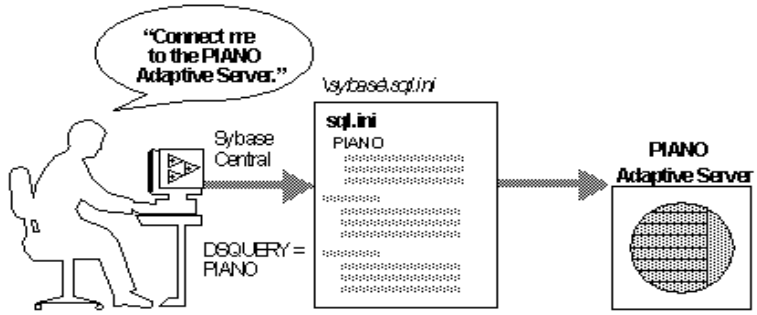

During installation, the installer adds entries to the *sql.ini* file for the new Adaptive Server, Backup Server, or XP Server.

Note You can use Windows File Replication to replicate *sql.ini* in the multiple locations. See the Microsoft documentation for information. You can also use Windows Registry to keep the interfaces file information.

Changing Adaptive Server configuration

For information about configuring Adaptive Server to your system's needs, see Chapter 3, "Default Adaptive Server Configuration."

For information about configuring languages, character sets, and sort orders, see Chapter 6, "Customizing Localization for Adaptive Server."

For information about configuring Adaptive Server to use high availability features, see *Using Sybase Failover in a High Availability Environment.*

For information about configuring Adaptive Server for distributed transaction management (two-phase commit), see the *Distributed Transaction Management User's Guide*.

Languages other than U.S. English

Many of the configuration tasks described in this manual require the use of the Server Config utility.

If you are running Server Config in a language other than U.S. English, make sure that any input you provide uses a character set that is supported by the us_english character set.

Note The us_english character set does not support accent marks, such as tildes (~) and umlauts (ü). This prevents Server Config from supporting the character sets that use these characters.

For more information about languages, character sets, and sort orders, see the Adaptive Server *Installation Guide* for your platform.

Adaptive Server specifications

Database specifications		
Databases per Adaptive Server	A maximum of 32,767 databases per server	
Maximum database size	• 2K page server – 4TB • 4K page server – 8TB • 8K page server – 16TB • 16K page server – 32TB	
Minimum allowable sybsystemprocs database	136MB	Required for an upgrade
Maximum size of a database device (disk partition)	2^{42} (4TB)	If the operating system supports file sizes up to 4TB, then Adaptive Server supports file system devices up to 4TB
Maximum number of database devices per server	2^{31}	
Maximum number of devices or device pieces per database	Unlimited	Limited by available memory
Maximum number of segments per database	31	

Maximum number of login IDs per server	2147516416
Maximum number of users per database	2146484223
Maximum number of groups per database	1032193

Table specifications

User objects per database	$2^{31} – 255$	
Indexes per table	250 (one clustered index)	
Rows per table	Limited by available storage	Maximum 2^{32}
Columns per composite index	31	
Creation of clustered index	1.2*(x + y) x = total data space in table, y = sum of space of all nonclustered indexes on table, and 20 percent overhead for logging	For sorted data, approximately 20 percent of the table size needed
Maximum size of object name	255	

Query specifications

Maximum number of tables participating in a query, for a query without a union, or for each side of a union in a query	64	Maximum of 50 user tables, including result tables, tables referenced by views (the view itself is not counted) correlations and self-joins; maximum of 46 worktables
Maximum number of tables in a “union” query	256	Includes up to 50 user tables and 14 work tables on every side of the union, for a total of up to 256 tables across all sides of the union.

Maximum number of databases participating in one transaction	Unlimited	Includes database where transaction began, all databases changed during transaction, and tempdb, if it is used for results or worktables
Practical number of databases participating in one query	16	Includes each occurrence of each database queried and tempdb, if it is used for results or worktables
Maximum number of tables with referential integrity constraints for a query	192	

Procedure specifications

Number of buffers and procedure buffers	Configurable	Limited by amount of RAM and maximum size of shared memory segment
Minimum memory required per stored procedure	2K	
Maximum number of parameters per stored procedure	2048	

Adaptive Server extended-limit capabilities vary by type of table and the database logical page size. Table 1-1 lists the column and row limits for allpages-locked (APL) tables.

Table 1-1: Allpages-locked (APL) tables

Maximum APL table limits	Number of columns	Column size 2K page	Column size 4K page	Column size 8K page	Column size 16K page
Fixed-length column	1024	1960 bytes	4008 bytes	8104 bytes	16296 bytes
Variable-length column	254	1948 bytes	3988 bytes	8068 bytes	16228 bytes

Table 1-2 lists the column and row limits for data-only-locked (DOL) tables.

Table 1-2: Data row and data page tables.

Maximum DOL table limits	Number of columns	Column size 2K page	Column size 4K page	Column size 8K page	Column size 16K page
Fixed-length column	1024	1958 bytes	4006 bytes	8102 bytes	16294 bytes
Variable-length column	1024	1954 bytes	4002 bytes	8098 bytes	16290 bytes

Database space requirements depend upon the logical page size of the server.If your model database is larger than the minim size listed below, then the minimum size of the database is equal to the model database. Table 1-3 lists the minimum size for each database.

Table 1-3: Database requirements for varying page sizes

Databases	2K page	4K page	8K page	16K page
Default database size	3MB	6MB	12MB	24MB
master database	13MB	26MB	52MB	104MB
model database	3MB	6MB	12MB	24MB
tempdb database	4MB	6MB	12MB	24MB
sybpcidb database	24MB	48MB	96MB	192MB

Larger logical page sizes can contain more data. Table 1-4 lists the maximum data for each logical page size.

Table 1-4: Data limits for tables according to page size

Tables	2K page	4K page	8K page	16K page
Bytes per index key	600	1250	2600	5300
User-visible row length DOL table	1958	4006	8102	16294
User-visible row length APL table	1960	4008	8104	16296

CHAPTER 2

Starting and Stopping Servers

This chapter describes how to start and stop Adaptive Server, Backup Server, and XP Server.

Overview

The methods described in this chapter are used to start Adaptive Server and Backup Server after a shutdown for database maintenance, because of an operating system crash, or for other reasons.

XP Server is not started by the installation process. XP Server is started only when any XP command is issued through isql.

You can use Sybase Central to start and stop servers manually or automatically. The ASE plug-in can also start Adaptive Servers monitored by Unified Agent if the agent is properly configured.

Requirements for starting servers

To start a server, your user account must have:

- Windows administrator privileges
- Access to the Adaptive Server distribution files
- Access to a *sql.ini* file entry for the server
- The system environment variables discussed in Chapter 1, “Introduction.”

- Access to SySAM licenses. For more information, see the *Sybase Software Asset Management Users Guide*.

The installation program creates the *sql.ini* file and system environment variables when you install servers on your computer.

Starting servers

Adaptive Server and Backup Server store their default start-up parameters in the Windows Registry file. This allows you to start and manage servers as Windows services, and allows servers to start automatically when you boot your computer.

Server start-up parameters

The default start-up parameters are stored under the Registry key \\HKEY_LOCAL_MACHINE\SOFTWARE\SYBASE\Server\ *server_name*\Parameters, where *server_name* is the name of the server you installed.

Backup Server server names are appended with "_BS".

Note You can install multiple servers, each with its own Registry key.

Start-up parameters are listed under Registry values named *Argn*, where *n* is a number from 0 to 8. The number of the argument indicates the order in which the server reads the parameter.

Table 2-1 lists the default start-up parameters for Adaptive Server.

Table 2-1: Default Adaptive Server start-up parameters

Parameter	Switch	Description
Arg0	-d *%SYBASE%\data\master.dat*	Location of the master device file
Arg1	-s *server_name*	Name of the Adaptive Server
Arg2	-e*%SYBASE%\%SYBASE_ASE%\ install\errorlog*	Location and name of the error log file
Arg3	-i *%SYBASE%\ini*	Directory containing the *sql.ini* file
Arg4	-M *%SYBASE%\%SYBASE_ASE%*	Directory that stores shared memory files

Parameter	Switch	Description
Arg5	-N *%SYBASE%\$SYBASE_ASE%\sysam\<srv_name>.properties*	Location and name of license cache file.

Changing start-up parameters

You cannot change any of these default start-up parameters unless you directly edit the Windows Registry values. However, you can use Server Config to specify additional start-up parameters.

Additional start-up parameters include any valid server command-line options listed for the sqlsvr and bcksrvr descriptions in the Adaptive Server *Utility Guide* for your platform.

❖ **Specifying additional start-up parameters**

1 Log in to Windows using an account with Windows administrator privileges.

2 Start Server Config.

3 Select the Adaptive Server or Backup Server icon to indicate which type of server you want to configure.

4 Select Configure Adaptive Server or Configure Backup Server to display a list of available servers on your system.

5 Select the name of the server to configure, and choose Continue.

6 If you are configuring Adaptive Server, enter the login name and password of a user with System Administrator privileges, and choose Continue.

7 If Adaptive Server is not running, Server Config asks you to start it now; choose Yes.

8 Select the Command Line button. Server Config displays the Command Line Parameters dialog box:

9 Edit the text in the Command Line Parameters box to include the additional start-up parameters and values you require.

Do not specify the default command line parameters listed in Table 2-1 on page 16. For details on available command-line parameters, see sqlsvr and bcksrvr in the Adaptive Server *Utility Guide* for your platform.

10 Choose OK.

11 Choose Save in the server's configuration dialog box.

12 Exit Server Config.

Starting and stopping servers using Unified Agent

You can shut down Adaptive Servers running either locally or remotely if you have the proper permission to do so. With ASE plug-in, connect to the Adaptive Server you want to shut down, and select File | Shut Down. If the Adaptive Server is monitored by Unified Agent, you do not have to connect first. Simply select the Adaptive Server and then select File | Shut Down.

Starting and stopping servers using the Control Panel

You can start, stop, and pause a server both automatically and manually from the Services applet in the Control Panel.

Note If you are running Adaptive Server and the Windows Process Viewer, and Adaptive Server is listed in the Process Viewer, you may not be able to restart Adaptive Server after you shut it down. This is because the Process Viewer holds some Registry resources, even after the viewed process is closed. Shut down the Process Viewer before you restart Adaptive Server.

Starting servers as an automatic service

This section describes how to configure your operating system for automatic restart of Adaptive Server and Backup Server.

In production systems, Adaptive Server and Backup Server should start automatically when you restart your computer. To do this, use the Control Panel to set up the server as an automatic service.

Note Do not place Adaptive Server devices on network drives. If Adaptive Server uses a device on a network drive, you cannot start the server as an automatic Windows service.

❖ **Setting up Adaptive Server as an automatic service**

1 Go to Windows Services at Start | Settings | Control Panel | Administrative Tools | Services.

2 Scroll through the list of available services until you find the listings for your Sybase servers.

Server names use the format:

"Sybase *type*Server _ *servername_suffix*"

where *servername* is the name of the Adaptive Server and *type* and *_suffix* represent the server type:

"SQL" for Adaptive Server

"BCK" and "_BS" for Backup Server

"XP" and "_XP" for XP Server

3 Double click on the Adaptive Server service entry, or right click on the Adaptive Server service entry and select Properties.

4 In "Startup type," select "Automatic."

5 Click the Close button to close the Services window.

The selected server now starts automatically each time you restart the computer. You can verify the status of the server by examining the status column in the Services applet.

See your Windows documentation or online help for more information on setting up automatic services.

Starting, stopping, and pausing servers manually

You can use the Control Panel to stop, start, and pause Adaptive Server manually.

1 Log in to Windows using an account with Windows administrator privileges.

2 Choose Start | Settings | Control Panel | Administrative Tools | Services. The Services window displays.

3 Scroll through the list of available services until you find the listings for your Sybase servers.

Server names use the format:

"Sybase *type*Server_*servername_suffix*"

where *servername* is the name of the Adaptive Server and *type* and *_suffix* represent the server type:

- "SQL" for Adaptive Server
- "BCK" and "_BS" for Backup Server
- "XP" and "_XP" for XP Server

4 Select the service name, then click Start, Stop, or Pause to confirm the choice.

5 Click the Close button to close the Services window.

You can verify the status of the server either by using Sybase Central or by examining the status column in the Services applet.

Stopping servers

Only the System Administrator has permission to issue a shutdown command. Using a shutdown command minimizes the amount of work that automatic recovery needs to do when the servers are restarted.

The preferred method of stopping Adaptive Server or Backup Server is to use the Transact-SQL shutdown command.

Stopping Adaptive Server

To shut down Adaptive Server:

1 Use isql to log in to an Adaptive Server account with System Administrator privileges:

```
isql -Usa -Ppassword -Sserver_name
```

2 Enter the following command to shut down the server:

```
1> shutdown
2> go
```

The default for the shutdown command is with wait. The with wait option allows Adaptive Server to finish executing SQL statements or procedures, perform a checkpoint in each database, disable new logins, and perform other shutdown tasks.

Issuing the shutdown command prints a message like this to the *stderr* file:

```
Server SHUTDOWN by request.The SQL Server is terminating
this process.
CT-LIBRARY error:
```

This is normal behavior.

If the message indicates that Adaptive Server is waiting for processes to complete, and you need to stop Adaptive Server immediately, you can use the shutdown with nowait command. shutdown with nowait does not wait for currently executing statements to finish, nor does it perform checkpoints in every database. Use the shutdown with nowait command only when necessary.

Stopping Backup Server

To shut down a Backup Server:

1 Use isql to log in to a server with System Administrator privileges:

```
isql -Usa -Ppassword -Sserver_name
```

2 Enter the following command to shut down the specified Backup Server:

```
1> shutdown SYB_BACKUP
2> go
```

After you shut down a Backup Server, you must wait at least 30 seconds before restarting it; otherwise, you see a message from the operating system that another process is using the disk.

Issuing the shutdown command prints a message similar to the following to the *stderr* file:

```
Backup Server: 3.48.1.1: The Backup Server will go down
immediately.
 Terminating sessions.
```

This is normal behavior. If a message indicates that Adaptive Server or Backup Server is waiting for processes to complete, and you need to stop Adaptive Server or Backup Server immediately, you can use the shutdown with nowait command. shutdown with nowait does not wait for currently executing statements to finish and does not perform checkpoints in every database.

Using shutdown with nowait for Backup Server can cause inconsistent or incomplete dumps and loads. Use this command only when necessary.

For more information on the shutdown command, see the *Reference Manual: Commands*.

Monitoring servers

There are two methods for checking a server's status: Unified Agent or through the Control Panel.

Unified Agent

You can monitor the Adaptive Server status either locally or remotely using Unified Agent, if Unified Agent is monitoring the Adaptive Server.

For more information about using Unified Agent to monitor Adaptive Server, see the *Unified Agent / Agent Management Console User's Guide*.

The Control Panel

You can use the Services option in the Control Panel to check a local server's status. The Services option is available under Start | Settings | Control Panel | Administrative Tools | Services. Check the Status column. If the server is:

- Running, the Status value is Started.
- Not running, the Status value is blank.

CHAPTER 3

Default Adaptive Server Configuration

When you install or upgrade Adaptive Server, it includes some default parameter settings and a few of its auxiliary programs.

After installing and testing this "default" Adaptive Server, configure it to your system's needs and install other optional features.

After installation, Adaptive Server default settings are as listed in Table 3-1. You may need to configure these settings to suit your computer and database needs.

Table 3-1: Defaults for Adaptive Server parameter settings

Item	Default value
Name	*AdaptiveServername*
Network support	TCP/IP
Socket number	5000
Named pipes	*\pipe\sybase\server*
Command line options	None
Error log path	*%SYBASE%\%SYBASE_ASE%\install/error log*
Event logging	Not configured
International Support (Localization):	
• Language	us_english
• Character set	cp850
• Sort order	Binary ordering
Login security mode	Standard

Table 3-2 lists the default settings for the Backup Server and XP Server.

Table 3-2: Defaults for the Backup and XP servers

Server	Item	Default value
Backup Server	Name	*AdaptiveServername*_BS
	Network support	Named Pipes, Windows Sockets (TCP/IP)
	Pipe name	*\pipe\sybase\backup*
	Socket number	5001
	Error log path	*%SYBASE%\%SYBASE_ASE\install\backup.log*
XP Server	Name	*AdaptiveServername*_XP
	Network support	Named Pipes, Windows Sockets (TCP/IP)
	Pipe name	*\pipe\sybase\xp*
	Socket number	5002
	Error log path	N/A

Starting Server Config for Adaptive Server

To change configuration settings for Adaptive Server, use the Server Config utility. You can run this program in one of two ways:

- By selecting the Server Config from within Windows. To run this utility from the Windows command prompt, run syconfig.exe.
- By running sp_configure from within isql. Use sp_configure to quickly and easily change single parameters and values. For more information, see sp_configure in the *Reference Manual: Procedures*.

This manual walks you through Adaptive Server configuration through the Server Config utility.

❖ **Starting Server Config**

1 Select Start Programs.

2 Select Start | Programs | Sybase | Adaptive Server Enterprise | Server Config.

3 Choose Server Config

4 When you have completed the necessary configuration changes, click Exit to quit Server Config.

For more information on how to configure:

- Adaptive Server, see “Configuring Adaptive Server” on page 25.

- For more information on how to configure Backup Server, see "Configuring Backup Server" on page 27.

Note The Adaptive Server 15.0.3 installer allows you to tune basic configuration settings during installation, instead of as a postinstallation task. See the *Installation Guide* for more details.

Configuring Adaptive Server

To change the Adaptive Server configuration, including its auxiliary programs and options:

1 Start Server Config.

2 Click the Adaptive Server icon, and click Configure Adaptive Server from the Configure Sybase Servers dialog box.

3 Select the name of the server to configure, and click Continue.

The Enter System Administrator Password dialog box appears.

4 Type the login name and password of an Adaptive Server user with System Administrator privileges, and click Continue.

5 Click Yes if the Adaptive Server is not running, and Server Config asks you if you want to start it.

The Configuring Adaptive Server Enterprise dialog box appears.

6 Select the option to be configured from the Change Options set of buttons:

- Command Line – see "Setting Adaptive Server parameters" on page 26.
- Default Backup Server – see "Changing the default Backup Server" on page 26.
- Default XP Server – see "Changing the default XP Server" on page 27.
- Two Phase Commit – see the Adaptive Server *Installation Guide* for your platform.
- Error Log Path – see "Setting error log paths" on page 92.

- Event Logging – see“Enabling and disabling Windows event logging” on page 94.
- Language – see the Adaptive Server *Installation Guide* for your platform.
- Login Security – see “Configuring login security” on page 133.

Setting Adaptive Server parameters

When you start Adaptive Server, you can configure the server to use certain configuration parameters that are not accessible through isql.

To set these configuration parameters:

1 Click Command Line from the Change Options box on the Configuring Adaptive Server Enterprise dialog box.

The Command Line Parameters dialog box appears.

2 Type in parameters and values that you want to set for Adaptive Server.

Type these parameters as you would at the command line. However, omit the command itself and any parameters that might vary.

3 Click OK to return to the Configure Adaptive Server Enterprise dialog box.

4 When you have completed the necessary configuration changes, click Exit to quit Server Config.

Changing the default Backup Server

During backup or recovery, the dump or load command uses the Backup Server named in the configuration for the selected Adaptive Server. You can name a different default Backup Server through the Adaptive Server configuration.

To name a different Backup Server to use as the default:

1 Click Default Backup Server from the Change Options buttons.

The Set Default Backup Server Name dialog box appears.

2 Type the name of the Backup Server as the default, and click OK.

For information about naming and configuring Backup Server, see “Configuring Backup Server” on page 27.

3 Click Save to return to the Configuring Adaptive Server Enterprise dialog box.

4 When you have completed the necessary configuration changes, click Exit to quit Server Config.

Changing the default XP Server

XP Server provides the extended stored procedures available through Adaptive Server.

When you install Adaptive Server, the program defines XP Server using the Adaptive Server name as a basis for the filename. For example, XP Server for an Adaptive Server named PIANO is named PIANO_XP.

You can change the configuration for the default XP Server for a particular Adaptive Server. See "Sybmail and extended stored procedures" on page 143.

Configuring Backup Server

Backup Server performs all Adaptive Server backup and recovery operations (dump and load).

When you install Adaptive Server, the program defines Backup Server using the Adaptive Server name as a basis for the file name. For example, Backup Server for an Adaptive Server named PIANO is named PIANO_BS.

To change the configuration for a Backup Server:

1 Start Server Config.

2 Click the Backup Server icon, and click Configure Backup Server from the Configure Sybase Servers dialog box.

3 Select the name of the server to configure from the Existing Servers dialog box, and click Continue.

The Configure Backup Server dialog box appears.

4 Change the path indicated in the Error Log Path area, if necessary.

For more information about the error log, see "Logging errors and events" on page 89.

5 Change the language indicated in the Language area that Backup Server will use for its messages, if necessary.

   For more information about languages, see the *Installation Guide*.

6 Change the server's character set in the Character Set area, if necessary.

   For more information about character sets, see the *Installation Guide*.

7 Click Save to return to the Configure Sybase Servers dialog box.

8 When you have completed the necessary configuration changes, click Exit to quit Server Config.

Configuring Job Scheduler and Self Management

Job Scheduler defines and schedules database administration and maintenance tasks. Self Management is Adaptive Server's ability to monitor it's state and adjust that state as necessary. Adaptive Server's ability to manage itself can be extended by creating and scheduling job that perform maintenance and tuning tasks.

You can only configure Job Scheduler and Self Management in resource file mode. To configure new Job Scheduler and Self Management, edit the sample resource file *%SYBASE%\%SYBASE_ASE%\sample\server\sybatch_js.res* and execute:

```
sybatch.exe -r sybatch_js.res
```

See the *Job Scheduler Users Guide* for more information.

CHAPTER 4

Network Communications Using *sql.ini*

Adaptive Server can communicate with other Adaptive Servers, Open Server applications, and client software across a network. Clients can communicate with one or more servers, and servers can communicate with other servers via remote procedure calls.

This chapter provides information about the connection process, the kinds of connections, and how to configure Adaptive Server to use *sql.ini* file connections.

For instructions on using Server Config to change the values that it can access, see “Configuring Adaptive Server” on page 25. For information on LDAP, see Chapter 5, “Lightweight Directory Access Protocol in Adaptive Server.”

Adaptive Server on Windows supports network connections using the Named Pipes, Sockets (TCP/IP), and IPX/SPX protocols. The default Adaptive Server uses TCP/IP and Named Pipes, since Named Pipes is always installed with Windows.

Two files control how clients find servers and drivers:

- The *sql.ini* file lists the server names, their network addresses, and the Net-Library driver to use to establish a connection.
- The library file, *libtcl.cfg*, lists the installed Net-Library drivers that are available to support each protocol (connection).

These files, which reside on both server and client machines, enable each Sybase product to find the other Sybase servers that are on the network. The installation program automatically creates, verifies, and appends these configuration files when you install Adaptive Server.

How clients connect to Adaptive Server

Client software performs the following steps to connect to Adaptive Server:

1 Determines the name of the Adaptive Server by finding the value of the DSQUERY environment variable.

2 Looks in the *sql.ini* file for an entry whose name matches the name of the server. If it cannot find a matching entry, the connection fails.

3 Looks in the *libtcl.cfg* file for an entry that matches the Net-Library driver name associated with the server entry in the *sql.ini*file. If the application cannot find such an entry, the connection fails.

4 Loads the specified Net-Library driver.

5 Uses the network connection information provided by the *sql.ini* file to connect to the server.

Figure 4-1 summarizes the client connection process.

Figure 4-1: Connecting to Adaptive Server

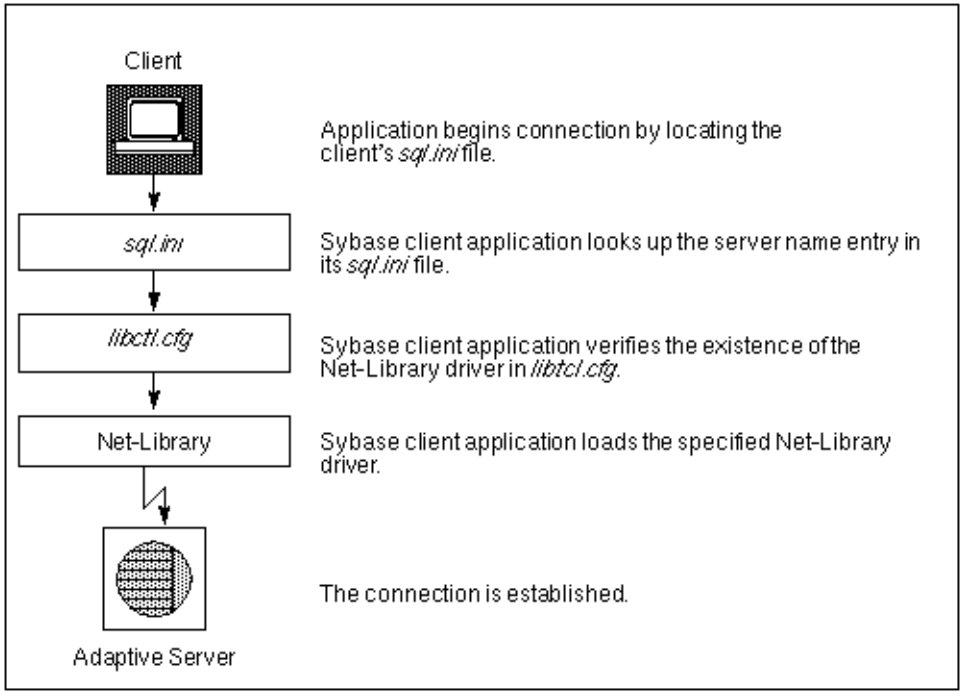

How Adaptive Server listens for client connections

Adaptive Server uses the *sql.ini* file to determine the address at which it should listen for clients. When you start, the Adaptive Server:

1 Determines the name of the Adaptive Server by finding the value of the DSLISTEN environment variable.

2 Looks in the *sql.ini* file for an entry that matches the specified server name.

3 Looks in the *libtcl.cfg* file for an entry that matches the Net-Library driver name associated with the server entry in the *sql.ini*file.

4 Loads the specified Net-Library driver.

5 Uses the information from the MASTER entry in the *sql.ini* file to determine the address at which it should listen for client connection requests.

How a client accesses Adaptive Server

The installation program provides a default *sql.ini* file in Adaptive Server. The file has MASTER and QUERY entries that use both the Named Pipes and Sockets (TCP/IP) drivers for all servers that were installed.

Enabling client access to a server

To enable a client to access a server on the network, create a *sql.in*i file on the client. In that file, include entries for all servers the client needs to access.

To create a new *sql.ini* file, see "Changing the server entries in sql.ini" on page 32.

Changing the server entries in *sql.ini*

To edit an existing *sql.ini* file on the server machine, or to create a new file on the client machine, use the Directory Services Editor utility, dsedit.

For more information about the components of a *sql.ini* file, see "Components in the sql.ini file" on page 33.

For more information about using dsedit, see the *Utilities Guide*.

For general information about the *sql.ini* file, see the *Open Client/Server Configuration Guide for Desktop Platforms*.

To start dsedit, select it either from the Sybase program group or from the Utilities group in Sybase Central.

To add an Adaptive Server to the *sql.ini* file:

1 Select Start | Programs | Sybase | Connectivity | Open Client Directory Service Editor.

The Select Directory Service dialog box appears.

2 Select a driver from the DS Name list, and click OK.

The DSEDIT - Interfaces Driver dialog box appears.

3 Select Server Object menu, and select Add.

The Input Server Name dialog box appears.

4 Type the name of the server to add, and click OK.

For information about valid server names, see "Server name" on page 34.

5 Select the new server name, which you have just added, from the Server list.

Steps 6 – 10 describe how to enter the server's address:

6 Select Server Address from Attributes box on the Interfaces Driver window.

7 Select the Server Object menu and select Modify Attribute.

The Network Address Attribute dialog box appears.

8 Click Add.

The Input Network Address For Protocol dialog box appears.

9 Choose the appropriate protocol from the drop-down list, enter the network address in the Network Address text box, and click OK.

For information about protocols, see "Network driver" on page 35.

For information about the formats of network addresses required by the different protocols, see "Server address" on page 35.

The Network Address Attribute dialog box reappears.

10 Click OK.

The dsedit utility creates MASTER and QUERY entries for the server. In the *sql.ini* file, the client ignores the MASTER entry.

11 Exit dsedit.

Components in the *sql.ini* file

This section provides useful background information for editing an *sql.ini* file. The *sql.ini* file will look similar to:

```
[PIANO_XP]
master=NLWNSCK,PIANO,5002
query=NLWNSCK,PIANO,5002


[PIANO]
master=NLWNSCK,PIANO,5000
query=NLWNSCK,PIANO,5000


[PIANO_BS]
master=NLWNSCK,PIANO,5001
query=NLWNSCK,PIANO,5001


[PIANO_JSAGENT]
master=NLWNSCK,PIANO,4900
query=NLWNSCK,PIANO,4900


[ws]
master=NLWNSCK,PIANO,8183
```

Server name

The server name is the name of the Adaptive Server to which clients will connect. Use the following rules to create an acceptable server name:

- Server names can be no more than 11 characters long. However, if you installed Adaptive Server on a FAT (file allocation table) partition, limit the server name to 8 characters.
- The initial character of a server name must be a letter (a–z, A–Z). The characters that follow can be letters, numbers, the underscore character (_), the pound sign (#), the at sign (@), or the dollar sign ($).
- The name cannot contain a period (.), a slash (/), a backslash (\), an accented letter, a character from a Japanese character set, or any other character that is invalid for Windows file names.
- Adaptive Server names are not case-sensitive. For example, “PRODUCTION,” “Production,” and “production” are interpreted as the same server name.

Network driver

The network driver specifies the name of the Net-Library driver to use for the connection. The driver name must correspond to a valid entry in the library (*libtcl.cfg*) file, which is located in the *ini* subdirectory of the Sybase installation directory.

The following example shows three driver entries in a *libtcl.cfg* file:

```
NLMSNMP=NLMSNMP Named Pipes Driver
NLWNSCK=NLWNSCK WinSock TCP/IP Driver
NLNWLINK=NLNWLINK NWLink SPX/IPX Driver
```

Note As drivers are added or removed, you can edit the *libtcl.cfg* file with a text editor or with the ocscfg.exe utility, located in the *bin* subdirectory of the Sybase installation directory.

Service type

The service type defines the Adaptive Server's service. The two service types are MASTER and QUERY:

- MASTER defines the service that Adaptive Server uses to listen to login requests from clients. This type defines a server machine.

  A MASTER entry is required only if you plan to use your computer as a server. It is not required in a *sql.ini* file for a computer that is running clients only.

- QUERY represents the service that a client application uses to log in to Adaptive Server. This type defines a client machine.

  A QUERY entry is required if you plan to use your computer to access a server. In general, since even dedicated servers need access to other servers, a QUERY entry is always required.

Server address

This value is the address at which Adaptive Server listens for client connections. The address requires the following information:

- Address format

- IP address
- Named Pipes format
- Windows Sockets format
- NWLink IPX/SPX format

Address format

The format of the server address depends on the network driver used by Adaptive Server.

The format for the server address can be:

- Named Pipes format
- Windows Sockets format
- NWLink IPX/SPX format

Use the following guidelines to define your server address:

- Some formats require a port, or socket number. Port numbers for MASTER and QUERY entries must be the same on server and client. For example, if a server is listening on 5000, the client workstation must be connecting on 5000.
- The server usually controls the port number, which means that you specify the same port number in the client's *sql.ini* file as that specified in the *sql.ini* file for the server to which it will connect.
- Port addresses must be unique to each server. The port address is determined by the port number provided in the *sql.ini* file in conjunction with the IP address.
- By default, the port number for Adaptive Server is 5000, for Backup Server, it is 5001.

Note Two Adaptive Servers on different computers can use the same port number because their IP addresses are different.

IP address

If you know a computer's IP address as well as its name, specify the IP address in the *sql.ini* file to ensure that the computer can be found on the network.

For example, the following entry, which uses Named Pipes, specifies a remote server's computer name and requires name resolution:

```
NLMSNMP,\\SMOKE\pipe\sybase\query
```

The following entry uses a remote server's IP address and does not require name resolution:

```
NLMSNMP,\\130.214.60.230\pipe\sybase\query
```

Named Pipes format

For the Named Pipes protocol, the network address consists of the unique pipe name for the server.

Use the following guidelines to create acceptable pipe names.

- Valid pipe names begin with *\pipe* and follow the same naming restrictions as MS-DOS file names. The default pipe name for Adaptive Server is *\pipe\sybase\query*.
- To avoid conflict, always use unique pipe names of the same "length" (levels) for all Sybase products on your computer. For example, you might select *\pipe\sybase\query* for Adaptive Server and *\pipe\backup\query* for Backup Server.
- Do not use pipe names such as *\pipe\sql* and *\pipe\sql\query*, because they do not ensure uniqueness.
- When adding a network entry to access a server on a remote network computer, such as on a client, preface the pipe name for the QUERY service with the following, where *machine_name* is the name of the computer that runs the server:

  ```
  \\machine_name
  ```

Warning! Server pipes must be local. Do not add *\\machine_name* if you are configuring a network entry for a server on a local computer. Additionally, do not preface the pipe name with this prefix when entering connection information for the MASTER service. If you include this prefix, you cannot restart Adaptive Server.

Windows Sockets format

For the Windows Sockets protocol, the server address consists of the TCP/IP host name or IP address of the Windows computer and a unique socket for the Adaptive Server, separated by a comma.

Keep the following guidelines in mind when creating the address:

- The TCP/IP host name is case-sensitive. For example, a possible entry for a TCP/IP host named "CENTAUR" is "CENTAUR, 5000".
- Adaptive Server uses the default socket number of 5000 to listen to connections from client workstations. Select a different socket number if another application on your computer already uses socket 5000.
- Valid socket numbers for Adaptive Server range from 1025 to 65535, in integers.

Increasing Windows Sockets connections

To support more than 64511 Windows Sockets (TCP/IP) connections to Adaptive Server, you may need to use the Windows Registry to increase the maximum number of sockets connections available on the server.

Warning! Do not modify a Registry value unless you are an Windows administrator and are familiar with the regedt32 utility. See your Windows operating system documentation for information on using regedt32.

❖ **Modifying an existing TcpNumConnections value**

1. Log in to Windows using an account with Windows administrator privileges, or use the default "sa" login.
2. Start the regedt32 utility from the run prompt.
3. Select the Registry window HKEY_LOCAL MACHINE.
4. Open the Registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters.
5. If the TcpNumConnections value exists, go to step 6.

   If the value does not exist, add and configure it by completing the steps under "To add a TcpNumConnections value."
6. Double-click the value.
7. In the DWORD Editor dialog box, select the Decimal option.

8 In the Data text box, enter the maximum number of connections to support.

9 Click OK to return to the Registry key dialog box.

10 If you have completed your tasks in regedt32:

1 Select Exit from the Registry menu to quit regedt32.

2 Restart your computer.

❖ **Adding a TcpNumConnections value**

1 Complete the Add Value dialog box as follows:

Value Name – TcpNumConnections

Data Type – select REG_DWORD from the drop-down list.

2 Click OK.

3 Complete the DWORD Editor dialog box as follows:

Data – enter the maximum number of TCP connections for the computer.

Radix – select the Decimal option button.

4 Click OK.

The utility adds the new value to the Registry key.

5 If you have completed your tasks in regedt32:

1 Choose Exit from the Registry menu to quit regedt32.

2 Restart your computer.

Using multiple TCP/IP network interface cards

When client workstations use multiple TCP/IP network interface cards, the Windows Administrator must edit the *lmhosts* file on the Windows server to accept connections from the clients.

To correctly enter the card information:

- There must be one entry for each network card.
- Each address must be unique.
- The entries in the *lmhosts* file match those in the *sql.ini* file.

For example, assume that a server named BALCLUTHA has three cards. Without an *lmhosts* entry and separate entries in *sql.ini*, the server listens on socket BALCLUTHA,5000 for all three cards. To provide unique addresses, set up *lmhosts* as follows:

```
130.214.10.248    NT0
130.214.11.248    NT1
130.214.12.248    NT2
```

In the *sql.ini* file, add entries for both QUERY and MASTER:

```
[BALCLUTHA]
query=NT0,5000
master=NT0,5000
query=NT1,5000
master=NT1,5000
query=NT2,5000
master=NT2,5000
```

Controlling the connection timeout

When an isql connection remains idle for several minutes, the next query may result in this error message:

```
Attempt to initiate a new SQL Server operation with
results pending.
```

This problem occurs when you use the Windows Sockets protocol, and you have a small value for Windows TcpKeepTries. To correct this problem, you must increase the value in Windows TcpKeepTries.

Warning! Do not modify Registry values unless you are an Windows Administrator and you are familiar with the regedt32 utility. See the Windows operating system documentation for information on using regedt32.

❖ **Increasing the TcpKeepTries value**

1 Start the regedt32 utility, and display the Parameters values.

2 Double-click the TcpKeepTries value.

3 Change the data value to at least 20, and choose OK.

4 If you have completed your tasks in regedt32:

   1 Choose Exit from the Registry menu to quit.

2 Reboot your computer.

NWLink IPX/SPX format

Before setting up Adaptive Server network support, configure the NWLink IPX/SPX software according to the instructions for your Windows operating system. Specify the correct network number (usually 0) and frame type during the configuration.

The frame type is generally mandated by the frame type of a NetWare file server on the network, usually 802.3. If your network does not use a NetWare file server, make sure all client and server computers use the same frame type.

Available NWLink IPX/SPX connection formats

Table 4-1 describes the available connection formats for NWLink IPX/SPX MASTER and QUERY entries.

Table 4-1: Connection information formats for IPX/SPX

Format	Connection information syntax	Example
1	*net_number*,*node_number*,*socket_number*	00000000,02608CDA1997,83BD
2	*cotmputer_name*,*socket_number*	piano,83BD
3	*computer_name*	piano

Keep the following items in mind when working with these formats:

- Any of the three formats is acceptable for the MASTER entry.
- Only Format 1 and Format 3 are acceptable for QUERY entries.
- Some formats are not acceptable for accessing a local Adaptive Server.

For more information, see “Selecting valid connection formats” on page 42.

In Table 4-1, *net_number* is the network number that you specified during the NWLink IPX/SPX configuration.

To find the network number:

1 Open Network and Dial-up Connections from Windows Control Panel.

2 In the Network and Dial-up Connections dialog box, right-click Local Area Connection, and then click Properties.

3 In the Local Area Connection Properties dialog box, double-click the entry: “NWLink IPX/SPX/NetBIOS Compatible Transport Protocol.”

4 The NWLink IPX/SPX/NetBIOS Compatible Transport Protocol dialog box is displayed. The current network number is the decimal number displayed in the "Internal network number" edit box

To determine the *node_number*, enter the net config command at the Windows command prompt. For example:

```
net config workstation
Computer name              \\PIANO
User name                  user1
Workstation active on      NBT_Elnk31 (00A0242EA892)
Software version           Windows 4.0
Workstation domain         AMERICAS
Logon domain               AMERICAS
COM Open Timeout (sec)     3600
COM Send Count (byte)      16
COM Send Timeout (msec)    250

The command completed successfully.
```

In the preceding example:

- The *node_number*, which is a 4-byte, hexadecimal number in the connection information string, appears in parentheses; "00A0242E".
- The *socket_number*, which can be any unused socket number on the computer, in 2-byte, hexadecimal format, appears with the *node_number*; "A892".
- The *computer_name* can be any unique name on the network. Use the local computer's name to ensure uniqueness.

Selecting valid connection formats

The NWLink IPX/SPX connection formats you use depend on whether you want to access Adaptive Server on a local computer or on a remote, network computer.

- When both Adaptive Server and the client program reside on the same computer, a local connection, use a Named Pipes connection.
- If you must use NWLink IPX/SPX for a local connection, follow these guidelines:
  - Use either Format 1 or Format 2 for the MASTER entry.
  - Use only Format 1 for the QUERY entry.

- If Adaptive Server and its clients reside on separate computers, a remote connection, you have two options:
  - Use Format 3 for both the MASTER and QUERY entries.
  - Use either Format 1 or Format 2 for the MASTER entry, but use Format 1 for the QUERY entry.

Sharing network configuration information

There are two ways to share identical network information across multiple systems:

- Create a master interfaces (*sql.ini*) file.
- Use Windows Registry as a directory service.

Creating a master *sql.ini* file

A master *sql.ini* file contains entries for all Sybase servers on the network. It can be used with every server and client connected to the network. By distributing copies of a master *sql.ini* file, you can ensure that all Sybase products on the network interact with one another.

To maintain consistency in the *sql.ini* files on a network, make the changes to one version of the file, and then copy that file to the rest of the computers on the network. For this task, you can use Windows File Replication to copy the file to many computers. For more information, see your Windows operating system documentation.

Using Windows Registry as a directory service

Another option is to use the Windows Registry as a directory service. Review the following Sybase product arrangements before settling on this method:

- Adaptive Server Enterprise only – you can deploy an application on multiple clients and enter the network information once in the Registry on the Adaptive Server computer without needing to create and maintain a *sql.ini* file on every client.

- Adaptive Server Enterprise and its bundled applications – the client applications that are bundled with Adaptive Server require a *sql.ini* file. Even if you are using the Registry for your own applications, you still need to maintain a *sql.ini* file if users are to connect from any of the Sybase client applications, such as Sybase Central.

The following instructions create server name keys under the Registry key specified for "ditbase" in *libtcl.cfg*, in the example in step 2, SOFTWARE\SYBASE\SERVER. It also stores the network information in the keys.

Both the Adaptive Server and client applications look in the Registry for network information before searching the *sql.ini* file.

You will need both the Open Client/Open Server Configuration and the dsedit utilities.

❖ **Using Windows Registry as a Directory Service**

1. Make sure the *ocscfg.dat* file is in your *d:\sybase\OCS-15_0\bin* directory.
2. Start the OC OS Config Utility.
    1. Select Start | Programs | Sybase | Connectivity | Open Client Directory Service Editor.
    2. Select the Directory Services tab.
    3. Click Add.
    4. Type REGISTRY for the DS Name.
    5. Type LIBDREG for the Directory Service Driver or select it from the drop-down list.
    6. Type *machine_name*:SOFTWARE\SYBASE\SERVER for the Directory Service Ditbase, where *machine_name* is the name of the computer that stores the network information.
    7. Click OK. The values you entered appear on the Directory Services dialog box.

    You can also use a text editor to add the following lines to the *libtcl.cfg* file:

```
[NT_DIRECTORY]
REGISTRY=LIBDREG ditbase=\\machine_name:SOFTWARE\SYBASE\SERVER
```

For information about using ocscfg, see the *Open Client/Server Configuration Guide for Desktop Platforms*.

3 Start dsedit.

   1 Select Start | Programs | Sybase | Connectivity | Open Client Directory Service Editor.

   2 In the Select Directory Service dialog box, select Registry from the DS Name list, and click OK.

   3 Follow the instructions in "How a client accesses Adaptive Server" on page 32 for creating server entries using dsedit.

Verifying server connections

After you configure your network connections, use the dsedit utility to verify that you can connect to a server. dsedit includes a network diagnostic utility that checks to see whether a process is listening at the specified address.

You can access this diagnostic utility in one of two ways:

- By choosing Server Object, then Server Ping from the dsedit menu, or
- By pressing the Ping (lightening bolt) key on your keyboard.

See "The dsedit Server ping utility" on page 175, for information about using dsedit to test connections.

Configuring ODBC connections

Some client applications do not connect to Adaptive Server directly through the Open Client software, but through the ODBC (Open Database Connectivity) driver instead.

For example, PowerDesigner™ connects through the ODBC driver. Other third-party applications may also require the ODBC driver.

For Adaptive Server versions previous to 12.5, the ODBC connections are built on top of the Open Client Client-Library, so you need to install the Open Client software on the clients where you install the ODBC Driver.

You can also install the driver separately on other client workstations on which you will be running third-party or developed products.

For more information about the ODBC driver, see the *ODBC Driver Reference Guide*.

To use ODBC connections, you need to configure the Adaptive Server ODBC driver to allow connection to Adaptive Server.

Configuring the ODBC driver

When you configure the ODBC driver to connect to Adaptive Server, you create an ODBC data source. You can configure more than one data source for Adaptive Server. For example, you might want one data source for each database.

❖ **Configuring a data source**

1 Start the ODBC Data Source Administrator (*odbcad32.exe*) from the Windows System program group.

For more information about ODBC, see your Windows operating system documentation.

2 Click on the System DSN tab to display the System Data Sources dialog box.

The dialog box appears with a list of sources you might have already defined.

3 Click Add to add a new Data Source.

The Create New Data Source dialog box appears.

4 Select Adaptive Server ODBC Driver as the driver you want to use for Adaptive Server, and click Finish.

The ODBC Adaptive Server Setup dialog box appears.

5 Complete the dialog box as follows:

Data Source Name – enter a short description of the Adaptive Server that is meaningful to you. For example, if you are creating the data source to connect to a specific Adaptive Server database, include the database name in the description.

Description (optional) – a long description of a data source name; for example, “Accounting database on Adaptive Server 3.”

6 Click the Options button to display the Login box.

7 Type the name of the database to which you want to connect in the Database Name text box:

   - For a PowerDesigner connection, you do not need to specify a database unless you want to reverse-engineer it. In this case, to "reverse-engineer" means to create a database and then determine its schema, rather than using the normal process of creating the schema first and then creating the database.

   You can fill in values for the other parameters in the box. For information about each parameter, see the online help or the *ODBC Driver Reference Guide* in Technical Library.

8 Click OK, and close the rest of the ODBC dialog boxes.

9 Exit the program.

You can now connect to Adaptive Server from applications that require connections through the ODBC Driver. When you start the application and it prompts you for an ODBC data source, choose the data source you have just named and configured.

IPv6 support

Adaptive Server supports IPv6 technology.

Understanding IPv6

IPv6 addressing terminology:

- Link-local address – an IPv6 address that is usable only over a single link.
- Site-local address – an IPv6 address that can be used within a single-site.
- Global address – an IPv6 address that can be used across the global Internet.

Note Interfaces files also provide IPv6 support.

IPv6 application types:

- IPv6-unaware – an application that cannot handle IPv6 addresses.

- IPv6-aware – an application that can communicate with nodes that do not have IPv4 addresses. In some cases, this might be transparent to the application, for instance when the API hides the content and format of the actual addresses.
- IPv6-enabled – an application that, in addition to being IPv6-aware, takes advantage of some IPv6 features.
- IPv6-required – an application that requires some IPv6 features and cannot operate over IPv4.

IPv6 Infrastructure:

IPv6 infrastructure

Dual Stack infrastructure implements both IPv4 and IPv6. This is the recommended infrastructure implementation for using Adaptive Server Enterprise as an IPv6-aware server.

Sybase applications are IPv6-aware. All code to turn Sybase™ Adaptive Server and the Open Client/Server components IPv6-aware was done using the IETF designed primitives.

The following matrix lists the platform run-time requirements and the specific product and its release version:

Table 4-2: IPv6 support

Platform	Adaptive Server IPv6 awareness	Open Client/Server IPv6 awareness
Sun Solaris 8 32- and 64-bit	12.5.3a and 15.0	12.5 and 15.0
HP-UX 11i(v1) 32- and 64-bit	12.5.3a and 15.0	12.5 and 15.0
Microsoft Server 2003	12.5.3a and 15.0	12.5 and 15.0
Linux RHEL 3.0	15.0	12.5 and 15.0

Many Sybase products that are Open Client/Server based like XP Server, Backup Server, Replication Server and Open Switch became automatically IPv6-aware due to the layered Open Client Transport Control Layer (CTlib->NETlib) which is IPv6-aware for network-socket operations. An important note is that any DBlib based Open Client product is not IPv6-aware.

For Adaptive Server Enterprise being IPv6-aware is a complex issue because some components within the ASE are 3rd party components and are not yet IPv6-aware. To understand how this impacts Adaptive Server Enterprise the following list shows all functional mechanisms of Adaptive Server Enterprise that are IPv6-aware with respect to the platform / release matrix above:

- Connection Handler
- RPC mechanisms
- Job Scheduler Task / Agent session connection
- Network Host API
- UDP Message support for sybsendmsg
- Component Integration Services connectivity
- Host / name resolving
- XML URL connection handler
- Auditing for client address data

The following functional mechanisms in Adaptive Server Enterprise do not support IPv6. These mechanisms in Adaptive Server Enterprise are IPv6-unaware. They will gradually (over time) be become IPv6-aware in follow-on releases:

- Java support
- License Management Server
- LDAP driver

Before starting Adaptive Server for IPv6-aware operations, make sure that your infrastructure is correctly set up. Once your operating system is correctly configured, an IPv6 connection handler can be configured and enabled. Configuring and enabling the IPv6 connection handler requires adding an additional DCL entry. A single Adaptive Server configuration can typically carry up to 32 connection handler assignments within the DCL.

For example if you have a Site-local setup with two domains administrated under the name server setup:

```
sybase.com - being responsible for all IPv4 networking applications
v6.sybase.com - being responsible for all IPv6 networking applications
```

The DCL entry for Adaptive Server to start named “SYBASE” on the host “revival” for port 17100 would typically look like:

```
SYBASE
```

```
master tcp ether revival.sybase.com 17100
query tcp ether revival.sybase.com 17100
master tcp ether revival.v6.sybase.com 17100
query tcp ether revival.v6.sybase.com 17100
```

In the above example, when Adaptive Server is started with IPv6-awareness it creates two connection handlers. One listens on port 17100 for incoming IPv4 Clients connection requests, and the other listens on port 17100 for incoming IPv6 Clients connection requests.

Note When you start Adaptive Server, you can set Trace flag 7815 to capture and log IPv4 and IPv6 client address connection requests and host / name lookups.

CHAPTER 5

Lightweight Directory Access Protocol in Adaptive Server

Adaptive Server uses directory services to establish client and RPC connections over the Internet. This chapter provides information about using LDAP directory services to establish connections.

Overview

Lightweight Directory Access Protocol (LDAP) is an industry standard for accessing directory services. Directory services allow components to look up information by a distinguished name (DN) from an LDAP server that stores and manages server, user, and software information that is used throughout the enterprise or over a network.

The LDAP server can be located on a different platform from the one on which Adaptive Server or the clients are running. LDAP defines the communication protocol and the contents of messages exchanged between clients and servers. Messages are operators, such as client requests for read, write and query, and server responses, including data-format information.

The LDAP server stores and retrieves information about:

- Adaptive Server, such as IP address, port number, and network protocol
- Security mechanisms and filters
- High availability companion server name

The LDAP server can be configured with these access restrictions:

- Anonymous authentication – all data is visible to any user.
- User name and password authentication – Adaptive Server uses the default user name and password from Windows – *%SYBASE%\%SYBASE_OCS%\ini\libtcl.cfg*

User name and password authentication properties establish and end a session connection to an LDAP server.

Note The user name and password that are passed to the LDAP server for user authentication purposes are distinct and different from those used to access Adaptive Server.

When an LDAP server is specified in the *libtcl.cfg* file the server information is accessible only from the LDAP server. Adaptive Server ignores the interfaces file.

If multiple directory services are supported in a server, then the order in which they are searched is specified in *libtcl.cfg*. You cannot specify the search order with the dataserver command-line option. See “Multiple directory services” on page 58.

LDAP directory services versus the Sybase interfaces file

The LDAP driver implements directory services for use with an LDAP server. LDAP directories are an infrastructure that provide:

- A network-based alternative to the traditional Sybase interfaces file
- A single, hierarchical view of information, including users, software, resources, networks, files, and so on

Table 5-1 highlights the differences between the Sybase interfaces file and an LDAP server.

Table 5-1: interfaces file versus LDAP directory services

interfaces file	Directory services
Platform-specific	Platform-independent
Specific to each Sybase installation	Centralized and hierarchical
Contains separate master and query entries	One entry for each server that is accessed by both clients and servers
Cannot store metadata about the server	Stores metadata about the server

LDAP directory services support more attributes than the Sybase interfaces file. These attributes can include server version, server status, and so on. See Table 5-2 for a list of attributes.

Note LDAP is only supported with reentrant libraries. You must use isql_r, instead of isql, when connecting to a server using LDAP directory services.

Table 5-2 lists the Sybase LDAP directory entries.

Table 5-2: Sybase LDAP directory definitions

Attribute name	Value type	Description
ditbase	*interfaces* file or *libtcl.cfg*	DIT base for object tree. If the *libtcl.cfg* file is specified, the *interfaces* file is ignored. The *libtcl.cfg* file can be overridden with ct_con_prop() for a specified connection.
dn	Character string	Distinguished name. Must be unique name that identifies the object.
sybaseVersion	Integer	Server version number.
sybaseServername	Character string	Server name.
sybaseService	Character string	Service type: Sybase Adaptive Server, or Sybase SQL Server.
sybaseStatus	Integer	Status: 1 = Active, 2 = Stopped, 3 = Failed, 4 = Unknown.
sybaseAddress	String	Each server address includes: • Protocol: TCP, NAMEPIPE, SPX DECNET (entry is case-sensitive). • Address: any valid address for the protocol type. **Note** dscp splits this attribute into Transport type and Transport address.
sybaseSecurity (optional)	String	Security OID (object ID).
sybaseRetryCount	Integer	This attribute is mapped to CS_RETRY_COUNT, which specifies the number of times that ct_connect retries the sequence of network addresses associated with a server name.

Attribute name	Value type	Description
sybaseRetryDelay	Integer	This attribute is mapped to CS_LOOP_DELAY, which specifies the delay, in seconds, that ct_connect waits before retrying the entire sequence of addresses.
sybaseHAservername (optional)	String	A secondary server for failover protection.

The traditional interfaces file with TCP connection and a failover machine looks like:

```
master tcp ether huey 5000
query tcp ether huey 5000
hafailover secondary
```

An example of an LDAP entry with TCP and a failover machine looks like:

```
dn: sybaseServername=foobar, dc=sybase,dc=com
objectClass: sybaseServer
sybaseVersion: 1500
sybaseServername: foobar
sybaseService: ASE
sybaseStatus: 4
sybaseAddress: TCP#1#foobar 5000
sybaseRetryCount: 12
sybaseRetryDelay: 30
sybaseHAServernam: secondary
```

All entries in the LDAP directory service are called entities. Each entity has a distinguished name (DN) and is stored in a hierarchical tree structure based on its DN. This tree is call the **directory information tree** (DIT). Client applications use a DIT base to specify where entities are stored. See “The libtcl.cfg file” on page 55.

In the example above, the entry describes an Adaptive Server named “foobar” listening on a TCP connection with a port number of 5000. This entity also specifies a retry count of 12 (times) and a retry delay of 30 (seconds). Once a client has found an address where a server responds, the login dialog between the client and the server begins.

You can find a complete list of the Sybase LDAP directory schema in Windows – *%SYBASE%\%SYBASE_OCS%\ini*.

In the same directory, there is also a file called *sybase-schema.conf*, which contains the same schema, but uses a Netscape-specific syntax.

Since LDAP supports multiple entries for each attribute, each address attribute must contain the address of a single server, including protocol, access type, and address. See sybaseAddress in Table 5-2.

For example, this is an LDAP entry for an Windows server listening on two addresses, with different connection protocols:

```
sybaseAddress = TCP#1#TOEJAM 4444
sybaseAddress = NAMEPIPE#1#\pipe\sybase\query
```

Note Each entry in the address field is separated by the # character.

You can edit these entries with dsedit. See "Adding a server to the directory services" on page 57.

To ensure cross-platform compatibility for all Sybase products, the protocol and address attribute fields should be in a platform- and product-independent format.

The *libtcl.cfg* file

You use the *libtcl.cfg* file to specify the LDAP server name, port number, DIT base, user name, and password to authenticate the connection to an LDAP server.

The purpose of the *libtcl.cfg* file is to provide configuration information such as driver, directory, and security services for Open Client/Open Server and Open Client/Open Server-based applications. 32-bit utilities such as dsedit and srvbuild, look up the *libtcl.cfg*.

The default *libtcl.cfg* file is located in *%SYBASE%|%SYBASE_OCS%\ini*.

If LDAP is specified in the *libtcl.cfg* file, the interfaces file is not used.

Note Open Client/Open Server applications that use the -I option at start-up override the *libtcl.cfg* file and use the interfaces file.

In its simplest form, the *libtcl.cfg* file is in this format:

```
[DIRECTORY]
ldap=libsybdldap.dll  ldapurl
```

where the *ldapurl* is defined as:

```
ldap://host:port/ditbase
```

The following LDAP entry, using these same attributes, is an anonymous connection and only works only if the LDAP server allows read-only access.

```
ldap=libsybdldap.dll ldap://seashore/d=sybase,dc=com
```

You can specify a user name and password in the *libtcl.cfg* file as extensions to the LDAP URL to enable password authentication at connection time.

Enabling LDAP directory services

To use a directory service, you must:

1 Configure the LDAP server according to the vendor-supplied documentation.

2 Add the location of the LDAP libraries to the PATH environment variable for your platform.

3 Configure the *libtcl.cfg* file to use directory services.

Use any standard ASCII text editor to:

- Remove the semicolon (;) comment markers from the beginning of the LDAP URL lines in the *libtcl.cfg* file under the *[DIRECTORY]* entry.
- Add the LDAP URL under the *[DIRECTORY]* entry. See Table 5-3 for supported LDAP URL values.

Warning! The LDAP URL must be on a single line.

```
ldap=libsybdldap.dll
ldap://seashore/dc=sybase,dc=com??one??bindname=uid=Manager,dc=sybase,
dc=com?password
```

For example:

```
[DIRECTORY]
ldap=libsybdldap.dll
ldap://seashore/dc=sybase,dc=com??one??bindname=uid=Manager,dc=sybase,
dc=com?password
```

"*one*" indicates the scope of a search that retrieves entries one level below the DIT base.

Table 5-3 defines the keywords for the *ldapurl* variables.

Table 5-3: ldapurl variables

Keyword	Description	Default
host (required)	The host name or IP address of the machine running the LDAP server	None
port	The port number that the LDAP server is listening on	389
ditbase (required)	The default DIT base	None
username	Distinguished name (DN) of the user to authenticate	NULL (anonymous authentication)
password	Password of the user to be authenticated	NULL (anonymous authentication)

4 Verify that the appropriate environment variable points to the required third-party libraries. The Netscape LDAP SDK libraries are located in *%SYBASE%\%SYBASE_OCS\lib3p*.The Windows PATH environment variable must point to this directory.

5 Add your server entry to the LDAP server using dscp or dsedit. See "Adding a server to the directory services" on page 57.

Adding a server to the directory services

Warning! Most LDAP servers have an ldapadd utility for adding directory entries. Sybase recommends you use dsedit instead since it has built-in semantic checks that generic tools do not provide.

Each server entry is made up of a set of attributes. When you add or modify a server entry, you are prompted for information about server attributes. Some attributes are provided by default, others require user input. When a default value is provided, it appears in brackets "[ ]". See Table 5-2 for accepted values.

You can use srvbuild to add entries, but not modify or delete them.

❖ **Adding a server entry to the directory service using *dsedit***

Before you can add, delete, or modify an LDAP server entry, you must add the LDAP URL to the *libtcl.cfg* file. See "The libtcl.cfg file" on page 55.

Use dsedit to add a server to the directory service:

1 From the Windows task bar, select Start | Programs | Sybase | Connectivity | Open Client Directory Service Editor.

2 Select LDAP from the list of servers, and click OK.

3 Click Add New Server Entry.

4 Enter:

- The server name – this is required.
- Security mechanism – optional. A list of security mechanism OIDs are located in *%SYBASE%\ini\objectid.dat*.
- HA server name – optional. This is the name of the high-availability failover server, if you have one.

5 Click Add New Network Transport.

- Select the transport type from the drop-down list.
- Enter the host name.
- Enter the port number.

6 Click OK two times to exit the dsedit utility.

To view the server entries, enter the following URL in Netscape:

```
ldap://host:port/ditbase??one
```

For example:

```
ldap://huey:11389/dc=sybase,dc=com??one
```

Note Microsoft Internet Explorer does not recognize LDAP URLs.

For more information about dscp, see the *Open Client/Server Configuration Guide*, in the 11.1.x Generic Collection at http://www.sybase.com/support/manuals.

Multiple directory services

Any type of LDAP service, whether it is an actual server or a gateway to other LDAP services, is called an LDAP server.

You can specify multiple directory services for high-availability failover protection. Not every directory service in the list needs to be an LDAP server.

For example:

```
[DIRECTORY]
ldap=libsybdldap.dll
ldap://seashore/dc=sybase,dc=com??one??bindname=uid=Manager,dc=sybase,
dc=com?password
```

In this example, if the connection to *test:389* fails, the connection fails over to the DCE driver with the specified DIT base. If this also fails, a connection to the LDAP server on *huey:11389* is attempted. Different vendors employ different DIT base formats.

Note For more information, see the *Open Client Client-Library/C Programmer's Guide* and the *Open Client Client-Library/C Reference Manual* at http://www.sybase.com/support/manuals.

Encrypting the password

Entries in the *libtcl.cfg* file are in human-readable format. Sybase provides a pwdcrypt utility for basic password encryption. pwdcrypt is a simple algorithm that, when applied to keyboard input, generates an encrypted value that can be substituted for the password. pwdcrypt is located in *%SYBASE%\%SYBASE_OCS%\bin*.

From the *%SYBASE%\%SYBASE_OCS%* directory, enter:

```
bin/pwdcrypt
```

Enter your password twice when prompted.

pwdcrypt generates an encrypted password. For example:

```
0x01312a775ab9d5c71f99f05f7712d2cded2i8d0ae1ce78868d0e8669313d1bc4c706
```

Copy and paste the encrypted password into the *libtcl.cfg* file using any standard ASCII-text editor. Before encryption, the file entry appears as:

```
ldap=libsybdldap.dll
ldap://seashore/dc=sybase,dc=com??one??bindname=uid=Manager,dc=sybase,
dc=com?password
```

Replace the password with the encrypted string:

```
ldap=libsybdldap.dll
ldap://seashore/dc=sybase,dc=com??one??bindname=uid=Manager,dc=sybase,dc=com
0x01312a775ab9d5c71f99f05f7712d2cded2i8d0ae1ce78868d0e8669313d1bc4c706
```

Warning! Even if your password is encrypted, you should still protect it using file-system security.

Performance

Performance when using an LDAP server may be slower than when using an interfaces file because the LDAP server requires time to make a network connection and retrieve data. Since this connection is made when Adaptive Server is started, changes in performance will be seen at login time, if at all. During normal system load, the delay should not be noticeable. During high system load with many connections, especially repeated connections with short duration, the overall performance difference of using an LDAP server versus the traditional interfaces file might be noticeable.

Migrating from the *sql.ini* file to LDAP

There is no direct method to upgrade an existing server using the *sql.ini* file to one that uses lightweight directory services. To upgrade a previous release of Adaptive Server to Adaptive Server version 15.0, see the *Installation Guide for Windows*.

If you have LDAP or other directory services defined in the *libtcl.cfg* file before configuring the server, the -i argument is not added to the *sql.ini* file.

If you do not have LDAP or other directory services defined in the *libtcl.cfg*, the -i argument is added to the Windows registry for your SYBASE server.

Once you have upgraded the server, you can configure your server to use LDAP service.

1 Shut down the server. See Chapter 2, “Starting and Stopping Servers.”

2 Edit the *%SYBASE%\%SYBASE_OCS%\ini\libtcl.cfg* file to add the directory service. See “Enabling LDAP directory services” on page 56.

3 Use dsedit and add the server entry to directory service. See “Adding a server to the directory services” on page 57.

4 Start the configuration utility. Select Start | Programs | Sybase | Adaptive Server Enterprise | Server Config.

5 Select Configure Adaptive Server.

6 Select the server for which you want to enable directory service, and click Continue.

7 Enter your log in name and password, and click Continue.

8 When prompted to start the server, select Yes.

9 On the Configure Adaptive Server screen, click “Cancel” or “Save”.

10 Exit Server Config.

Alternatively, you can add or remove the -i argument which specifies the interfaces (*sql.ini* on Windows) file directly from the Windows registry.

1 Select Start | Run and enter, regedt32.

2 Select the HKEY_LOCAL_MACHINE view.

3 Select SOFTWARE\Sybase\Server*server_ name*\Parameters

4 Remove the -i argument from the line that ends with ...*Sybase**ini**sql.ini*

CHAPTER 6

Customizing Localization for Adaptive Server

This chapter provides information about Sybase localization support for international installations. It also includes information for reconfiguring localization.

This chapter provides only the information that you need to know to configure languages, character sets, and sort order. For more information, see the *System Administration Guide: Volume 1*.

Overview of localization support

Localization is setting up an application to run in a particular language or country environment, including translated system messages and correct formats for date, time, and currency. Adaptive Server supports localization for international customers and for customers with heterogeneous environments.

This support includes:

- Data processing support – Adaptive Server comes with character set and sort-order definition files it uses to process the characters used in different languages.

  Sybase provides support for the major languages in:

  - Western Europe

    - Eastern Europe
    - Middle East
    - Latin America
    - Asia
- Translated system messages – Adaptive Server includes language modules for:
    - Brazilian Portuguese
    - Chinese (Simplified)
    - French
    - German
    - Japanese
    - Korean
    - Polish
    - Spanish
    - Thai
- Translated documentation – translated documentation is available in:
    - Chinese (Simplified)
    - French
    - German
    - Japanese
    - Korean
    - Polish
    - Spanish

Language modules

Adaptive Server stores its localized software messages in separate language modules.

When you install a language module, the installation program loads the messages, character set, and sort-order files that support the new language in the correct locations.

When you install Adaptive Server and Backup Server, system messages in English are installed by default.

Default character sets for servers

The default character set is the character set in which data is encoded and stored on the Adaptive Server databases.

By default, when Adaptive Server and Backup Server are installed on Windows systems, the installation installs the character set files for CP 850 which supports the Western European languages.

Changing the default character set for servers

You can select any character set as the default on Adaptive Server, including character sets that are not the platform default character sets. Keep the following guidelines in mind when selecting a new default character set:

- To avoid conversion errors or overhead, determine the default character set based on the character set used by your clients.

  For example, if most of your clients use ISO 8859-1, you can minimize the amount of data conversion that has to occur by specifying ISO 8859-1.

- If your server is operating in a heterogeneous environment, choose a character set that will work with all the character sets needed. Often, this is Unicode (UTF-8).

Warning! Make all changes to the default character set and sort order for a new Adaptive Server before creating any user databases or making any changes to the Sybase-supplied databases. Changing the character set and sort order after data or data structures have been added to Adaptive Server can cause incorrect behavior. To change the character set or sort order after you have added data, see the *System Administration Guide: Volume 1*.

Supported character sets

The following language, scripts and character sets are supported by Adaptive Server:

- Arabic – See Table 6-1 on page 67.
- Baltic – See Table 6-2 on page 67.
- Chinese, Simplified – See Table 6-3 on page 67.
- Chinese, Traditional – See Table 6-4 on page 67.
- Cyrillic – See Table 6-5 on page 68.
- Eastern European – See Table 6-6 on page 68.
- Greek – See Table 6-7 on page 68.
- Hebrew – See Table 6-8 on page 68.
- Japanese – See Table 6-9 on page 69.
- Korean – See Table 6-10 on page 69.
- Thai – See Table 6-11 on page 69.
- Turkish – See Table 6-12 on page 69.
- Unicode – See Table 6-13 on page 69.
- Vietnamese – See Table 6-14 on page 69.
- Western European – See Table 6-15 on page 70.

The tables define each character set and indicate information on whether it requires Unilib™ conversion (Unilib Required column).

- Checkmark (x) – the character set requires Unilib conversion.
- No checkmark – the character set may use either the Unilib conversion or the built-in conversion.

For more information see "Character set conversion" on page 70.

Table 6-1 lists the Arabic character set:

Table 6-1: Arabic character sets

Character set	Unilib required	Description
cp864	X	PC Arabic
cp1256	X	Microsoft Windows Arabic
iso88596	X	ISO 8859-6 Latin/Arabic

Table 6-2 lists the Baltic character set:

Table 6-2: Baltic character sets

Character set	Unilib required	Description
cp1257	X	Microsoft Windows Baltic

Table 6-3 lists the simplified Chinese character set:

Table 6-3: Simplified Chinese character sets

Character set	Unilib required	Description
eucgb	X	EUC GB encoding = Simplified Chinese character sets
cp936	X	Microsoft Simplified Chinese character sets
gb18030	X	RC 18030 standard

Table 6-4 lists the traditional Chinese character set:

Table 6-4: Traditional Chinese character set

Character set	Unilib required	Description
cp950	X	PC (Microsoft) Traditional Chinese
euccns	X	EUC CNS encoding = Traditional Chinese with extensions
big5	X	Big 5 Traditional Chinese
big5hk	X	Big 5 with HKSCS extensions

Table 6-5 lists the Cyrillic character set:

Table 6-5: Cyrillic character sets

Character set	Unilib required	Description
cp855		IBM PC Cyrillic
cp866		PC Russian
cp1251		Microsoft Windows 3.1 Cyrillic
iso88595		ISO 8859-5 Latin/Cyrillic
koi8		KOI-8 Cyrillic
mac_cyr		Macintosh Cyrillic
kz1048		Kazakhstan Cyrillic

Table 6-6 lists the Eastern European character set:

Table 6-6: Eastern European character sets

Character set	Unilib required	Description
cp852		PC Eastern Europe
cp1250		Microsoft Windows 3.1 Eastern European
iso88592		ISO 8859-2 Latin-2
mac_ee		Macintosh Eastern European

Table 6-7 lists the Greek character set:

Table 6-7: Greek character sets

Character set	Unilib required	Description
cp869		IBM PC Greek
cp1253		MS Windows Greek
greek8		HP GREEK8
iso88597		ISO 8859-7 Latin/Greek
macgrk2		Macintosh Greek

Table 6-8 lists the Hebrew character set:

Table 6-8: Hebrew character sets

Character set	Unilib required	Description
cp1255	X	Microsoft Windows Hebrew
iso88598	X	ISO 8859-8 Hebrew

Table 6-9 lists the Japanese character set:

Table 6-9: Japanese character sets

Character set	Unilib required	Description
cp932	X	IBM J-DBCS:CP897 + CP301 (Shift-JIS)
deckanji		Digital UNIX JIS encoding
eucjis		EUC-JIS encoding
sjis		Shift-JIS (no extensions)

Table 6-10 lists the Korean character set:

Table 6-10: Korean character sets

Character set	Unilib required	Description
eucksc	X	EUC KSC Korean encoding = CP949

Table 6-11 lists the Thai character set:

Table 6-11: Thai client character sets

Character set	Unilib required	Description
tis620	X	TIS-620 Thai standard
cp874	X	Microsoft Windows Thai

Table 6-12 lists the Turkish character set:

Table 6-12: Turkish character sets

Character set	Unilib required	Description
cp857		IBM PC Turkish
cp1254		Microsoft Windows Turkish
iso88599		ISO 8859-9 Latin-5 Turkish
macturk		Macintosh Turkish
turkish8		HP TURKISH8

Table 6-13 lists the Unicode character set:

Table 6-13: Unicode character set

Character set	Unilib required	Description
utf8	X	Unicode UTF-8 encoding

Table 6-14 lists the Vietnamese character set:

Table 6-14: Vietnamese character set

Character set	Unilib required	Description
cp1258	X	Microsoft Windows Vietnamese

Table 6-15 lists the Western European character set:

Table 6-15: Western European character set

Character set	Unilib required	Description
ascii8	X	US ASCII, with 8-bit data, ISO 646
cp437		IBM CP437 - U.S. code set
cp850		IBM CP850 - European code set
cp860	X	PC Portuguese
cp863	X	IBM PC Canadian French code page
cp1252	X	Microsoft Windows US (ANSI)
iso_1		ISO 8859-1 Latin-1
mac		Standard Macintosh coding
roman8		HP ROMAN8
iso 885915	X	ISO 8859-15 Latin-1 with Euro support

Character set conversion

Backup Server passes messages to Adaptive Server in the client's language and in the Adaptive Server character set. Adaptive Server then converts the messages and issues them in the client's language and character set. Keep the following requirements in mind when selecting a character set:

- In a heterogeneous environment, Adaptive Server and Backup Server may need to communicate with clients running on different platforms and using different character sets. To maintain data integrity, the server converts the code between the character sets.
- To use the built-in conversion, you need to install the character set definition files on the server for all the character sets being used by your clients. Built-in conversion support is available for many character sets.
- Unilib conversion support is available for all character sets supported by Sybase. To enable Unilib conversion, you must use sp_configure and turn enable unicode conversions on. See the *System Administration Guide: Volume 1*.

If either Adaptive Server or Backup Server does not support a client's language or character set, that server issues a warning message. Errors also occur when the Backup Server character set is not compatible with the Adaptive Server character set. By default, Unicode conversion is enabled.

Character set conversion is supported only between character sets for the same language or between character sets in the same language group.

For example, automatic character set conversion is supported between the character sets for the Western European languages: ASCII 8, CP 437, CP 850, CP 860, CP 863, CP 1252, ISO 8859-1, ISO 8859-15, Macintosh Roman and ROMAN8. Similarly, conversion is supported between the character sets for Japanese: CP 932, EUC-JIS, Shift-JIS, and DEC-Kanji.

However, code conversion is not supported between any of the Western European language character sets and the Japanese character sets. For more information about supported conversions, see the *System Administration Guide: Volume 1*.

Conversions between server and client

If Adaptive Server does not support the client's language or character set, the client can connect with the server, but no character conversions will occur.

When a localized client application connects to Adaptive Server, the server checks to see if it supports the client's language and character set.

- If Adaptive Server supports the language, it automatically performs all character set conversions and displays its messages in the client's language and character set.
- If Adaptive Server does not support the language, it uses the user's default language or Adaptive Server's default language.
- If Adaptive Server does not support the character set, it issues a warning to the client, turns conversion off, and sets the language to U.S. English.

Sort orders

Each character set comes with one or more sort orders (collating sequences), which are located in the sort-order definition files (*.srt* files). These files accompany the character set definition files and can be found in the same directory.

You can select a sort order for your data according to the needs at your site. Keep in mind that the server can support only one sort order at a time, so select a sort order that will work for all of your clients.

Warning! Make all changes to the default character set and sort order for a new Adaptive Server before creating any user databases or making any changes to the Sybase-supplied databases. Changing the character set and sort order after data or data structures have been added to Adaptive Server may cause incorrect behavior. To change the character set or sort order after you have added data, see the *System Administration Guide: Volume 1*.

Available sort orders

The sort order determines the collating sequence Adaptive Server uses to order, compare, and index character data. Each character set comes with one or more sort orders.

Sort orders are located in sort order definition files (*.srt* files) that accompany your character set definition files.

Note Available sort orders vary according to the character set installed on Adaptive Server.

You can see the available sort orders for your character set by looking in the *.srt* file for your language. Sort orders are stored in the following path:

%SYBASE%\charsets\<charset_name>.srt*

For more information about localization files, see "Localization directories" on page 75.

Table 6-16 describes the sort orders that you can specify at installation time or at a later time using the syconfig utility.

Table 6-16: Sort orders available in Adaptive Server

Sort order name	Description
Binary order	Sorts all data according to numeric byte values for that character set. Binary order sorts all ASCII uppercase letters before lowercase letters. Accented or ideographic (multibyte) characters sort in their respective standards order, which may be arbitrary. All character sets have binary order as the default. If binary order does not meet your needs, you can specify one of the other sort orders either at installation or at a later time by, using the syconfig utility.

Sort order name	Description
Dictionary order, case-sensitive, accent-sensitive	case-sensitive. Sorts each uppercase letter before its lowercase counterpart, including accented characters. Recognizes the various accented forms of a letter and sorts them after the associated unaccented letter.
Dictionary order, case-insensitive, accent-sensitive	Case-insensitive dictionary sort order. Uppercase letters are equivalent to their lowercase counterparts and are intermingled in sorting results.
Dictionary order, case-insensitive, accent-insensitive	Case-insensitive dictionary sort order. Diacritical marks are ignored.
Dictionary order, case-insensitive with preference	Case-insensitive dictionary sort order, with case preference for collating purposes. A word written with uppercase letters is equivalent to the same word written with lowercase letters. Uppercase and lowercase letters are distinguished only when you use an order by clause. The order by clause sorts uppercase letters before it sorts lowercase. **Note** Do not select this sort order unless your installation requires that uppercase letters be sorted before lowercase letters in otherwise equivalent strings for order by clauses. Using this sort order may reduce performance in large tables when the columns specified in an order by clause match the key of the table's clustered index.
Alternate dictionary order, case-sensitive	Case-sensitive alternate dictionary sort order with lowercase variants sorted before uppercase. Use with several of the Western European languages.
Alternate dictionary order, case-insensitive, accent-insensitive	Case-insensitive and accent-insensitive alternate dictionary sort order. Use with several of the Western European languages.
Alternate dictionary order, case-insensitive, uppercase preference	Case-insensitive alternate dictionary sort order with uppercase preference. Use with several of the Western European languages.
Spanish dictionary order, case-sensitive	Case-sensitive Spanish dictionary sort order. Use with Spanish and for most Latin American locales.
Spanish dictionary order, case-insensitive	Spanish case-insensitive dictionary sort order. Use with Spanish and for most Latin American locales.
Spanish dictionary order case-insensitive, accent-insensitive	Spanish case-insensitive and accent-insensitive dictionary sort order. Use with Spanish and for most Latin American locales.
Scandinavian dictionary order, case-sensitive	Case-sensitive dictionary sort order. Use with Scandinavian languages.
Scandinavian dictionary order, case-insensitive, uppercase preference	Case-insensitive and accent-insensitive dictionary sorting, with uppercase preference. Use with Scandinavian languages.

To see the sort orders that are available, use Server Config to display the sort orders for the character sets you plan to use.

Language modules

If you want Adaptive Server error messages to be displayed in a language other than U.S. English (us_english), you must install the appropriate language module.

When you install a new language module, installation automatically loads the language into the Sybase installation directory to support the new language. For information about directories, see “Localization directories” on page 75.

Installing a new language module

A full install of Adaptive Server installs all the language components automatically. If you did not select a full install, you need to install additional language modules manually.

To install a new language module:

1 Load the language module software from the distribution media. You must load this software into the same directory in which you loaded Adaptive Server.

2 Reconfigure the language and, if necessary, the character set and sort order for Adaptive Server. For instructions, see “Changing the localization configuration” on page 79.

Message languages

For messages, U.S. English is installed as the default language in Adaptive Server. The following rules apply to language modules:

- During Adaptive Server installation or reconfiguration, you can specify a default language other than U.S. English. However, you must have installed the language module for the language you specify.

- If your clients require Adaptive Server messages in a language other than U.S. English, you must load the language module for those languages. Then, you can configure Adaptive Server to the language used by your clients.
- If Adaptive Server does not support messages in a client's language, these clients receive messages in the server default language.

  For example, if your client's language is Latin, the Spanish language module is installed, and Spanish is specified as the Adaptive Server default language, the client receives messages in Spanish.

Localization

By default, the Adaptive Server and Backup Server configurations use the English locale settings, which include:

- Character set definition files for Western European character sets
- Sort-order definition files for Western European character sets
- U.S. English system message files

During the installation process or through reconfiguration, you can specify a different language, character set, and sort order.

Localization directories

Sybase localization configuration involves the following directories:

- *locales*
- *charsets*

The table below illustrates the structure of the localization files. It does not show a complete list of all the files.

%SYBASE%/ or *$SYBASE/*	*charsets*	*charset_name*	**.srt* files
		charset_name...	*charset.loc*
		unicode	**.uct* files
	locales	*language_name*	charset_name
		language_name...	*charset_name...*
		locales.dat	
		message	*language_name*
			language_name...

About the directory

The *%SYBASE%\locales* or directory contains a subdirectory for each available language. Each language subdirectory contains a subdirectory for each character set available with that language.

- The *.loc* files in these subdirectories enable Adaptive Server or Backup Server to report errors in a specific language, encoded in a specific character set.

  There are a variety of *.loc* files in each subdirectory. Most of these files contain translated error messages for a specific product or utility.
- The *common.loc* file in the "utf8" subdirectory for each language contains localized information, such as local date, time, and currency formatting, that is used by all products.
- The *locales.dat* file contains entries that associate platform-specific locale names with Sybase language and character set combinations.

About the *charsets* directory

The files in *%SYBASE%\charsets\charset_name* contain information related to each particular character set, such as the definition of the character set and any sort orders available for that character set.

About the *locales.dat* file

You can edit the *locales.dat* file to:

- Change the default language or character set for a platform, or

- Add new associations between platform locale names and Sybase language and character set names.

Format of *locales.dat* file entries

Each entry in the *locales.dat* file links a platform-specific locale definition to a Sybase language and character set combination. Each entry has the following format:

```
locale = platform_locale, syb_language, syb_charset
```

where:

- *platform_locale* is the platform-specific keyword for a locale. For acceptable values, see your operating system documentation.

  When the locale being defined is the default for the site, *platform_locale* is "default".

- *syb_language* is the name of the language directory to be used from within *%SYBASE%\locales\language_name*.

- *syb_charset* is the character set name that determines the character set conversion method and identifies the directory location of the message files for clients from within *%SYBASE%\locales\language_name\charset_name*.

For example, the following entry specifies that the default locale uses us_english for the language and iso_1 for the character set:

```
locale = default, us_english, iso_1
```

How client applications use *locales.dat*

Client applications use the *locales.dat* file to identify the language and character set to use. The connection process follows these steps:

1 When a client application starts, it checks the operating system locale setting and then checks the *locales.date* file to see if that setting is appropriate for Adaptive Server. For example, a locale entry for French can look like the following:

```
locale = fr_FR, french, iso_1
```

2 When the client connects to Adaptive Server, the language and character set information is passed to Adaptive Server in the login record.

3 Adaptive Server then uses:

- The character set information, for example, iso_1, to identify the client's character set and verify whether it can convert character data to this character set
- The language (in the preceding example, French) and character set information to see if it has messages in the client's language

Note Adaptive Server software comes with some locale entries already defined in the *locales.dat* file. If these entries do not meet your needs, you can either modify them or add new locale entries.

Editing the *locales.dat* file

Before beginning the edit, make a copy of the original file, in case you have problems with the resulting edited version.

To edit the *locales.dat* file:

1 Open the *locales.dat* file copy in a text editor such as Notepad.

2 Find the section for Windows, which is enclosed in brackets *[NT]*.

3 Make sure the section contains an entry for the language (*syb_language*) and character set (*syb_charset*) combination that you want to use.

- If an entry does not exist, continue with step 4.
- If an entry does exist, continue with step 5.

Note The value for *platform_locale* must match the value required by your operating system. If the locales definitions in your system configuration files do not match the Sybase locale definitions, your applications will not run properly.

For example, if you want your Open Client messages to appear in French, and Adaptive Server is using the ROMAN8 character set, you would check the *locales.dat* entries for your platform and look for the following entry:

```
locale = fr_FR, french, roman8
```

4 Add the required entry or modify an existing entry.

5 Save the changes, if any, and exit the text editor.

Changing the localization configuration

By default, the Adaptive Server and Backup Server configurations use the English locale settings localization, which include:

- Character set definition files for Western European character sets
- Sort order definition files for Western European character sets
- us_english system message files

During the installation process and through reconfiguration, you can specify a different language, character set, and sort order.

❖ **Changing the localization configuration for Adaptive Server and Backup Server**

1. Start Server Config at Start | Programs | Sybase | Adaptive Server Enterprise | Server Config.
2. Click the icon for the server for which you want to change configuration, and click its corresponding Configure button in the Configure Sybase Servers dialog box.
3. From the Existing Servers screen, select the name of the server you want to configure, and click Continue.
4. Log in, if necessary.

   When configuring a Backup Server, you may already be logged in. If you are, go to step 6.

   When configuring an Adaptive Server, you need to log in first.

   1. Enter the login name and password of a user with System Administrator privileges, and click Continue.
   2. Click Yes if the Adaptive Server is not running. Server Config asks you to start it now.
5. Depending on the type of server you are configuring, go to “For Adaptive Server,” on page 79, or “For Backup Server” on page 81.

For Adaptive Server

Each language uses about 2MB of database space per module. If necessary, use the alter database command to increase the size of the master database before adding another language.

Note If you want to install more than one language on Adaptive Server, and the master database is not large enough to manage more than one language, the transaction log may become too full. You can expand the master database only on the master device. For more information, see the *System Administration Guide: Volume 2*.

❖ **Configuring localization for Adaptive Server on the server, start *syconfig***

1 Choose Language from the Configure Adaptive Server Enterprise dialog box.

Server Config displays the Language Options dialog box.

Note If you change the sort order or default character set, you must reconfigure existing databases to work with the new data requirements. See the *System Administration Guide: Volume 1* for more information.

2 To add or remove a language or character set, click the appropriate Add/Remove option.

Server Config displays the Install Languages or Install Character Sets dialog box, depending on your choice.

The languages and character sets displayed in the Selected list are already installed and available for Adaptive Server to use.

Server Config permits you to configure only those languages for which it finds message files. Since we do not provide message files for all possible languages, some languages cannot be installed using the Server Config utility. If your language does not appear as one of the available languages you must exit Server Config and follow the instructions in “Installing a new language module” on page 74.

To add or remove a language or character set:

1 Select a language or character set from the Available list, and click Add or Remove.

2 Click OK. The Configure Adaptive Server dialog box redisplays.

Note The Japanese language cannot coexist with any other installed language. If you install the Japanese language on Adaptive Server, you must make it the default language.

3 To change the default language, character set, or sort order, click the appropriate Set Default button in the Language Options dialog box.

   Server Config displays the appropriate Change dialog box:

   - Change Default Language
   - Change Default Character Set
   - Change Default Sort Order

   Adaptive Server supports only one sort order at a time, so the Sort Order heading provides only the Set Default option.

4 For languages or character sets:

   1 Select an option from the Available list, and click Add.

   2 Click OK.

   For sort orders:

   1 Select a sort order from the Available Sort Orders list.

   2 Click OK.

5 In the Language Options dialog box, click OK.

6 In the Configure Adaptive Server dialog box, choose Save to save the localization settings and return to the Configure Sybase Servers dialog box.

7 When you have completed the necessary configuration changes, click Exit to quit Server Config.

For Backup Server

When you select the Backup Server to configure, Server Config displays the Configure Backup Server dialog box.

❖ **Configuring localization for Backup Server**

1 From the Configure Backup Server dialog box, choose:

   1 The default language for Backup Server from the Language drop-down list

   2 The default character set from the Character Set drop-down list

2 Click Save to save the changes and return to the Configure Sybase Servers dialog box.

3 When you have completed the necessary configuration changes, click Exit to quit Server Config.

Configuring Adaptive Server for other character sets

To configure Adaptive Server with the character set and sort order for your language, complete the following steps. Your system messages appear in the default language, English.

1 Use the charset utility to load the default character set and sort order.

To use the charset, the server must be running and you must have System Administrator privileges. Use the *file name* of the sort order:

```
%SYBASE%\%SYBASE_ASE%\bin\charset -Usa -Ppassword -
Sserver_name sort_order_file character_set
```

Replace *sort_order_file* with the name of the sort order file. See Table 6-17 on page 83. Replace *character_set* with the Sybase name for your character set. See Table 6-18 on page 85.

2 Use the charset utility to load any additional character sets. See "charset utility" on page 86 for more about this utility.

If you plan to use the Adaptive Server built-in character set conversions, you must load the character set definition files for all the characters set on your client platforms. If you are using the Unilib character set conversions, you do not need to do this.

3 Using isql, log in to your server as "sa" and select the master database.

```
1> use master
2> go
```

4 Use the *ID* of the sort order to configure your server for the new character set and sort order.

```
1> sp_configure "default sort_order_id",
2> sort_order_id, "character_set"
3> go
```

Replace *sort_order_id* with the ID for your sort order. See Table 6-17 on page 83. Replace *character_set* with the Sybase name for your character set. See Table 6-18 on page 85.

5 Shut down the server to start the reconfiguration process.

6 Restart the server. Use Windows Service Manager from your Sybase Program Group or from a command prompt, invoke RUN_*server_name.bat* from *%SYBASE%\%SYBASE_ASE%\install*.

7 The server starts, rebuilds all the system indexes, then shuts down. Restart a second time to bring the server up in a stable state.

Sort orders

Table 6-17 describes the available sort orders. If your language does not appear, then there is no language-specific sort order for your language—use a binary sort order.

Table 6-17: Available sort orders

Language or script	Sort orders	File name	ID
All languages	Binary order	*binary.srt*	50
Cyrillic	Dictionary order, case-sensitive, accent-sensitive	*cyrdict.srt*	63
	Dictionary order, case-sensitive, accent-sensitive	*cyrnocs.srt*	64
English	Dictionary order, case-sensitive, accent-sensitive	*dictiona.srt*	51
French	Dictionary order, case-insensitive, accent-sensitive	*nocase.srt*	52
German	Dictionary order, case-insensitive, accent-sensitive, with preference	*nocasepr.srt*	53
These sort orders work with all Western European character sets.	Dictionary order, case-insensitive, accent-insensitive	*noaccent.srt*	54
English	Alternate dictionary order, case-sensitive	*altdict.srt*	45
French	Alternate dictionary order, case-sensitive, accent-insensitive	*altnoacc.srt*	39
German	Alternate dictionary order, case-sensitive, with preference	*altnocsp.srt*	46
These sort orders work only with CP 850.			
Greek	Dictionary order, case-sensitive, accent-sensitive	*elldict.srt*	65
This sort order works only with ISO 8859-7.			
Hungarian	Dictionary order, case-sensitive, accent-sensitive	*hundict.srt*	69
These sort orders work only with ISO 8859-2.	Dictionary order, case-insensitive, accent-sensitive	*hunnoac.srt*	70
	Dictionary order, case-insensitive, accent-insensitive	*hunnocs.srt*	71
Russian	Dictionary order, case-sensitive, accent-sensitive	*rusdict.srt*	58
This sort order works with all Cyrillic character sets except for CP 855.	Dictionary order, case-insensitive, accent-sensitive	*rusnocs.srt*	59

Language or script	Sort orders	File name	ID
Scandinavian	Dictionary order, case-sensitive, accent-sensitive	*scandict.srt*	47
These sort orders work only with CP 850.	Dictionary order, case-insensitive, with preference	*scannocp.srt*	48
Spanish	Dictionary order, case-sensitive, accent-sensitive	*espdict.srt*	55
	Dictionary order, case-insensitive, accent-sensitive	*espnocs.srt*	56
	Dictionary order, case-insensitive, accent-insensitive	*espnoac.srt*	57
Thai	Dictionary order	*dictionary.srt*	51
Turkish	Dictionary order, case-sensitive, accent-sensitive	*turdict.srt*	72
These sort orders work only with ISO 8859-9.	Dictionary order, case-insensitive, accent-insensitive	*turnoac.srt*	73
	Dictionary order, case-insensitive, accent-sensitive	*turnocs.srt*	74

Character sets

Table 6-18 lists the supported character sets and their Sybase name.

Table 6-18: Sybase character set names

Character sets	Sybase name
ASCII 8	acsii_8
Big 5	big5
CP 437	cp437
CP 850	cp850
CP 852	cp852
CP 855	cp855
CP 857	cp857
CP 860	cp860
CP 863	cp863
CP 864	cp864
CP 866	cp866
CP 869	cp869
CP 874	cp874
CP 932	cp932
CP 936	cp936
CP 950	cp950
CP 1250	cp1250
CP 1251	cp1251
CP 1252	cp1252
CP 1253	cp1253
CP 1254	cp1254
CP 1255	cp1255
CP 1256	cp1256
CP 1257	cp1257
CP 1258	cp1258
DEC Kanji	deckanji
EUC-CNS	euccns
EUC-GB	eucgb
EUC-JIS	eucjis
EUC-KSC	eucksc
GREEK8	greek8
ISO 8859-1	iso_1
ISO 8859-2	iso88592

Character sets	Sybase name
ISO 8859-5	iso88595
ISO 8859-6	iso88596
ISO 8859-7	iso88597
ISO 8859-8	iso88598
ISO 8859-9	iso88599
ISO 8859-15	iso885915
Koi8	koi8
Kazakhstan Cyrillic	kz1048
Macintosh Cyrillic	mac_cyr
Macintosh Central European	mac_ee
Macintosh Greek	macgrk2
Macintosh Roman	mac
Macintosh Turkish	macturk
ROMAN8	roman8
Shift-JIS	sjis
TIS 620	tis620
TURKISH8	turkish8
UTF-8	utf8

charset utility

Use the charset utility to load character sets and sort orders into Adaptive Server. If you are using charset to load the default character set and sort order, this should be done only at the time of installation.

To change the default character set and sort order of Adaptive Server, see the *System Administration Guide: Volume 1*.

Syntax

```
charset
[ -U username ]
[ -P password ]
[ -S server ]
[ -I interfaces ]
[ -v version ]
sort_order
[charset]
```

Table 6-19: Keywords and options for charsets

Keywords and options	Description
-U	If you are not already logged in to your operating system as "sa", you must specify `-Usa` in the command line.
-P	Specifies the "sa" password on the command line. If not specified, the user is prompted for the "sa" password.
-S	Specifies the name of the server. If not specified, charset uses the DSQUERY environment variable to identify the server name. If there is no DSQUERY environment variable, charset attempts to connect to a server named "SYBASE."
-I	Specifies the *interfaces* file to use. If not specified, charset uses the *interfaces* file in the *SYBASE* directory.
-v	Causes the Sybase version string to be printed, then exits. Use with no other options specified.
sort_order	When charset is used to load the default character set and sort order, *sort_order* is a mandatory parameter specifying the name of the sort order file to be used by Adaptive Server. When loading additional character sets, use *charset.loc* to indicate the name of the character set files.
charset	Specifies the directory of the character set to be used by Adaptive Server.

CHAPTER 7

Logging Error Messages and Events

This chapter describes how to use the error logging features of Adaptive Server for Windows.

Logging errors and events

Adaptive Server for Windows supports two types of message logging:

- Adaptive Server error logging
- Windows event logging

Adaptive Server error logging

Adaptive Server begins to write information to a local error log file, called the Adaptive Server error log each time Adaptive Server starts: *%SYBASE%\%SYBASE_ASE%\install\errorlog*

This error log file:

- Stores information about the success or failure of each start-up attempt.

- Logs error and informational messages generated by the server during its operations.
- Remains open until you stop the server process.
- Retains its contents until you rename, move, or empty the file.

If the error log become too large, you can:

- Dynamically change the path of the error log using sp_errorlog. Once the older errorlog is not being user by Adaptive Server, you can move it, and make space available.
- Stop the Adaptive Server and delete logged messages.

See "Diagnosing System Problems," in the *System Administraton Guide, Volume 1* for a description of the error log format.

Enabling and disabling error logging

Logging to the Adaptive Server Error Log is always enabled. However, when you create or modify a specific user-defined message, you can set it to be omitted from the log. See "Logging user-defined messages" on page 96.

Types of information logged

The Adaptive Server Error Log stores the following types of messages:

- Start-up messages from Adaptive Server
- Backtraces and stack traces from Adaptive Server
- Fatal error messages (severity level 19 and higher)
- Kernel error messages
- Informational messages

Windows event-logging

Adaptive Server also logs error messages in the Windows event log, if event logging is enabled.

Using the Windows event-logging feature, you can:

- Manage Adaptive Server error messages in the same way that you manage error messages for other Windows applications and services

- Set up a central event-logging site to store error messages from multiple Adaptive Servers

For information about centralized event logging, see "Using a central logging site" on page 99.

Setting up Windows event-logging for use by Adaptive Server

By default, Windows event logging of Adaptive Server messages is enabled, but you can disable it. You can also specify that logging of specific messages always be enabled.

For information about controlling logging of Adaptive Server messages to the Windows Event Log, see "Enabling and disabling Windows event logging" on page 94.

To make Windows event logging available to Adaptive Server, ensure that the following conditions are true in the Windows Event Log Settings box:

- The Overwrite Events as Needed option is selected
- The Maximum Log Size is set to at least 2048 bytes

Use the Windows Event Viewer to confirm or change these settings:

1 Choose Programs from the Start menu, choose Administrative Tools, and choose Event Viewer.

2 Choose Log Settings from the Log menu.

   The Event Log Settings dialog box appears. Make sure that the System Log is selected.

3 Change the Maximum Log Size to 2048, if necessary.

4 Click the Overwrite Events as Needed button to toggle the feature on or off.

5 Click OK.

6 Choose Exit from the Log menu.

Types of information logged

Adaptive Server for Windows logs the same messages in the Windows event log as in its Adaptive Server error log, with the exception of normal start-up messages. Some start-up messages are recorded only in the Windows event log, but all messages are logged in the local Adaptive Server Error Log.

Optionally, you can specify the recording of successful and unsuccessful logins to Adaptive Server in the Adaptive Server error log and the Windows event log. See“Logging user-defined messages” on page 96.

Managing the logs

Table 7-1 names the parameters, options, and system procedures for enabling and disabling event and error logging and indicates whether they affect the two logs. It also lists the pages in this chapter that contain instructions on using these elements to refine message logging.

Table 7-1: Methods for enabling/disabling error and event logging

Method	Affects event log	Affects error log	See page
error logging configuration parameter	Yes	No	95
event log computer name configuration parameter	Yes	No	95, 98
Server Config Event Logging option	Yes	No	95
Server Config Error Log Path option	No	Yes	93, 104
sp_altermessage system procedure	Yes	Yes	96
sp_addmessage system procedure	Yes	Yes	96
log audit logon success configuration parameter	Yes	Yes	97
log audit logon failure configuration parameter	Yes	Yes	97
xp_logevent system extended stored procedure	Yes	No	98

Setting error log paths

The installation program sets the error log location in the Sybase installation directory when you configure a new Adaptive Server. Backup Servers have their own error logs.

The default location for each server’s error log is:

- Adaptive Server: *%SYBASE%\%SYBASE_ASE%\install directory*
- Backup Server: *%SYBASE%\%SYBASE_ASE%\install directory*

At start-up, you can reset the name and location of the Adaptive Server Error Log file from the command line. Use the -e start-up parameter to set the name and location for the error log file.

To change the default error log path or file name:

- For Adaptive Server, see “Setting the Adaptive Server error log path” on page 93.
- For Backup Server, see “Setting the Backup Server error log path” on page 94.

Note Multiple Adaptive Servers cannot share the same error log. If you install multiple Adaptive Servers, specify a unique error log file name for each server.

Setting the Adaptive Server error log path

Use the Server Config utility to change the path:

1 Select Start | Programs | Sybase | Adaptive Server Enterprise | Server Config.

2 Click the Adaptive Server icon from the Products box in the Configure Sybase Servers dialog box.

3 Click the Configure Adaptive Server button in the Adaptive Server Enterprise box.

4 Select the name of the server to configure in the Existing Servers box, and click Continue.

5 Type the login name and password of an Adaptive Server user with System Administrator privileges in the Enter System Administrator Password dialog box.

6 Click Continue.

7 Click Yes if Adaptive Server is not running, and Server Config prompts you to start it.

8 Click the Error Log Path button in the Configure Adaptive Serve Enterprise dialog box.

Server Config displays the Error Log Installation Path dialog box.

9 Type the full path name to an error log file that is not on a network drive, and click OK.

10 In the Configure Adaptive Server dialog box, click Save to save the new error log setting.

11 Click Exit to quit Server Config.

Setting the Backup Server error log path

Use the Server Config utility to change the path:

1 Select Start | Programs | Sybase | Adaptive Server Enterprise | Server Config.

2 Click the Backup Server icon from the Products box in the Configure Sybase Servers dialog box.

3 Click the Configure Backup Serve button in the Backup Server box.

4 Select the name of the server to configure in the Existing Servers box, and click Continue.

5 Type the full path name to an error log file that is not on a network drive in the Configure Backup Server dialog box.

6 Click Save to save the new error log setting.

7 Click Exit to quit Server Config.

Enabling and disabling Windows event logging

By default, Adaptive Server enables the logging of its messages to the Windows event log at start-up. This section explains how to disable and enable logging of Adaptive Server messages to the Windows event log.

There are two ways to control event logging:

- Using Server Config
- Using sp_configure

Using Server Config

Use the Server Config utility to control event logging:

1 Select Start | Programs | Sybase | Adaptive Server Enterprise | Server Config.

2 Click the Adaptive Server icon, and click the Configure Adaptive Server button.

3 Select the name of the server to configure in the Existing Servers dialog box, and click Continue.

4 Type the login name and password of an Adaptive Server user with System Administrator privileges in the Enter System Administrator Password dialog box.

5 Click Continue.

6 Click Yes if the Adaptive Server is not running, and Server Config asks you to start it now.

7 Click Event Logging in the Configure Adaptive Server Enterprise dialog box.

Server Config displays the Event Logging dialog box.

8 Click the Use Windows Event Logging button to enable or disable Adaptive Server error message logging to the Windows Event Log.

9 In the Event Log Computer Name text box:

- To send messages to a remote computer log, type its name.
- To send messages to a local computer log, let the value default to LocalSystem.

10 Click OK.

11 Click Save to save your changes in the Configure Adaptive Server dialog box.

12 Click Exit to quit Server Config.

Using *sp_configure*

You can enable Adaptive Server message storage in the Windows event log by using sp_configure to set the event logging configuration parameter. Possible values are:

- 1 – to enable logging of Adaptive Server messages

```
sp_configure "event logging", 1
```

- 0 – to disable logging of Adaptive Server messages

```
sp_configure "event logging", 0
```

Note Restart Adaptive Server after enabling logging with sp_configure; disabling does not require a server restart.

For information about the event logging configuration parameter and sp_configure in general, see the *System Administration Guide: Volume 1*.

Managing messages

When event logging is enabled, you can manage its functions in the following ways:

- Use sp_addmessage to add a user message, or sp_altermessage to control whether a specific message is logged in both the Adaptive Server error log and in the Windows event log.

  For the complete syntax for the sp_addmessage and sp_altermessage system procedures, see the *Reference Manual: Procedures*.

- Use configuration parameters to specify whether auditing events are logged. Auditing events pertain to a user's success, log audit logon success, or failure, log audit logon failure, in logging in to Adaptive Server.

- Use the xp_logevent extended stored procedure to set up logging of user-defined events in the Windows event log in Adaptive Server.

Logging user-defined messages

You can specify whether a user-defined message is logged to the Adaptive Server error log as well as to the Windows event log. Adaptive Server lets you make this determination for:

- New messages (sp_addmessage)
- Existing messages (sp_altermessage)

For more information about these commands and their parameters, see sp_addmessage and sp_altermessage in the *Reference Manual: Procedures*.

New messages

Include the with_log option in sp_addmessage when you add a new user-defined message to sysusermessages. This parameter sets the Adaptive Server to log the message each time that the message appears.

Existing messages

Include the with_log option in sp_altermessage to change an existing user-defined message. This parameter alters the reporting status of that message:

- TRUE – to enable logging
- FALSE – to disable logging

Logging auditing events

By default, Adaptive Server does not log auditing events. However, you can use sp_configure parameters to specify whether Adaptive Server is to log auditing events, such as logins, to the Adaptive Server error log and to the Windows event log.

Possible parameters and values are:

- log audit logon success at 1 – to enable logging of successful Adaptive Server logins:

  ```
  sp_configure "log audit logon success", 1
  ```

- log audit logon failure at 1 – to enable logging of unsuccessful Adaptive Server logins:

  ```
  sp_configure "log audit logon failure", 1
  ```

- Either parameter at 0 – to disable logging of that message type:

  ```
  sp_configure "log audit logon success", 0
  sp_configure "log audit logon failure", 0
  ```

For more information about sp_configure, see the *System Administration Guide: Volume 1*.

Logging user-defined events

You can arrange to have user-defined events logged to the Windows event log from within Adaptive Server. For example, you can create a “database imported” event that is generated after a database has been imported successfully.

Using the xp_logevent extended stored procedure (ESP), you can arrange to log the event. This ESP allows you to specify the following:

- The message that is to appear in the event description field of the event viewer when the event is logged
- Whether the event should be characterized as informational, warning, or error

For more information, see xp_logevent in the *Reference Manual: Procedures*.

Using a remote log

By default, if event logging is enabled, Adaptive Server logs messages to the Windows event log on the local computer system.

To change the destination computer for logging messages:

1 Set event log computer name with sp_configure on the local computer. Use either:

   - sp_configure, as in the following command line, or:

     ```
     sp_configure "event log computer name", 0, user1
     ```

   - Enter the name of the target computer in the Event Log Computer Name box on the Event Logging dialog box.

   To display the name box, see “Using Server Config” on page 94.

2 Start the server from a Domain Administrators account.

   a Choose Start/Settings/Control Panel/Services.

   b Select the remote server to use from the list.

   c Click Startup.

   d Click This Account in the Log On As box.

   e Open the drop-down list to display the Add Users dialog box.

f Double-click an account name that is in the Domain Administrators group, and click OK.

g Click OK at the Service dialog box.

h Click Start to exit the utility and enable the server.

Regardless of how you specify the destination computer, be sure that it is configured to record Adaptive Server error messages. To configure the destination computer, see “Using a central logging site” on page 99.

Using a central logging site

You can record messages from multiple Adaptive Servers in the Windows event log of a central, network computer. The recording computer does not need to run Adaptive Server.

Figure 7-1 illustrates a central logging site.

Figure 7-1: Diagram of a central logging site

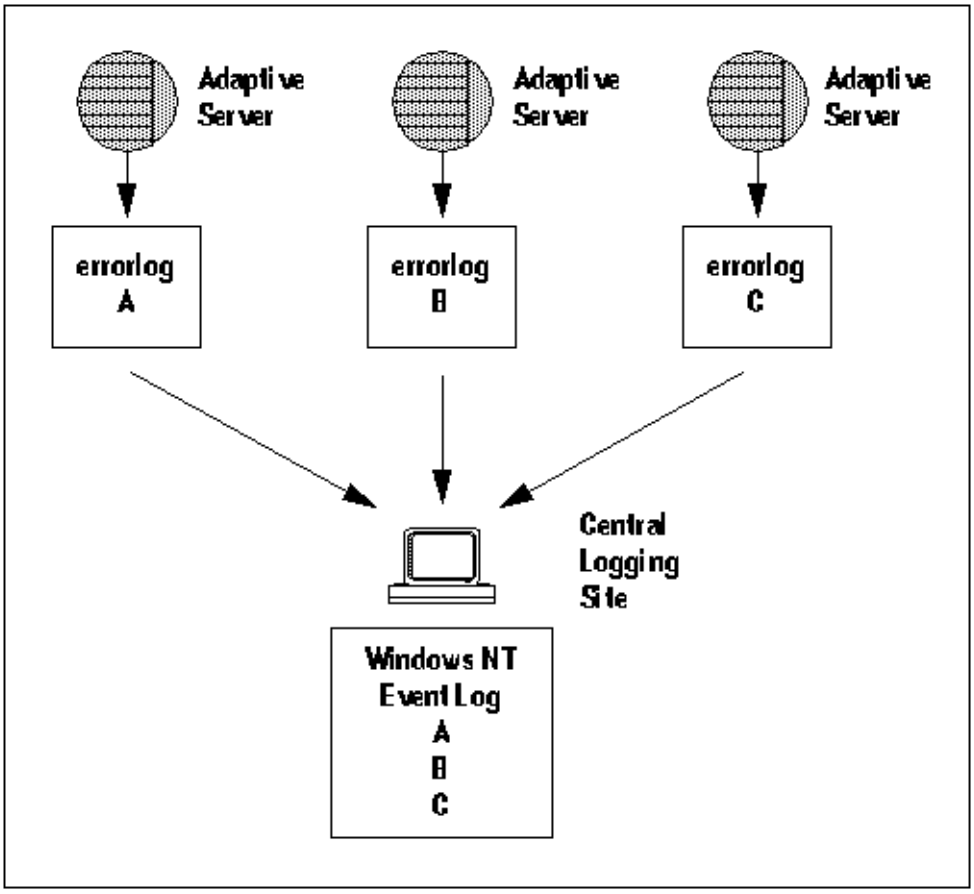

Using a central logging site provides added flexibility in managing multiple Adaptive Servers. For example:

- A System Administrator can monitor the status of all Adaptive Servers on the network by examining the central event log.
- Users of individual Adaptive Servers can view either the local Adaptive Server error log file or the central event logging site to examine error messages.

Logging messages from multiple Adaptive Servers

To log messages from multiple Adaptive Servers, the central logging computer must have:

- Access to the *sybevent.dll* file
- A Registry key for each Adaptive Server that will log messages on the central computer
- A set of four key values that define each Registry key for Adaptive Server

Setting up a local central logging site

An event-logging computer uses a Registry key to define each message-sending Adaptive Server and is unable to log messages from servers for which it has no key.

To set up a computer as a central logging site, you must create and define a Registry key for each Adaptive Server that is to log messages into the site.

To create and define a Registry key

Use the *sybevent.dll* file and the regedt32 utility.

To create and define a Registry key:

1 Log in to Windows using an account with Windows administrator privileges.

2 Copy the *sybevent.dll* file from an Adaptive Server machine if it does not exist on the local computer.

The *sybevent.dll* file is stored in the *dll* subdirectory of the Sybase installation directory (*\sybase\dll*, by default). The actual location of *sybevent.dll* on the logging computer is not important, however, you must record a fixed location for the file in the Windows Registry.

Note You can use the same *sybevent.dll* file on the event-logging computer, as long as all Adaptive Servers are at the same Version level; for example, 11.5.1.

3 Start the Windows regedt32 utility.

For instructions and screens on using this utility with Adaptive Server, see "Increasing Windows Sockets connections" on page 38.

4 Complete the steps in "Creating a Registry key" on page 102 to create a key for a single Adaptive Server.

5 Complete the steps in "Defining a Registry key" on page 102 to define the key that you just created.

6 Repeat steps 4 and 5 for each Adaptive Server that is to send messages to the logging site computer.

❖ **Creating a Registry key**

1 In the regedt32 utility, select the Registry window named HKEY_LOCAL_MACHINE.

2 Open the levels until you reach the Registry key named:

HKEY_LOCAL_MACHINE\SYSTEM\
CurrentControlSet\Services\EventLog\Application

3 From the Edit menu, choose Add Key to display the Add Key dialog box.

4 Complete the dialog box as follows:

Key Name – type the name of the Adaptive Server computer that is to store the messages at the central logging site.

Class – leave this box blank. You do not have to specify a class for the new key.

5 Verify that you have entered the new Registry key correctly.

6 Click OK.

7 Complete the steps in "Defining a Registry key" on page 102 to define the key that you just created.

❖ **Defining a Registry key**

1 In the regedt32 utility, open the Registry key that you just created.

2 From the Edit menu, choose Add Value.

3 Type an event-logging value name as shown in Table 7-2 for the new Registry key. Enter the value name exactly as it is shown in the table; value names are case-sensitive.

Table 7-2: Registry values for a central logging PC

Value name	Datatype	String	Notes
CategoryCount	REG_DWORD	0x6	Do not change the data value. Make sure the string value is hexadecimal (Hex).
CategoryMessageFile	REG_SZ	*%SYBASE%\%SYBASE_ASE%\dll directory*	Replace *%SYBASE%\%SYBASE_ASE%\dll directory* with the path to the *sybevent.dll* file.
EventMessageFile	REG_SZ	*%SYBASE%\%SYBASE_ASE%\dll directory*	Replace *%SYBASE%\%SYBASE_ASE%\dll directory* with the path to the *sybevent.dll* file.
TypesSupported	REG_DWORD	0xff	Do not change the data value. Make sure the string value is hexadecimal (Hex).

Note Be sure to enter the correct path to the *sybevent.dll* file for the CategoryMessageFile and EventMessageFile values.

4 Select the data type for the value as named in Table 7-2 from the drop-down list.

5 Verify that you have entered the new key value and datatype correctly, and click OK.

6 Type the data or string in the Editor box, and click OK.

7 Repeat steps 2–6 for the remaining three values in each Registry key.

8 To create another key begin again with“Creating a Registry key” on page 102.

9 Once you have created a Registry key for each Adaptive Server, choose Exit from the Registry menu in the Registry Editor dialog box to quit regedt32.

Viewing the messages

You need the Windows Event Viewer and a text editor to display the error messages and events that have been logged.

In the Windows event log

Use the Windows Event Viewer in the Administrative Tools group.

To examine Adaptive Server messages recorded in the Windows event log:

1 Choose Programs from the Start menu, choose Administrative Tools, and choose Event Viewer.

   The Viewer displays a list of Adaptive Server messages.

2 Double-click on a message to display its Event Detail dialog box.

   The Description list box defines the Adaptive Server message number as a number and text.

In the Adaptive Server error log

Use a text editor, such as Notepad, on the logging computer to open the file and view the messages in the Adaptive Server error log.

If you cannot find the error log file:

1 Choose Start/Programs/Sybase/Server Config.

2 Click the Adaptive Server icon, and click the Configure Adaptive Server button.

3 Select the name of the server whose error log you want to examine from the Existing Servers dialog box, and click Continue.

4 Type the login name and password of an Adaptive Server user with System Administrator privileges in the Enter System Administrator Password dialog box.

5 Click Continue.

6 Click Yes if the Adaptive Server is not running, and Server Config asks you to start it now.

7 Click Error Log Path from the Configure Adaptive Server dialog box.

   Server Config displays the Error Log Installation Path dialog box, which supplies the current path to the error log.

For detailed information on interpreting the information in the error log, see the *System Administration Guide: Volume 1*.

CHAPTER 8

Using Security Services with Windows LAN Manager

This chapter describes how to use Adaptive Server security services with the Windows LAN Manager to authenticate users and provide data integrity.

Security services with Windows LAN Manager

When using Adaptive Server on Windows, you can enable the security services provided by Windows LAN Manager to authenticate users, clients, and servers to one another.

Figure 8-1 shows a client application that is using LAN Manager to ensure a secure connection with Adaptive Server.

Figure 8-1: Establishing secure connections between LAN Manager and Adaptive Server

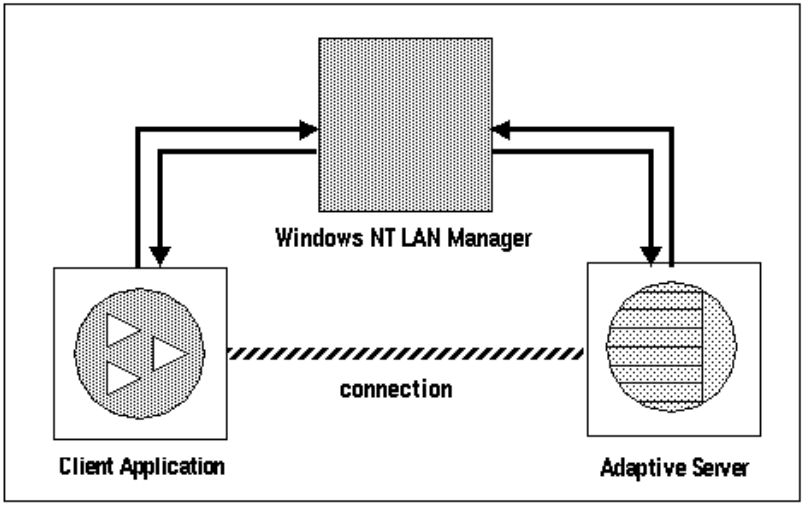

The secure connection between LAN Manager and a server can be used to provide a unified login to Adaptive Server. Through this login, the LAN Manager authenticates users *once* and does not require them to supply a name and password every time they log into Adaptive Server.

The secure connection also can support one or more of the following security services:

- Message integrity to verify that data communications have not been modified
- Replay detection to verify that data has not been intercepted by an intruder
- Out-of-sequence check to verify the order of data communications

How login authentication works

When a client requests authentication services, the following steps occur:

1 The client validates the login with LAN Manager. LAN Manager returns a *credential*, which contains security-relevant information.

2 The client sends the credential to Adaptive Server and informs Adaptive Server that it wants to establish a secure connection.

3 Adaptive Server authenticates the client's credential with LAN Manager.

When the credential is valid, Adaptive Server establishes a secure connection between itself and the client.

Administering security services using LAN Manager

Table 8-1 describes a process for using Adaptive Server's unified login capability with LAN Manager.

Warning! Adaptive Server must be installed before completing the steps in Table 8-1.

Table 8-1: Process for administering network-based security

Step	Description	See
1. Set up the configuration files: • *libtcl.cfg* • *sql.ini*	Use a text editor to modify the *libtcl.cfg* file. Use dsedit to specify security mechanisms in the *sql.ini* file or a Directory Service.	"Modifying configuration files for a unified login" on page 108
2. Make sure the security administrator for LAN Manager has created logins for each user and for the Adaptive Server and Backup Server.	The security administrator for LAN Manager must add names and passwords for users and servers.	"Identifying users and servers to LAN Manager" on page 111, and your Windows documentation
3. Configure security for the installation.	Use sp_configure to enable the use of security services.	"Configuring Adaptive Server for LAN Manager security" on page 111
4. Restart Adaptive Server.	Activates the use security services parameter.	"Initiating the new security services" on page 116
5. Add logins to Adaptive Server to support enterprise-wide login.	Use sp_addlogin to add users. Optionally, specify a default secure login with sp_configure.	"Adding logins to support unified login" on page 117
6. Connect to the server.	Use isql with the -V option or use Open Client Client-Library to connect to Adaptive Server and specify the security services to use. **Note** If you use the isql utility, you do not have to supply a user name or password.	"Defining the connection to a server for security services" on page 118 *Open Client/Server Configuration Guide for Desktop Platforms* "Security Features" topics page in the *Open Client Client-Library Reference Manual*

Modifying configuration files for a unified login

Configuration files are created during installation at a default location in the Sybase directory structure. Table 8-2 provides an overview of the configuration files required for LAN Manager to use unified login and security services.

Table 8-2: Names and locations for configuration files

File name	Description	Directory
libtcl.cfg	This driver configuration file contains information pertaining to directory, security, and network drivers and any required initialization information.	*%SYBASE%\ini*
objectid.dat	This object identifiers file maps global object identifiers, such as the LAN Manager, to local names for character set, collating sequence, and security mechanisms.	*%SYBASE%\ini*
sql.ini	The *sql.ini* file contains connection and security information for each server that it lists.	*%SYBASE%\ini*

For a detailed description of the configuration files, see the *Open Client/Server Configuration Guide for Desktop Platforms*.

Setting up drivers for network-based security

The *libtcl.cfg* file stores information about the following driver types:

- Network (Net-Library)
- Directory Services
- Security

A **driver** is a Sybase library that provides an interface to an external service provider. Adaptive Server dynamically loads drivers so you can change the driver used by an application without relinking the application.

Entries for network drivers

The syntax for a network driver entry in the *libtcl.cfg* file is:

driver=protocol description

where:

- *driver* is the name of the network driver.
- *protocol* is the name of the network protocol.

- *description* is a description of the entry. This element is optional.

Note You can comment out the network driver entry by placing a semicolon at the beginning of the line; then, Adaptive Server will use a driver that is compatible with your application and platform.

Entries for Directory Services

Entries for Directory Services apply if you want to use a Directory Service instead of the *sql.ini* file.

For information about directory entries, see "Sharing network configuration information" on page 43.

Warning! Client applications bundled with Adaptive Server require a *sql.ini* file for effective processing. Eliminating this file with a Directory Service may limit Adaptive Server functionality.

Entries for security drivers

The syntax for a security driver entry in the *libtcl.cfg* file is:

provider=driver

where:

- *provider* is the local name for the security mechanism. *objectid.dat* defines the mapping of the local name to a global object identifier. The default local name for Windows LAN Manager on Windows and Windows 95 (for clients only) is "LIBSMSSP".

  Note If you use a provider name other than the default, you must also change the local name in the *objectid.dat* file. For an example, see "Checking the LAN Manager's local name" on page 110.

- *driver* is the name of the security driver. The Windows LAN Manager driver is named "LIBSMSSP." The default location of all drivers is *%SYBASE%\%SYBASE_OCS%\dll*.

Editing the *libtcl.cfg* file

Use the ocscfg utility to edit the *libtcl.cfg* file. This utility displays the file's contents in a dialog box with section headings in the form of tabs for easy perusal.

For information on using the ocscfg utility, see the *Open Client/Server Configuration Guide for Desktop Platforms*.

The following text is a sample *libtcl.cfg* file for desktop platforms:

```
[NT_DIRECTORY]
 ntreg_dsa=LIBDREG  ditbase=software\sybase\serverdsa
[DRIVERS]
 NLWNSCK=TCP  Winsock TCP/IP Net-Lib driver
 NLMSNMP=NAMEPIPE  Named Pipe Net-Lib driver
 NLNWLINK=SPX  NT NWLINK SPX/IPX Net-Lib driver
 NLDECNET=DECNET  DecNET Net-Lib driver
[SECURITY]
 NTLM=LIBSMSSP
```

Checking the LAN Manager's local name

The *objectid.dat* file maps global object identifiers to local names.

Note You need to change this file only if you have changed the local name of the LAN Manager in the *libtcl.cfg* file.

The file contains sections such as [CHARSET] for character sets and [SECMECH] for security services. Of interest here is the security section.

The following example is a security section excerpt from the *objectid.dat* file:

```
[secmech]
      1.3.6.1.4.1.897.4.6.3    = NTLM
```

You can specify only one local name for LAN Manager. Use any text editor to edit this file.

Warning! Do not change the "1.3.6.1.4.1.897.4.6.3" identification.

Specifying security information for Adaptive Server

You can use the *sql.ini* file or a Directory Service to provide information about the servers in your installation.

To use either the *sql.ini* file or a Directory Service, run the dsedit utility. This utility provides a graphical user interface for specifying server attributes such as the server version, name, and security mechanism.

For information about using dsedit, see the *Open Client/Server Configuration Guide for Desktop Platforms*.

For more information about using directory services with Adaptive Server on Windows, see “Sharing network configuration information” on page 43.

Identifying users and servers to LAN Manager

The security administrator for LAN Manager must define *principals* (defined users) to the security mechanism. Use LAN Manager’s User Manager utility to identify all users for the system.

You do not need to enter the Adaptive Server name as a principal to LAN Manager. However, the Windows user account that you use to start Adaptive Server must be defined as a valid principal to LAN Manager. For example, to use an Windows account named “servadmin” to start Adaptive Server, you must define “servadmin” as a principal to LAN Manager.

This rule applies whether you start Adaptive Server through Sybase Central or as an Windows service. See the *Installation Guide*.

For detailed information about the User Manager utility, see your Windows documentation.

Configuring Adaptive Server for LAN Manager security

Adaptive Server uses several configuration parameters to administer unified login and security services through LAN Manager. To set these parameters, you must be a System Security Officer.

All parameters for unified login and security through LAN Manager are part of the "Security-Related" configuration parameter group. Use the configuration parameters to:

- Enable the use of external security services (LAN Manager)
- Require unified login
- Require one or more message integrity security services

Enabling and disabling external security services

To reset the status of LAN Manager security services, use sp_configure with the use security services configuration parameter:

- 1 – to enable services with LAN Manager.
- 0 – the default, to disable network-based security services.

The syntax is:

```
sp_configure "use security services", [0|1]
```

For example, to enable services with LAN Manager, execute:

```
sp_configure "use security services", 1
```

Managing unified login

You can use configuration parameters to:

- Require unified login
- Establish a default secure login

Because all the parameters for unified login are dynamic, they take effect as soon as you change them. You must be a System Security Officer to set the parameters.

Requiring unified login

The unified login required configuration parameter controls the type of login that is acceptable to Adaptive Server. The possible values are:

- 1 – to require all users who request a connection to Adaptive Server to be authenticated by LAN Manager.

- 0 – the default, to let Adaptive Server accept both traditional login names and passwords and authenticated credentials.

The syntax is:

sp_configure "unified login required", [0|1]

For example, to require all logins to be authenticated by a security mechanism, execute:

```
sp_configure "unified login required", 1
```

Establishing a secure default login

When a user with a valid credential from LAN Manager logs in to Adaptive Server, the server checks to see whether the name is listed as a user in master..syslogins. If it is, Adaptive Server accepts that user name.

For example, a user logs in to LAN Manager as “ralph”, and “ralph” is listed in master..syslogins. Adaptive Server uses all roles and authorizations as defined for “ralph” on that server.

As an alternative example, a user with a valid credential logs in to Adaptive Server, but is unknown to the server. Adaptive Server accepts the login only when a *secure default login* has been defined with sp_configure. Adaptive Server uses the default login for any user who is not defined in master.syslogins, but who is validated by LAN Manager.

To set up a secure login, use the following syntax:

sp_configure "secure default login", 0, *login_name*

where *login_name* is a user name. The default value for the secure default login parameter is “guest”.

The login used for this parameter must be a valid login in master..syslogins. For example, to set the login “gen_auth” to be the default login.

1 Use sp_addlogin to add the login as a valid user in Adaptive Server:

```
sp_addlogin gen_auth, pwgenau
```

This procedure sets the initial password to “pwgenau”.

2 Use sp_configure to designate the login as the security default:

```
sp_configure "secure default login", 0, gen_auth
```

Adaptive Server then uses this login for a user who, although validated by LAN Manager, is unknown to Adaptive Server.

Note This user does not have a unique identity in Adaptive Server. That is, more than one user can assume the suid (system user ID) associated with the secure default login. You might want to activate auditing for all activities of the default login. Instead of using the secure default login, consider using sp_addlogin to add all users to the server.

For more information about adding logins, see "Adding logins to support unified login" on page 117.

Mapping LAN Manager login names to server names

All login names in Adaptive Server must be valid identifiers. However, external security mechanisms, such as LAN Manager, may allow login names that are not valid in Adaptive Server.

For example, login names that are longer than 30 characters or that contain special characters such as !, %, *, and & are invalid names in Adaptive Server.

Table 8-3: Conversion of invalid characters in login names

Invalid character	Converts to
Ampersand &	Underscore _
Apostrophe '	
Backslash \	
Colon :	
Comma ,	
Equals sign =	
Left single quotation mark ‘	
Percent sign%	
Right angle bracket >	
Right single quotation mark ’	
Tilde ~	

Invalid character	Converts to
Caret ^	Dollar sign $
Curly brackets { }	
Exclamation point !	
Left angle bracket <	
Parentheses ( )	
Period .	
Question mark ?	
Asterisk *	Pound sign #
Minus sign -	
Pipe |	
Plus sign +	
Quotation marks " "	
Semicolon ;	
Slash /	
Square brackets [ ]	

For more information about identifiers, see the *Reference Manual: Blocks*.

Requiring data integrity check

You can use the following configuration parameters with LAN Manager. These parameters cause Adaptive Server to check one or more types of data integrity for all messages.

- msg integrity reqd – set this parameter to 1 to force a check for general tampering in all messages.

  If the parameter is set to 0 (the default), message integrity is not required. However, the client can establish this check if the security mechanism supports it.

- msg out-of-seq checks reqd – set this parameter to 1 to force a check for sequence changes in all messages.

  If the parameter is set to 0 (the default), sequence checking is not required. However, the client can establish this check if the security mechanism supports it.

- msg replay detection reqd – set this parameter to 1 to force a check for replay or interception in all messages.

If the parameter is set to 0 (the default), replay detection is not required. However, the client can establish this check if the security mechanism supports it.

Ensuring adequate memory for security services

The value of the total memory configuration parameter specifies the number of 2K blocks of memory that Adaptive Server requires at start-up. To make sure that there is sufficient memory when using unified login and security services through LAN Manager, allocate approximately 6K of additional memory per connection.

For example, if the maximum number of unified logins that occur at the same time is expected to be 150, increase the total memory parameter by 450. This increase expands memory allocation by 450 2K blocks.

The syntax is:

sp_configure total memory, *value*

where *value* is the sum of the current memory and the memory you are adding.

For example, to supply Adaptive Server with 25,000 2K blocks of memory, including the increased memory for network-based security, enter:

```
sp_configure total memory, 25000
```

The minimum requirement for this parameter is specific to the operating system.

For information about estimating and specifying memory requirements for Adaptive Server, see the *System Administration Guide: Volume 2*.

Initiating the new security services

Changes to the security services are static. You must restart Adaptive Server to activate the security services.

For instructions on starting and stopping Adaptive Server, see the *Installation Guide*.

Adding logins to support unified login

When a user logs in to Adaptive Server with an authenticated credential, Adaptive Server follows these steps, as needed:

1 Checks that user is a valid user in master..syslogins.

- If the user name appears, Adaptive Server accepts the login without requiring a password.
- If the user name does not appear, Adaptive Server performs step 2.

2 Checks that a default secure login is defined in master..syslogins.

- A default login definition allows the user to log in successfully.
- The absence of a default login definition causes Adaptive Server to reject the login.

Therefore, consider whether to allow only users who are defined as valid logins to use Adaptive Server or to allow any user with the default login to use Adaptive Server.

Note You must add the default login in master..syslogins and use sp_configure to define the default. For more information, see "Establishing a secure default login" on page 113.

General procedure for adding logins

To add logins to the server and, optionally, to add users with appropriate roles and authorization to one or more databases, follow the general procedure described in Table 8-4.

Table 8-4: Adding logins and authorizing database access

Task	Required role	Command or procedure	See
1. Add a login for the user.	System Security Officer	sp_addlogin	*Security Administration Guide*
2. Add the user to one or more databases.	System Security Officer, System Administrator, or Database Owner	sp_adduser Enter this procedure from within the database.	*Security Administration Guide*

Task	Required role	Command or procedure	See
3. Add the user to a group in a database.	System Security Officer, System Administrator, or Database Owner	sp_changegroup Enter this procedure from within the database.	*Security Administration Guide* *Reference Manual: Procedures*
4. Grant system roles to the user.	System Administrator or System Security Officer	grant role	*Security Administration Guide* *Reference Manual: Commands*
5. Create user-defined roles and grant the roles to users.	System Security Officer	create role grant role	*Security Administration Guide* *Reference Manual: Commands*
6. Grant access to database objects.	Database object owner	grant [select | insert | delete | update| references | execute]	*Security Administration Guide*

Defining the connection to a server for security services

Use the following options to define an Adaptive Server for network-based security services such as Windows LAN Manager through the isql and bcp utilities:

- -R *remote_server_principal* – to specify the principal name for Adaptive Server.
- -V *security_options* – to specify network-based user authentication.
- -Z *security_mechanism* – to specify the name assigned to LAN Manager.

For more information about Adaptive Server utilities, see the *Utility Guide* for your platform.

Specifying the principal name

Use -R *remote_server_principal* to specify the principal name for the server as defined for LAN Manager.

By default, a server's principal name matches the server's network name, which is specified by either the -S option or the DSQUERY environment variable. You must use the -R option when the server's principal name and network name are not the same.

Specifying network-based user authentication

Use -V *security_options* to specify network-based user authentication.

With this option, the user must log in to Windows LAN Manager before running the utility. In this case, if a user specifies the -U option, the user must supply the network user name known to the security mechanism, and any password supplied with the -P option is ignored.

-V can be followed by a *security_options* string of key-letter options to enable additional security services. The key letters are:

- i – to enable data integrity service. This option verifies that data communications have not been modified.
- r – to enable data replay detection. This option verifies that data has not been intercepted by an intruder.
- q – to enable out-of-sequence detection. This option verifies the order of data communications.

You can specify additional security options by including them immediately following the -V option. For example, to use isql with network-based user authentication, replay detection, and out-of-sequence detection, enter:

```
isql -Vrq
```

Specifying the name assigned to LAN Manager

The -Z *security_mechanism* specifies the name assigned to LAN Manager in the *libtcl.cfg* configuration file; "LIBSMSSP", by default.

When the line does not supply a *security_mechanism* name, the command uses the default mechanism.

For more information about security mechanism names, see the *Open Client/Server Configuration Guide for Desktop Platforms*.

Note When you log in to LAN Manager and then log in to Adaptive Server, you do not need to specify the -U (user) option on the utility because Adaptive Server gets the user name from LAN Manager.

Determining the status of security services

To determine whether security services are enabled for the current session, use show_sec_services. In the following example, the results indicate that unified login is enabled, and, therefore, so are the security services:

```
select show_sec_services()
 go
-----------------------------------------------------
unifiedlogin
(1 row affected)
```

Configuration parameters used in security services

This section summarizes the configuration parameters that the unified login and security services use through LAN Manager. These parameters provide the following security checks:

- msg integrity reqd – to check data integrity.
- msg out-of-seq checks reqd – to check message sequence.
- msg replay detection reqd – to detect interception or replay.
- secure default login – to specify a default login.
- unified login required – to control user authentication.

For general information on configuration parameters, see the *System Administration Guide: Volume 1*.

Checking data integrity

Summary information	
Name in pre-11.0 version	N/A
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System Security Officer

The msg integrity reqd parameter controls the checking of all messages to ensure data integrity. The use security services parameter must be set to 1 (enabled) for message integrity checks to occur.

Checking message sequence

Summary information	
Name in pre-11.0 version	N/A
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate
Required role	System Security Officer

The msg out-of-seq checks reqd parameter controls the checking of all messages to ensure that the sequence is correct. The use security services parameter must be set to 1 (enabled) for sequence checks to occur.

Detecting interception or replay

Summary information	
Name in pre-11.0 version	N/A
Default value	0 (off)
Range of values	0 (off), 1 (on)
Status	Dynamic
Display level	Intermediate

Summary information	
Required role	System Security Officer

The msg replay detection reqd parameter controls the checking of all messages to detect whether they have been intercepted (detect replay). The use security services parameter must be set to 1 (enabled) for replay detection checks to occur.

Specifying a login

Summary information	
Name in pre-11.0 version	N/A
Default value	0
Range of values	0 (followed by another parameter naming the default login)
Status	Dynamic
Display level	Intermediate
Required role	System Security Officer

The secure default login parameter specifies a default login for all users who are preauthenticated, but do not have a login in master..syslogins.

Use the following syntax to establish the secure default login:

sp_configure "secure default login", 0, *default_login_name*

where *default_login_name* is the name of the default login for a user who, although unknown to Adaptive Server, has already been authenticated by a security mechanism. This name must be a valid login in master..syslogins.

For example, to specify “dlogin” as the secure default login, execute:

```
select sp_configure "secure default login", 0,
    dlogin
```

Controlling user authentication

Summary information	
Name in pre-11.0 version	N/A
Default value	0
Range of values	0, 1

Summary information	
Status	Dynamic
Display level	Intermediate
Required role	System Security Officer

The unified login required parameter controls authentication of all users who log into Adaptive Server by means of a security mechanism. The use security services parameter must be set to 1 (enabled) to use the unified login security service.

Managing login security on an Windows computer

This section discusses how to use the login security features of Adaptive Server for Windows.

For more information on system security, see the *Security Administration Guide*.

Overview of security features

You can use Adaptive Server security features alone or in combination with the Windows security features.

Adaptive Server security

As a standalone product, Adaptive Server ensures security by:

- Storing login information for all database users in the master.dbo.syslogins table. Passwords stored are encrypted.
- Requiring client applications to specify the login name and password of a database user, either programmatically or with a command-line option.
- Checking the user name and password against the information in syslogins, and accepting or rejecting the login accordingly.

Combined Adaptive Server and Windows login security

Adaptive Server increases security by integrating the default Adaptive Server login process with Windows security features. The resulting integrated security modes add the following conveniences for users:

- Authorized users do not have to maintain separate login passwords for Adaptive Server and Windows.
- System Administrators can take advantage of Windows security features such as encrypted passwords, password aging, domain-wide user accounts, and Windows-based user and group administration.

Trusted connections and combined login security

Combined login security operates only over network protocols that support authenticated connections between clients and servers. Such connections are referred to as *trusted connections*.

Trusted connections are limited to client applications that access Adaptive Server by using the Named Pipes protocol.

Note Other network protocols, such as TCP/IP sockets and IPX/SPX, do not support authenticated connections, so clients on these protocols are handled according to the standard Adaptive Server login mechanism.

A System Administrator must use sp_grantlogin to assign permissions to Windows users and groups. Using sp_grantlogin, the System Administrator has the following additional options:

- Assigning one or more Adaptive Server roles to Windows users and groups
- Designating that the user or group should receive the default database object permissions assigned by the grant command

If the System Administrator does not use sp_grantlogin to assign user or group permissions, users cannot log in through trusted connections. For more information, see “Permitting trusted connections” on page 127.

Note Adaptive Server does not permit trusted connections for Windows users named “sa.” The user name “sa” is reserved for the default Adaptive Server System Administrator account.

Understanding login security modes

Adaptive Server provides the following modes for configuring login security:

- Standard
- Integrated
- Mixed

Standard mode

When operating in Standard mode, Adaptive Server manages its own login validation process for all connections by:

- Ignoring the Windows network user name and checking the supplied Adaptive Server user name and password against the information in the syslogins table
- Providing valid users with Adaptive Server connections and allowing valid users to receive the permissions and roles that were assigned to them with the grant command

For a description of the login security features of Adaptive Server, see the *Security Administration Guide*.

Integrated mode

When operating in Integrated mode, Adaptive Server uses Windows-based authentication mechanisms for all connections by:

- Allowing only trusted connections, using Named Pipes, to connect to Adaptive Server
- Ignoring any Adaptive Server login name and password that is submitted in the login request. Instead, it checks the mapped Windows network user name against the information in the syslogins table.

  If no matching login name exists, and the login process includes a default user name, Adaptive Server substitutes the default login name, for example, “guest”, to complete the connection. For more information, see “Default login” on page 128.
- Providing authorized users, when they log in, with permissions and roles as described in “Permitting trusted connections” on page 127.

- Following the Windows Domain structure for the use of computers. Windows must authenticate each user, either through trust relationships or through explicitly assigned permissions on each server.

Note If you bypass the Windows login security for Adaptive Server authentication, that is, if you opt for Adaptive Server security only, it does not matter to which user or group you assign the computers. The only requirement is that the protocol you use allows the client and server to communicate.

Mixed mode

When operating in Mixed mode, Adaptive Server allows both trusted, as with Named Pipes, and "untrusted" connections. It first examines the requested login name as specified by the client application, then handles the login depending on the information supplied.

Adaptive Server processes the login:

- When the login name matches the mapped network user name, is null, or is composed of spaces, Adaptive Server treats the login attempt as a trusted connection and uses the rules for Integrated mode.
- When the user supplies a different login name, Adaptive Server treats the login attempt as an untrusted connection and uses the rules for Standard mode.

Mixed mode offers users the convenience of login security integration without forcing all clients and applications to use that integration.

- Existing applications that embed a hard-coded login name and password for all users continue to operate as before.
- Other operating system clients, such as Apple Macintosh clients and UNIX-based workstations, also can access an Adaptive Server in Mixed mode.
- Users accessing Adaptive Server over trusted connections can avoid a separate Adaptive Server password validation by omitting the user name and password in their login request.

Note Applications can be designed to send an empty login name field in the connection request, thereby avoiding a separate login step.

Managing the login security features

Use the following elements to manage login security in Integrated or Mixed mode:

- Trusted connections
- Windows Registry parameters

Permitting trusted connections

When operating under Integrated or Mixed Login Mode, Adaptive Server assigns permissions to trusted user connections by checking the user's network or Windows group name. This check determines whether the Security Administrator, using sp_grantlogin, has assigned an Adaptive Server role, or the default value, to that name, and Adaptive Server acts accordingly.

- When no permissions were assigned to the name, and Adaptive Server is operating in:
  - Integrated mode, Adaptive Server refuses the connection.
  - Mixed mode, Adaptive Server treats the connection as an untrusted connection. Then, the login process continues under the Standard mode rules.
- When one or more Adaptive Server roles have been assigned to the user's network name or to the user's Windows group, the user receives those roles and permissions that were assigned by the Security Administrator through the grant statement.
- When only the default value has been assigned to the user's network name or Windows group, the user receives only the permissions and roles that were assigned by the Security Administrator through the grant statement.

The most important point to remember is that Windows users or their associated Windows groups must have permissions that were assigned with sp_grantlogin.

For examples of this system procedure, see "Assigning trusted connection permissions" on page 131.

For more information about sp_grantlogin, see the *Security Administration Guide*.

Windows Registry parameters

When you install Adaptive Server and other Sybase products on your computer, the installation program configures several parameters to help you to manage the login security features while in Integrated or Mixed mode.

This sections describes the following management parameters:

- Default login
- Default domain
- SetHostName
- Character mappings

To modify the parameter values, see "Modifying the parameter values" on page 130.

Default login

Adaptive Server uses the default login parameter to specify the Adaptive Server login name that an authorized user can enter when a network user name does not appear in the syslogins table. Standard mode does not use this value.

When there is no value for default login, Adaptive Server denies access to users who do not have a network user name in syslogins.

Default domain

Adaptive Server uses the default domain parameter to specify the Windows or LAN Manager domain name for matching network user names to Adaptive Server login names.

Because two different domains can define the same network user name for two different users, the following rules apply:

- Adaptive Server can authorize access to both distinct users, but it must be able to distinguish between the two names in the login process for a trusted connection.
- For user names defined in domains other than the parameter's default value, Adaptive Server adds the domain name and a domain separator, a backslash character (\), to the network user name before looking for the user name in the syslogins table.

For example, the domain MARKETING is the Adaptive Server default definition, and two different users employ the network user name "john", one in the MARKETING domain and the other in the ENGINEERING domain.

- John in MARKETING accesses Adaptive Server with the login name of "john" over a trusted connection.
- John in ENGINEERING accesses the same Adaptive Server with a login name of "ENGINEERING\john" to which his name was mapped before the software looked it up in syslogins.

- When your server computer participates in a specific domain, set the default domain parameter to that domain name. Otherwise, set default domain to the server's computer name.

SetHostName

The SetHostName parameter determines whether the host name from the client login record is replaced with the Windows network user name for users under integrated security mode.

- 1 (enabled) – to include the network user name in the results of the sp_who system procedure.
- 0 (disabled) – the default, to omit the network user name from the results of the sp_who system procedure.

To modify the SetHostName value, which is located in the following Registry path: *HKEY_LOCAL_MACHINE\SOFTWARE\Sybase\ Server\server_name*, you must use the regedt32 utility.

For general information about regedt32, see your Windows operating system documentation.

Character mappings

Certain characters that are valid for Windows user names are not valid for Adaptive Server login user names. Such characters include the following:

- Domain separator (\)
- Space ( )
- Hyphen (-)
- Period (.)
- Single quotation mark (')

- Exclamation point (!)
- Percent sign (%)
- Caret (^)
- Ampersand (&)

Character mapping lets you determine how these invalid characters can be converted into characters that are valid for Adaptive Server.

For example, the Windows user name “t-johns” contains a dash character (-), which is invalid in Adaptive Server. You can map the dash character to a valid “at” sign (@) to make the user name compatible with Adaptive Server, as “t@john”. The mapping stores the dash as an “at” sign, but displays it as a dash.

When you first install Adaptive Server, the installation program maps a few invalid characters to the valid characters that are listed in Table 8-5.

Table 8-5: Default mapping values

Invalid character	Valid mapped character
Domain separator (\)	Underscore (_)
Hyphen (-)	Pound sign (#)
Space ( )	Dollar sign ($)

Modifying the parameter values

To modify the values for the default login, default domain, and SetHostName parameters, use one of the following utilities:

Note You can change the SetHostName value only through regedt32.

- Use the Server Config utility to modify the value only for Adaptive Server.

  For general steps on using Server Config, see “Changing login security options” on page 135.

- Use the regedt32 utility to change the value directly for use throughout your Windows operating system.

  For steps on using regedt32 to affect your operating system, see your Windows operating system documentation.

Administering login security using system procedures

You can administer integrated security under Windows in the following ways:

- Assign trusted connection permissions – sp_grantlogin
- Display Adaptive Server integrated login configuration– sp_loginconfig
- Display permissions and user names – sp_logininfo
- Revoke permissions – sp_revokelogin

For the full syntax for these procedures, see the procedure names in the *Reference Manual: Procedures*.

Assigning trusted connection permissions

To assign permissions to Windows users and groups that access Adaptive Server over trusted connections:

- Use sp_grantlogin when Adaptive Server is running under Integrated mode or Mixed mode, and the connection is Named Pipes.
- Use the grant command when Adaptive Server is running under Standard mode or Mixed mode with a connection other than Named Pipes.

The sp_grantlogin permissions can include either one or more Adaptive Server roles or just the default parameter. This parameter indicates that Adaptive Server provides the user with the default permissions as assigned by the grant command.

To use the sp_grantlogin, grant, and default parameters in an example:

1 To assign the System Administrator and System Security Officer roles to all members of the Windows group named Administrators, enter:

```
sp_grantlogin "Administrators", "sa_role sso_role"
```

2 Then, to assign “select” permissions on the sales table to the Windows user, “hasani”, enter:

```
sp_grantlogin "hasani", "default"
grant select on sales to hasani
```

Note If you do not specify a role or a value with sp_grantlogin, the procedure automatically assigns the default value.

Displaying the current Registry values

To display the current settings for the Registry values, use sp_loginconfig as discussed under “Windows Registry parameters” on page 128.

For example, executing sp_loginconfig on a newly installed Adaptive Server displays a list similar to the following:

```
name                    config_item
 ---------------------- ----------------------
 login mode             standard
 default account        NULL
 default domain         EAST
 set host               false
 key _                  domain separator
 key $                  space
 key @                  space
 key #                  -
```

Displaying permissions and user names

To display the current permissions and mapped user names for both Windows users and groups, use sp_logininfo. The following list describes this sample display:

```
account name             mapped login name
          type                  privilege
          -------------- ---------------------------------------------
BUILTIN\Administrators   BUILTIN\Administrators
          group                 'sa_role sso_role oper_role'
WEST\chantal             WEST_chantal
          user                  'default'
EAST\chantal             chantal
          user                    'sa_role'
```

- Three roles were assigned to the Windows administrators group: sa_role, sso_role, and oper_role.
  - The group names are prefaced by “BUILTIN\” to indicate that the entry refers to a built-in Windows group (a default group on all servers), rather than a group that is created by the user.
  - The domain separator in a group name is not mapped to a valid Adaptive Server character.

You do not need to add a login or grant further permissions to an Windows group, but you do need to add a login for each user in that group.

- The first Windows user, named “chantal”, has the default parameter assigned as a permission. “chantal” is a member of the WEST domain, and her mapped Adaptive Server login name is “WEST_chantal”.

  “WEST_chantal” is the name the System Administrator should use when assigning an Adaptive Server login name and permissions to this user.

- The second Windows user, also named “chantal”, logs in from the EAST domain. Her mapped user name is simply “chantal”, since EAST has been set as Adaptive Server’s default domain (see the second item in this list).

To change or revoke the displayed users, groups, and permissions use the sp_grantlogin and sp_revokelogin procedures.

Revoking permissions granted with *sp_grantlogin*

To revoke permissions that were granted with sp_grantlogin use either:

- The sp_revokelogin command when Adaptive Server is running under Integrated Security mode or under Mixed mode, and the connection is Named Pipes.
- The revoke command when Adaptive Server is running under Standard mode or under Mixed mode, and the connection is other than Named Pipe.

The following command line revokes all permissions from the Windows group named Administrators:

```
sp_revokelogin Administrators
```

Configuring login security

This section provides general guidelines and suggestions for configuring Adaptive Server login security. Although you can complete the configuration tasks in a variety of ways, it is easiest to follow the steps in the order shown.

Create Windows users and groups

To create the user accounts and user groups that will access Adaptive Server over trusted connections, run the User Manager from the Administrative Tools (Common) menu. To access this menu, choose Start/Programs.

Keep the following guidelines in mind when creating groups and users:

- Make sure that Windows users and groups exist *before* you assign permissions to them in Adaptive Server.
- Create the accounts with a user name other than "sa".

  Note Some functions that were assigned to the "sa" user in earlier versions of Adaptive Server are now divided between the sa_role and sso_role. You may want to assign both roles to Adaptive Server system administrators to provide the same permission level on an upgraded system. For more information, see the *System Administration Guide: Volume 1*.

- Begin planning the permission levels you want to assign to the users and groups.

  Although it may seem intuitive to grant the sa_role to the Windows Administrators group, the choice ultimately depends on the security requirements for your site.

When using integrated security features for the first time, consider restricting the permission level to a small group of Windows users. After you become more experienced with administering integrated security, you can expand the permission levels to include Windows groups.

Configure mapping and default domain values

To set the mapping and default domain options, follow the instructions under "Changing login security options" on page 135.

Configure these values *before* adding accounts to Adaptive Server in step 4, as these values affect the format of entries in syslogins.

For example, a user named "joseph" in the WEST domain is to log in to Adaptive Server over a trusted connection. If you set the Map_ value to the domain separator (\) and the default domain value to NULL, the name "WEST_joseph" must appear in the syslogins table. However, if you later change the default domain value to WEST, the login name "joseph" would need to be in syslogins instead of "WEST_joseph".

Set login security mode

To set the security mode to either Integrated or Mixed, follow the instructions under "Changing login security options" on page 135.

When using login security features for the first time, consider using Mixed mode. If, for some reason, you cannot connect over a trusted connection, Mixed mode allows you to log in to Adaptive Server using standard Adaptive Server user names and passwords, such as the user name "sa".

Add network login names to *syslogins*

To add a login name for each Windows user who will access Adaptive Server over a trusted connection, use sp_addlogin. Include any nondefault domain names and the correct mapping characters in the login name.

If you are not sure what to enter as the login name, experiment with a sample user to clarify your options:

1 Use sp_grantlogin to assign a role to a sample user on the network.

2 Enter sp_logininfo to determine what the format of entries in syslogins should look like.

3 Use the entries listed in the mapped login name column as templates for the login names you create with sp_addlogin.

Assign roles

To assign roles or "default" permissions to Windows users or groups, use sp_grantlogin. When performing this step, keep in mind that assigning permissions to Windows groups generally provides more flexibility than assigning permissions to individual users.

After you have configured several groups with the correct permissions, you can use the User Manager to manage individual user's access to Adaptive Server.

Changing login security options

When you install a new Adaptive Server, the installation program sets it to operate in Standard mode. Use the Server Config tool to change the following settings:

- The login security mode (Standard, Integrated, or Mixed)
- The name of the default login account
- The name of the default domain

❖ **Selecting a login security mode**

1. Log into Windows using an account with Windows administrator privileges
2. Start the Server Config utility
3. Complete the initial steps to configure Adaptive Server.

   For instructions, see"Starting Server Config for Adaptive Server" on page 24.
4. Click the Login Security button in the Configure Adaptive Server Enterprise dialog box.
5. Continue with "To enable Standard login security mode" or "For integrated or mixed login security mode," depending on the login mode.

❖ **Enabling Standard login security mode**

1. Click the Standard option button to display Standard Current Login Security Mode box, then click OK.
2. Click Save in the Configure Adaptive Server dialog box.
3. Click Exit to quit Server Config.

❖ **Enabling Integrated or Mixed login security mode**

1. Click the Integrated option button to display Integrated in the Current Login Security Mode box, and click Continue
2. Set the login security mode:

   For Integrated Mode, click the Automatic Login for Trusted Connections (Named Pipes) Only option.

   For Mixed mode, click the Trusted First and Adaptive Server Login for Excluded (i.e., Netware, TCPIP) option.
3. Enter the values to use as defaults:

   In the default login box, type the name of the default user account to use for trusted connections. Adaptive Server uses this value when it cannot locate the user name in syslogins. For more information, see "Default login" on page 128.

   In the default domain box, type either the default domain name or the workstation's network name. For more information, see "Default domain" on page 128.

4 Click the Map Characters button to configure Adaptive Server mappings under an Integrated security mode.

   Server Config displays the Character Mapping dialog box.

5 Use the drop-down lists to select the invalid character to be mapped to each valid Adaptive Server character.

   For more information, see “Character mappings” on page 129.

6 Click OK to save the character mapping configuration and return to the Integrated Login Options dialog box.

7 Click OK in the Integrated Login Options dialog box.

8 Click OK in the Login Security Options dialog box.

9 Click Save in the Configure Adaptive Server dialog box to save the new configuration.

10 Click Exit to quit Server Config.

CHAPTER 9

Using E-mail with Adaptive Server

Adaptive Server can send and receive e-mail messages through Sybmail, the Sybase messaging facility, and can take advantage of Windows Mail. This chapter provides instructions for using and configuring Sybmail to work with Windows Mail.

Sybmail messages

Adaptive Server for Windows can send, receive, and process e-mail messages.You can set Adaptive Server to manage these messages by using:

- A set of extended stored procedures (ESPs) that the user must run manually, or
- A system procedure that invokes the ESPs automatically by using procedural language code, rather than Transact-SQL statements.

Sending messages

Messages from Adaptive Server (outgoing messages) can be one of two types:

- Text
- Formatted query results

Adaptive Server's capability for e-mail greatly increases the potential usefulness of a stored procedure or trigger. For example:

- A user-defined stored procedure that registers a new employee in the company database can include commands that send e-mail messages to a new employee and to other departments that need to be aware of the new hire, such as facilities, human resources, and training.
- A trigger on an inventory table can send an e-mail message to inform the purchasing department that an item needs to be reorder when an update causes the number of items on hand to fall below a certain level.
- A weekly report generated from a database query can be produced automatically and sent to a mailing list.

Receiving messages

Adaptive Server's ability to process incoming mail allows users to send queries and receive results using e-mail, rather than a traditional client/server connection.

Sybmail flexibility allows a user to send queries to Adaptive Server from any computer, and, at a later time, to check e-mail for the results from either the same or a different computer.

Preparing Windows Mail for Sybmail

Sybmail takes advantage of the Windows Mail facility, so you need to prepare the Windows Mail system before you can use Sybmail. You must:

1 Connect to a post office.

2 Create a mailbox.

3 Create a mail profile for Adaptive Server.

The following sections provide a general outline for setting up Adaptive Server in the Windows Mail system.

For detailed instructions on working with Mail on your system, see your Windows operating system documentation or the *Microsoft Windows Resource Kit*.

Connecting to a post office

An Windows post office holds messages until all of the recipients have retrieved them.

The computer that is running Adaptive Server must have access to an Windows post office on the network. You can:

- Create a new post office, if one does not exist for your domain, or
- Connect to an existing workgroup post office.

When connecting to an existing post office, be prepared to supply its path.

Creating a mailbox for Adaptive Server

After connecting to a post office, create a mailbox for Adaptive Server in the destination post office.

Note Only the Windows post office administrator can add a new mailbox.

Be sure to supply a mailbox name and password for the mailbox.

- The password will be useful later when you establish a Sybmail user account on Adaptive Server.

  Make sure that the password meets the requirements for Adaptive Server passwords:

  - Must be at least 6 bytes.
  - Must be enclosed in quotation marks if the password does not begin with an alphabetic character.
- The mailbox name creates the association between the mailbox and the Adaptive Server mail profile that you will create in the next step.

Creating a mail profile for Adaptive Server

After you have added a mailbox for Adaptive Server, use the mailbox information to create a mail profile that is associated with the mailbox.

Note Each mail profile is associated with a single mailbox, although a single mailbox may be associated with several mail profiles.

The mail profile must have a password and be associated with a mailbox name.

- The password must be the same as Adaptive Server's mailbox password.
- The mailbox name must be the same as the mailbox name specified when you created the mailbox for Adaptive Server.

In the Mail Login Properties window, make sure the check box labeled "When logging on, automatically enter password" is selected (checked).

Creating an Adaptive Server login for Sybmail

After setting up an Adaptive Server profile in Windows Mail, create a login for Sybmail on Adaptive Server. When creating this user account, make sure that the following conditions are true:

- The *loginame* parameter is "sybmail".
- The *fullname* parameter has the same value as the Profile Name for Adaptive Server's mail profile.

  Adaptive Server uses this value as its MailUserName.
- The *password* parameter has the same value as the password for the mailbox that is associated with the server's mail profile.

  This value becomes Adaptive Server's MailPassword.

These values are the defaults for starting up an Adaptive Server mail session with the extended stored procedure xp_startmail, as discussed in "Managing a mail session" on page 144.

You can use either of the following methods to create a login for Adaptive Server:

- sp_addlogin from isql:

```
sp_addlogin "sybmail", "wrtyzz2c", @fullname="sqlserver"
```

or

- The Add Login facility in Sybase Central or Adaptive Server Manager.

Figure 9-1 summarizes the relationships between the values that you supplied to prepare an account for Sybmail.

Figure 9-1: User-defined values relationships in Sybmail

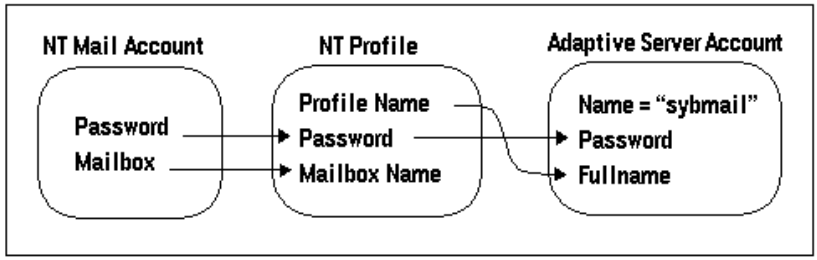

Sybmail and extended stored procedures

Adaptive Server uses XP Server, an Open Server application, to execute all of its extended stored procedures (ESPs), including the system ESPs that implement Sybmail.

By default, XP Server configuration uses the System Account (LocalSystem) as its start-up account. However, to use Sybmail, you must configure XP Server to start under a user account.

❖ **Configuring XP Server for a user account**

1 Start the Server Config tool.

2 Complete the initial steps to configure Adaptive Server.

For instructions, see "Configuring Adaptive Server" on page 25.

3 Click the Configure Default XP Server button in the Configure Adaptive Server Enterprise dialog box.

4 Click This Account to enable the option, and type a valid Windows user account and password for the server. Make sure that the account has the right to log in as a service.

If you do not have an existing user account with the right to log in as a service, you can grant a user this right from the Windows User Manager:

1. Open User Manager from the Administrative Tools (Common) menu in the Start menu.
2. Select the User name to act as the service.
3. Choose User Rights from the Policies menu.
4. In the User Rights Policy dialog box, select the Show Advanced User Rights check box.
5. In the Right drop-down list, select "Log on as a service", and click OK.
6. Exit the User Manager.

5. Click OK.
6. Click Save in the Configure Adaptive Server Enterprise dialog box.
7. Click Exit to quit Server Config.

Managing a mail session

You must initiate an Adaptive Server mail session before any messages can be sent or received.

Note Only one Sybmail session at a time can be running on an Adaptive Server.

Starting a session

When Adaptive Server starts a session, the mail user is represented by the MailUserName and the MailPassword that you defined when you created the Adaptive Server login for Sybmail.

You can initiate an Adaptive Server mail session in one of two ways:

- Call the xp_startmail extended stored procedure explicitly each time you start Adaptive Server.

You can override the previously mentioned login default by passing another user name and password to xp_startmail. You might want to do this if there are multiple profiles associated with Adaptive Server's mailbox, and you want to use an alternative profile.

- Arrange to start a mail session automatically when Adaptive Server starts up.

  For automatic start-up of an Adaptive Server mail session for subsequent Adaptive Server sessions, set the start mail session configuration parameter to 1.

  With the automatic start-up, you do not need to use xp_startmail to begin a mail session the next time that you start Adaptive Server.

Starting Sybmail without parameters

You can start Sybmail with xp_startmail and no parameters (default configuration), but only in the following situations:

- The Sybmail user account exists and the Start mail session parameter was configured to 1 when Adaptive Server was started, or
- The Sybmail user account exists, and you enter the following command to automatically start Sybmail:

```
sp_configure "start mail session", 1
```

In both of these situations, do not restart XP Server before issuing the command to start Sybmail with its default configuration. Once you restart XP Server, it drops the default settings.

Stopping a mail session

A mail session stops automatically when Adaptive Server shuts down. You also can explicitly stop an Adaptive Server mail session at any time with the xp_stopmail ESP.

For syntax and parameters for xp_startmail and xp_stopmail, see the *Reference Manual: Procedures*.

Note Stop the current Adaptive Server mail session with xp_stopmail before using xp_startmail to start another mail session for a different profile name. Until you stop the first session, the second session cannot access resources that are considered to be still in use by the first session.

Stored and extended procedures for handling messages

Table 9-1 summarizes the procedures that are available for processing e-mail for Adaptive Server.

Table 9-1: Procedures for processing mail

Procedure	Description
xp_deletemail	Deletes a message from the Adaptive Server message inbox.
xp_findnextmsg	Retrieves the message identifier of the next message in the Adaptive Server message inbox.
xp_readmail	Reads a message from the Adaptive Server message inbox.
xp_sendmail	Sends a message from Adaptive Server.
xp_startmail	Starts an Adaptive Server mail session.
xp_stopmail	Stops an Adaptive Server mail session.
sp_processmail	Reads, executes, responds to, and deletes messages submitted to Adaptive Server by e-mail.

Sending messages

An outgoing message can consist of text or the formatted results of a query or batch of queries. You can send a message directly through isql from either a stored procedure or a trigger that uses the xp_sendmail ESP.

Keep the following concepts in mind when managing outgoing messages:

- To send query results, input the query, or a stored procedure containing the query, to xp_sendmail. The query results are sent to the recipients of the message.

- When the message consists of query results, you can specify whether you want the results to be sent in the body of the e-mail message or as an attachment.
- When the message consists of text, use the *message* parameter to xp_sendmail.
- When the message consists of the results of a query, use the *query* parameter, and pass the quoted text of the query or the quoted execute command with its stored procedure name.

For syntax and parameters for xp_sendmail, see the *Reference Manual: Procedure*.

Text messages

The trigger in the following example sends e-mail to "purchasing" when an update causes the number of items on hand (onhand) in an inventory table (part) to fall below a certain level (min_onhand).

```
1> create trigger reorder
2> on part
3> for update as
4> if update(onhand)
5> if (select onhand - min_onhand
6> from inserted <= 0
7> begin
8> execute xp_sendmail
9> @subject="Inventory Notice"
10> @recipient="purchasing"
11> @message="Parts need to be reordered."
12> end
```

Query result messages

In response to the e-mail message generated by the trigger listed in the previous examples, the purchasing department can send the Adaptive Server mailbox a query to determine which parts should be reordered.

Note For a diagram of the process, see Figure 9-1 on page 143.

Adaptive Server then reads the query into a variable, named *received_mess*, as in the following example, with xp_readmail, and uses xp_sendmail to execute it and return the results:

```
declare @received_mess varchar(255)
execute xp_sendmail @recipient = "purchasing"
@query = @received_mess, @dbname = "inventory"
@dbuser ="sa"
```

Another example of mailing query results, a user-defined stored procedure, named usp_salesreport, in the salesdb database, is run at the end of the month to report on monthly sales activity. By invoking this procedure inside a call to xp_sendmail, you can automatically send the results of the procedure through e-mail to a mail group.

The following example sends the results of the usp_salesreport stored procedure as an attachment to an e-mail message addressed to “sales”, with copies to “mitchell” and “hasani”. The procedure is executed in the salesdb database with the privileges of the database owner of salesdb.

```
execute xp_sendmail @recipient = "sales",
@copy_recipient = "mitchell"; "hasani",
@subject = "Monthly Sales Report",
@query = "execute usp_salesreport",
@attach_result = true,
@dbname = "salesdb",
@dbuser = "dbo"
```

Receiving messages

Adaptive Server expects incoming e-mail messages to be in the form of Transact-SQL statements. Incoming mail can consist of a single statement or a batch of statements, delimited by an end-of-batch indicator.

Note Messages containing multiple statements must follow the rules for batches, as described in the *Transact-SQL Users Guide*.

Sybmail provides ESPs to process incoming messages, including the following:

- xp_findnextmsg
- xp_readmail

- xp_deletemail

These ESPs are briefly described below. For syntax and parameters, see the *Reference Manual: Procedures*.

Finding the next message

xp_findnextmsg returns the message identifier of the next message in the Adaptive Server inbox. Use the unread_only parameter to specify the messages for consideration:

- true – to consider only unread messages.
- false – to consider all messages.

You need the message identifier that is returned by xp_findnextmsg to pass to subsequent procedures that read and delete messages.

Reading a specific message

You can read a specific message by passing its message identifier to xp_readmail.

To read the first message in the inbox, or the first unread message, depending upon the unread_only parameter, do not specify a message identifier.

xp_readmail places the contents of the message in its *message* output parameter.

Other output parameters that store the remaining attributes of the message include *originator* (message sender), *date_received* (message received date), *subject* (message subject), and *recipients* (message addressees).

Deleting a message

After reading Adaptive Server's mail with xp_readmail, you can remove the message from Adaptive Server's inbox by passing the message identifier to xp_deletemail.

If you do not specify a message identifier, xp_deletemail deletes the first message in the inbox.

Processing incoming mail

You can process Adaptive Server's incoming e-mail queries manually by:

1 Calling the ESPs xp_findnextmsg, xp_readmail, and xp_deletemail individually for each message

2 Using xp_sendmail to execute the query in each message and send the e-mail results back to the requestor

However, it is much easier to use sp_processmail, which invokes these ESPs automatically.

sp_processmail reads and responds to the unread messages in the Adaptive Server inbox. You can determine which messages to process by passing a value for the *originator* parameter and/or the *subject* parameter, as shown in Table 9-2.

Table 9-2: Selecting messages by sender or subject

When you specify	sp_processmail processes
originator	Only mail from the specified sender
subject	Only mail with the specified subject header
originator and *subject*	Only mail by the specified sender with the specified subject header
Neither *originator* nor *subject*	The unread mail in the inbox

sp_processmail uses default parameters when invoking xp_sendmail, but you can override the *dbname*, *dbuser*, and *separator* defaults by passing these values to sp_processmail. For the syntax for sp_processmail and xp_sendmail, see the *Reference Manual: Procedures*.

The following example processes all the unread mail sent to Adaptive Server by the e-mail sender "admin":

```
sp_processmail @originator = "admin",
@dbuser = "sa", @dbname = "db1"
```

The procedure executes the queries in the db1 database in the System Administrator's context and returns the results an e-mail attachment to "admin" and to all the copied and blind-copied recipients of the original incoming message.

Using Sybmail security

To prevent unauthorized users from accessing privileged Adaptive Server data through e-mail, you must set:

- The execution privileges on the ESPs that process mail
- The security context for executing queries

Use the xp_sendmail or sp_processmail procedures to set these values.

Setting execution privileges

The ESPs that process mail, such as xp_findnextmsg, xp_readmail, xp_sendmail, and xp_deletemail, are database objects owned by the System Administrator.

Limit execution permission of these procedures to users with the sa_role or to a very small group of users to prevent unauthorized users from accessing Sybmail to execute queries that they would normally not be able to execute.

Setting the execution context

When you use xp_sendmail to execute a query that has been submitted by e-mail, the procedure causes Adaptive Server to execute the query with the privileges of a particular Adaptive Server login in a particular database. This login/database combination is the *execution context.* By default, the login is "sybmail" and the database is master.

You can set the execution context for individual messages by passing different login and database values to xp_sendmail or sp_processmail with the following optional variables:

- *dbuser* – to reset the login name.

  The login must represent a valid Adaptive Server account on the target Adaptive Server.

- *dbname* – to reset the database name.

The following sections describe the execution context when the procedure specifies one, both, or neither of the optional variables.

Naming both the user and the database

Specify both *dbuser* and *dbname* to control how Adaptive Server executes the query. These variables can affect the process:

- In the user context of the specified login when that login is a valid user in the specified database
- In the user context of "guest" when the login is not a valid user in the specified database

When the specified database is a system database, a "guest" account always exists. However, when the specified database is a user database, the database owner must have ensured that:

- The entity represented by the *dbuser* login is a valid database user, or
- There is a "guest" user in the database that can map to any login and execute queries with minimal permissions.

Naming the user but not the database

Specify only *dbuser* to name a user but cause Adaptive Server to execute the command, xp_sendmail or sp_processmail, in the master database.

When the login specified by *dbuser* is not a valid user in the master database, Adaptive Server executes the query in the user context of "guest".

Naming the database but not the user

Specify only *dbname* to set the default *dbuser* as "sybmail" and to cause Adaptive Server to execute any query under the user context of "guest".

When the specified database is a system database, a "guest" account always exists. However, when the specified database is a user database, the database owner must have ensured that there is a "guest" user in the database that can map to any login and execute queries with minimal permissions.

Naming neither the user nor the database

Specify neither parameter to retain the default *dbuser* as "sybmail" and the default database as master. Adaptive Server executes the e-mail query as "guest" in the master database.

CHAPTER 10

Managing Adaptive Server Databases

The administration of Adaptive Server databases includes both routine tasks and performance and tuning considerations.

- The *System Administration Guide: Volume 1* and *Volume 2* discusses most of the administrative tasks in detail.
- The *Performance and Tuning Series* provides in-depth explanations of performance issues.

This chapter discusses some of the tasks described in these books that may require different handling for Windows.

Managing database devices

The term **database device** refers to a disk or a portion of a disk that stores Adaptive Server databases and database objects.

Device requirements

The size and number of Adaptive Server devices depend on the following constraints:

- The maximum device size is 4TB.
- Each database can have up to 2G - 1devices.

- The maximum database size is 4 – 32 TB (dependent upon page size.)

Although some operating systems can designate an entire hard disk to use as a database device, Windows accepts only an operating system file (*.dat* file) as a database device.

When you install Adaptive Server, the program creates a *.dat* file in the *\data* directory of the Sybase installation directory. To use a *.dat* file as a database device, you can either use the default *d:\sybase\data* directory or create a device and a directory in which to store it.

Creating *.dat* files for database devices

If you choose to create a new device, use the disk init command to specify the drive, path, and file name of the database device.

Warning! Do not place Adaptive Server devices on network drives, as this causes unpredictable system behavior. Also, if your Adaptive Server uses a network drive, you cannot start the server as an automatic Windows service.

❖ **Create a database device using the file d:\devices\user1.dat**

1 If the *d:\data* directory does not exist, create it from the Windows command-prompt:

```
d:\> mkdir data
```

2 Start isql and connect to Adaptive Server using the "sa" account:

```
d:\sybase\bin> isql -Usa -Ppassword -Sserver_name
```

3 Create the device using a disk init statement similar to the following example:

```
1> disk init
2> name = "user_device1",
3>physname = "d:\data\user1.dat",
4>size = 4M
5> go
```

This example creates a 4MB device, without an actual device number. To use a specific number, run sp_helpdevice to determine the number of an available device, and enter that number using "vdevno = (number)".

For more information about sp_helpdevice and the disk init command, see the *System Administration Guide: Volume 2* and the *Reference Manual: Commands*.

Note Raw partitions for database devices provide little performance advantage over files as database devices and might have been favored in past releases for cache coherence and security. However, because the Windows file system now addresses these concerns, it is recommended that you do not use raw partitions.

Backing up and restoring data

Sybase supports tape drives and hard disks for backing up and restoring databases.

- The dump command backs up databases and transaction logs.

  To back up your databases, follow the instruction for "Using a tape drive" on page 155 or "Using a hard disk" on page 158, depending on which media you plan to use for the dump.

- The load command restores databases and transaction logs.

  To copy Sybase-supplied databases, see the *Installation Guide*.

Note Always use the Adaptive Server dump database and load database commands, rather than the Windows backup and restore utilities, to back up and restore Adaptive Server databases. Using the Adaptive Server commands ensures database integrity.

For more information about backing up and restoring databases, see the *System Administration Guide: Volume 2*.

Using a tape drive

Sybase software can back up and restore databases to tape drives that are compatible with Windows, including:

- 1/4-inch cartridge

- 4-mm and 8-mm digital audio tape (DAT) formats

To back up a database to a tape drive:

1 Install the tape drive according to the manufacturer's instructions.

   This task includes installing an Windows-compatible driver for the tape drive by using the Add/Remove buttons in the Tape Devices dialog box from the Control Panel. For instructions, see your tape drive and Windows operating system documentation.

2 Start isql, and connect to Adaptive Server:

```
d:\sybase\bin> isql -Usa -Ppassword -Sserver_name
```

3 Use the Windows tape device name with isql statements to name the tape drive.

For more information about using the dump and load commands, see "Examples of backing up and restoring databases" on page 159.

Windows tape drive names

Windows tape devices use the format "TAPE*n*", where *n* is the tape drive number, in its physical device names. Windows assigns the names as follows:

- TAPE0 is assigned to the tape drive with the lowest SCSI ID, then
- TAPE1 is assigned to the drive with the next highest SCSI ID, and so on until all devices have been assigned names

For example, to dump a database directly to the first tape drive, substitute the following value for the *stripe_device* parameter in the dump database command:

```
\\.\tape0
1> dump database pubs2 to "stripe_device"
2> capacity = 10000
3> go
```

The Windows setup program uses these device names to create logical device names to refer to the Windows tape devices; for example, TAPEDUMP1 and TAPEDUMPS2 (logical names) "for TAPE0 and TAPE1 (tape device names), respectively.

Note On your local computer, you can use the logical names TAPEDUMP1 and TAPEDUMP2 to refer to the associated tape devices. However, when you run the backup on a remote Backup Server, be sure to use the actual tape device names, rather than the logical names. See also "Setting the maximum capacity for a tape drive" on page 157.

To create a new, logical device name, use the sp_addumpdevice system procedure.

Setting the maximum capacity for a tape drive

To run properly, the dump command needs to know the maximum capacity of the destination tape drive. It determines this capacity in one of two ways, depending on the tape device name that you use:

- The physical device name – you must include the capacity parameter in the dump command. This parameter specifies the maximum number of bytes to write to a tape device.

  Check your tape's capacity, and keep the following in mind:

  - The minimum value that the capacity parameter can accept is 5 databases pages, 2K each.
  - The maximum value that the capacity parameter can accept is 4,294,967,295K.
  - The actual capacity value should be 70 to 80 percent of the true capacity of the tape.
  - If you omit the capacity parameter for Windows, Backup Server writes the maximum number of bytes for the specified tape device.

- The logical device name – the command uses the size parameter stored in the sysdevices system table.

  You can override that value by using the capacity parameter as described in the preceding list item.

Using a hard disk

Sybase software can back up data to any existing directory on a mounted Windows volume.

To back up a database to a hard disk:

1 Select a volume that has enough free space to hold the database.

2 To place the database file in a new directory on the volume, use the mkdir command to create the directory.

3 Start isql and connect to Adaptive Server:

```
d:\sybase\bin> isql -Usa -Ppassword -Sserver_name
```

4 Use the full drive, path, and file name designation to name the dump device.

For more information about using the dump and load commands, see "Examples of backing up and restoring databases" on page 159.

Dumping across a network

Backup Server may issue an "Access denied" message when you try to dump to a device mounted from across a network.

By default, all Windows services are started by using the "LocalSystem" user account, which does not allow the service to access network-mounted drives, for example, NFS, NetWare, or NTFS mounts from other machines.

To work around this restriction, configure Backup Server to start with a regular user account, rather than the Windows default account. The user should have the permission to access remote drives.

To start Backup Server with a regular user account:

1 Double-click the Services icon from the Control Panel.

2 Select the Backup Server to configure, and click the Startup button.

3 In the Log On As area, name the user in the This Account box to activate that option, type the user's password, and confirm that password.

4 Click OK to exit the Services dialog box.

5 Click Close to exit Services.

Examples of backing up and restoring databases

Following are examples of using the dump and load commands for backup and recovery of Adaptive Server database on Windows. For more information, see the *System Administration Guide: Volume 2*.

User databases

The following sections provide examples for backing up and restoring user database.

Specifying a database and device

This section provides examples on using a tape drive and a *.dat* file as the backup and recovery resources.

Using a tape drive

In the commands in this section, the physical device name TAPE0 replaces the *stripe_device* variable.

To use the first tape device to back up and load a database:

```
1> dump database pubs2 to "\\.\TAPE0"
2> go
1> load database pubs2 from "\\.\TAPE0"
2> go
```

Using a *.dat* file

To back up and load the pubs2 database using a *.dat* file:

```
1> dump database pubs2 to "d:\backups\backup1.dat"
2> go
1> load database pubs2 from "d:\backups\backup1.dat"
2> go
```

Specifying a remote Backup Server

To back up to and restore from the first tape drive on a remote Windows Backup Server named REMOTE_BKP_SERVER:

```
1> dump database pubs2 to "\\.\TAPE0" at REMOTE_BKP_SERVER
2> go
1> load database pubs2 from "\\.\TAPE0" at REMOTE_BKP_SERVER
2> go
```

Naming a backup file

To back up a transaction log and create a default backup file name:

```
1> dump tran publications to "\\.\TAPE0"
```

```
2> go
```

To restore the log using the default file name in the file clause:

```
1> load tran publications from "\\.\TAPE0"
2> with file = "cations930590E100"
3> go
```

Note The dump command uses the last 7 characters in the database name publications to create the transaction log backup file *930590E100*. See the *System Administration Guide*.

In the following example, as directed by the user, the 15-character file name, *personnel97sep111800* records the following backup information:

- The database name (personnel)
- The date (*97sep11*) – September 11, 1997
- The time (*1800*) – 18:00 or 6:00 p.m.

To back up the personnel database using the file clause to create the file name:

```
1> dump database personnel to "\\.\TAPE0"
2> with file = "personnel97sep111800"
3> go
```

To restore the personnel database by advancing the tape automatically to *personnel97sep111800* before restoring:

```
1> load database personnel from "\\.\TAPE0"
2> with file = "personnel97sep111800"
3> go
```

Note The file names in the preceding examples are valid only for systems that use the NTFS file system. If you are using a FAT-based file system, file names are limited to 8 characters with a 3-character extension.

Specifying additional dump devices

To back up the database to three devices using the stripe on parameter and *three* devices:

```
1> dump database personnel to "\\.\TAPE0"
2> stripe on "\\.\TAPE1"
3> stripe on "\\.\TAPE2"
4> go
```

To restore the database using the stripe on parameter and *two* devices:

```
1> load database personnel from "\\.\TAPE0"
2> stripe on "\\.\TAPE1"
3> go
```

To back up a database using three devices, each attached to the remote Backup Server, REMOTE_BKP_SERVER:

```
1> dump database personnel
2> to "\\.\TAPE0" at REMOTE_BKP_SERVER
3> stripe on "\\.\TAPE1" at REMOTE_BKP_SERVER
4> stripe on "\\.\TAPE2" at REMOTE_BKP_SERVER
5> go
```

Tape handling options

To initialize two devices to overwrite the existing contents with the new transaction log backups:

```
1> dump transaction personnel to "\\.\TAPE0"
2> stripe on "\\.\TAPE1" with init
3> go
```

Getting information about files

To return header information for the first file on the tape:

```
1> load database personnel from "\\.\TAPE0"
2> with headeronly
3> go
```

To return header information for the file *personnel9229510945*:

```
1> load database personnel from "\\.\TAPE0"
2> with headeronly, file = "personnel9229510945"
3> go
```

System databases

You can back up system databases the same way you back up user databases. It is not necessary to back up the tempdb database, as it is re-created every time the server restarts.

For more information, see the *System Administration Guide: Volume 2* and the *Transact-SQL Users Guide*

Optimizing Adaptive Server performance and tuning

You can make changes to your Windows system to improve Adaptive Server performance. The Windows utilities let you monitor Adaptive Server's use of operating system resources—disk, memory, and I/O—to see if you need to make any changes to your system.

For more information, see *Performance and Tuning Series*.

Using dedicated Adaptive Server operation

Installing Adaptive Server on a dedicated computer improves performance, because the software does not have to share system resources with file and print server applications. However, Adaptive Server is not a foreground application, because it runs as an Windows service. Increasing the priority of Adaptive Server increases the CPU time available for the server.

To increase the priority of Adaptive Server:

1 Start the Server Config tool either from the Sybase menu or from the Sybase Central Utilities panel.

2 Select Configure Adaptive Server.

3 Select the server to configure, then click Continue.

4 If the server needs to be started, click Yes, and enter an "sa" login and password when prompted.

5 Select Command Line Parameters.

6 Enter `-P` in the parameter entry field.

7 Click OK.

When the server restarts, it picks up the new command line parameter.

Using disk drives

The overall performance in an I/O-bound application is determined by the number of disk drives on a system, not by the amount of space available. A single disk drive might not be able to deliver the number of I/Os per second that are needed for your Adaptive Server application.

To achieve your performance objectives for an application, you must have enough disk drives to give the necessary number of I/Os per second.

Note Your disk drive requirements may not be directly related to the size of your database. Depending on the amount of I/O you need, you may have free space on your disk drives.

Monitoring disk usage

Sybase recommends that you distribute data in heavily used databases across multiple disks. To do this effectively, you need to monitor disk usage.

If one or more disks are consistently very busy, distribute the database objects on those disks to other devices. This strategy spreads out the work among disks and allow for greater data throughput.

You can use stored system procedures on Adaptive Server to monitor the disk space:

- To determine which devices a specific database is using, run sp_helpdevice or sp_helpdb.

  For more information, see sp_helpdevice and sp_helpdb in the *Reference Manual: Procedures*; also see the *System Administration Guide: Volume 2*.

- To check for disk space usage rates and I/O contention, run sp_sysmon.

  For more information, see sp_sysmon in the *Reference Manual: Procedures*; see also the *Performand and Tuning Series: Monitoring Adaptive Server with sp_sysmon*.

Monitoring Adaptive Server statistics with Windows Performance Monitor

You can use the Windows Performance Monitor to monitor Adaptive Server statistics.

Note You must be running the running the 32-bit version of Adaptive Server on the 32-bit Windows operating system for sybperf to work. For example, the 32-bit version of sybperf does not work with Windows 2008.

To support performance monitor integration, Adaptive Server must be registered as an Windows Service. This registration occurs automatically in the following situations:

- When you use the Services option through the Control Panel
- When you have configured Windows to start Adaptive Server as an automatic service

To enable performance monitoring, make sure that the SQL Perfmon Integration configuration parameter is set to 1. If necessary, use the sp_configure system procedure to reset this parameter.

Note After you set this parameter, you must restart Adaptive Server for the setting to take effect.

To monitor selected Adaptive Server statistics from Windows Performance Monitor:

1 Start the Windows Performance Monitor (*perfmon.exe*) from its program group.

2 Choose Add to Chart from the Edit menu.

The Add to Chart dialog box appears.

3 Select the computer to monitor, if necessary.

- For a local computer, skip this step and go to step 4.
- For a remote computer, click the drop-down list button on the Computer text box, select the computer you are monitoring from the Select Computer dialog box, and click OK.

4 Select the Adaptive Server Counter group that contains the counter to monitor from the Object drop-down list.

5 Select the counter that you want to monitor from the Counter list for the selected group.

   For an explanation of a particular counter, select the counter and click the Explain button. The bottom of the dialog box displays the explanation.

6 If selecting a counter displays numbers in the Instance box, select the instance that you want to monitor.

7 Click Add to activate the counter on the Performance Monitor display.

For general information on the Windows Performance Monitor, see your Windows documentation.

CHAPTER 11

Adding Optional Functionality to Adaptive Server

This chapter provides instructions for adding the following optional functionality for Adaptive Server:

- Auditing – tracks security-related system activity in an audit trail, which can be used to detect penetration of the system and misuse of resources.

Topics include:

Installing auditing

Auditing is an important part of security in a database management system. Security-related system activity is recorded in an audit trail, which can be used to detect penetration of the system and misuse of resources. By examining the audit trail, the system security officer can inspect patterns of access to objects in databases and can monitor the activity of specific users. Audit records can be traced to specific users, enabling the audit system to act as a deterrent to users who are attempting to misuse the system.

A System Security Officer manages the audit system and is the only user who can start and stop auditing, set up auditing options, and process audit data.

Audit system devices and databases

The audit system includes several components. The major components are:

- The sybsecurity device and the sybsecurity database, which stores audit information
- The audit trail, which is composed of several audit devices and tables that you determine at configuration time
- The syslogs transaction log device, which stores transaction logs

The *sybsecurity* device and database

The sybsecurity device stores the sybsecurity database. The sybsecurity database is created as part of the auditing configuration process. It contains all the system tables in the model database as well as a system table for keeping track of server-wide auditing options and system tables for the audit trail.

Tables and devices for the audit trail

Adaptive Server stores the audit trail in system tables, named sysaudits_01 through sysaudits_08. For example, if you have two audit tables, they are named sysaudits_01 and sysaudits_02. At any given time, only *one* of the audit tables is *current*. Adaptive Server writes all audit data to the current audit table. A System Security Officer can use sp_configure to set or change which audit table is current.

When you configure Adaptive Server for auditing, you determine the number of audit tables for your installation. You can specify up to eight system tables (sysaudits_01 through sysaudits_08). Plan to use at least two or three system tables for the audit trail and to put each system table on its own device, separate from the master device. If you do this, you can use a threshold procedure that archives the current audit table automatically, before it fills up and switches to a new, empty table for subsequent audit records.

Device for *syslogs* systems table

When you configure for auditing, you must specify a separate device for the syslogs system table, which contains the transaction log. The syslogs table, which exists in every database, contains a log of transactions that are executed in the database.

Pre-installation tasks for auditing devices

Determine the location of the raw devices you need for the sybsecurity, syslogs, and sysaudits table devices. You will need to provide this information later.

It is recommended that you configure your system with the minimum number of auditing devices you require—you must configure for at least three devices. You can add more auditing devices later with sp_addaudittable. For information, see the *Reference Manual: Procedures*.

Sybase recommends:

- Installing auditing tables and devices in a one-to-one ratio

  Tables that share the same device will share the same upper threshold limit. These tables cannot be used sequentially when a device fills up, because they both reside on the same device.

- Installing each auditing table on its own device

  This enables you to set up a smoothly running auditing system with no loss of auditing records.

  With two auditing tables, when one fills up, you can switch to the other. With a third auditing table, if one device fails, the System Security Officer can install a new threshold procedure that changes the device rotation to skip the broken device until the device is repaired.

- Making the device larger than the table

  When you use only three auditing tables and devices, the size of the table and the size of the device can be similar, because you can obtain more auditing capacity by adding more auditing tables and devices (up to eight). When you are working toward the upper table and device limit (six to eight), you may want to make the device considerably larger than the table. Then, you can expand the table size later towards the upper size of the device when a larger auditing capacity is desired, and few or no device additions are available.

Installing Auditing

The Adaptive Server auditing feature records information about the use of the server. By default, the auditing feature is not installed, but you can install it by using the instructions in this section. For more information about the auditing features, see the *Security Administration Guide*.

The basic steps to install auditing include:

- Create auditing devices.
- Create the auditing database.
- Run the instsecu script to populate the database tables.

To install auditing:

1 Open a Command Prompt window.

2 Start the isql program as user “sa”:

```
isql -Usa -Ppassword -Sserver_name
```

3 Determine the next available device number to use for the auditing device using statements similar to the following:

For the auditing database itself:

```
1> declare @devno int
2> select @devno = max(low/16777216)+1 from
sysdevices
3>disk init
4> name = "auditdev",
5> physname = "%SYBASE%\data\sybaud.dat",
6> vdevno = @devno,
7> size = 5120
8> go
```

For the auditing database log:

```
1>declare @devno int
2> select @devno = max(low/16777216)+1 from
sysdevices
3> disk init
4> name = "auditlogdev",
5> physname = "%SYBASE%\data\sybaudlg.dat",
6> vdevno = @devno,
7> size = 1024
8> go
```

4 At the isql prompt, use the disk init command to create the auditing devices.

5 Create the auditing database:

```
1> create database sybsecurity on auditdev
2> log on auditlogdev
3> go
```

6 Exit isql:

```
exit
```

7 Change to the *scripts* directory:

```
cd %SYBASE%\ASE-15_0\scripts
```

8 Set the DSQUERY environment variable:

```
set DSQUERY = server_name
```

9 Start the isql program as user "sa" with the instsecu script as the input file:

```
isql -Usa -Ppassword -Sserver_name -iinstsecu
```

10 Restart Adaptive Server.

After auditing is installed, no auditing occurs until a System Administrator or System Security Officer enables auditing with the auditing system procedures. See the *Security Administration Guide* for information about enabling auditing features.

Installing online help for Transact-SQL syntax

This section provides instructions for installing online help for Transact-SQL syntax.

Online syntax help: *sp_syntax*

The *%SYBASE%\%SYBASE_ASE%\scripts* directory contains scripts for installing the syntax help database, sybsyntax. You can retrieve this data with the sp_syntax system procedure. For more information on sp_syntax, see the *Reference Manual: Procedures*.

All Adaptive Server installations receive the ins_syn_sql script. This script includes syntax information for Transact-SQL, the system procedures, and the Sybase utilities. When you execute this script, you install the SQL portion of the sybsyntax database.

You can install any of these scripts, depending on the need for Sybase information on your server. The first script you execute creates the sybsyntax database and the needed tables and indexes. Any scripts that you execute after the first one add to the existing information in the database. If you execute a script that was executed previously, the previously installed rows of information are deleted from the table in the database and then reinstalled.

Default device for the *sybsyntax* database

The sybsyntax database requires space on the device that is at least as large as the model database. By default, the sybsyntax installation scripts install the sybsyntax database on the device that is designated as the default database device.

If you have not used sp_diskdefault to change the status of the master device (which is installed as the default disk) or to specify another default device, the scripts install sybsyntax on the master device. This configuration is not recommended because sybsyntax uses valuable space, which is best left available for future expansion of the master database.

To avoid installing sybsyntax on the master device, do one of the following:

- Use sp_diskdefault to specify a default device other than the master device. For information about sp_diskdefault, see the *Reference Manual: Procedures*.
- Modify each sybsyntax installation script that you plan to execute to specify a different device, as explained in the following section.

Installing *sybsyntax*

For each sybsyntax installation script you want to execute:

1 Determine the type (raw partition, logical volume, operating system file, and so on) and location of the device where you plan to store the sybsyntax database. You will need to provide this information later.

2 Make a copy of the original script. Be sure you can access this copy, in case you experience problems with the edited script.

3 Use a text editor to edit the script, if necessary, to change the default device from the master device to the device created in step 1. For information on the default device, see “Default device for the sybsyntax database” on page 172.

   - Comment out the following section, which specifies the default device:

```
/* create the database, if it does not exist */
 if not exists (select name from sysdatabases
 where name = "sybsyntax")
 begin
     /* create the sybsyntax table if it doesn't exist */
     /* is the space left on the default database
```

```
        devices > size of model? */
        if (select sum (high-low +1) from sysdevices where status
        & 1 = 1) - (select sum(size) from sysusages, sysdevices
            where vstart >= sysdevices.low
            and vstart <= sysdevices.high
            and sysdevices.status &1 = 1) >
            (select sum(sysusages.size) from sysusages
            where dbid = 3)
        begin
            create database sybsyntax
        end
        else
        begin
            print "There is not enough room on the default
            devices to create the sybsyntax database."
        return
        end
    end
```

- After you have commented out this entire section, add a line like this to the script:

```
create database sybsyntax on device_name
```

where *device_name* is the name of the device where you want to install sybsyntax.

4 Execute the script with a command like the following:

```
isql -Usa -Ppassword -Sservername <
%SYBASE%\%SYBASE_ASE%\scripts\ins_syn_sql
```

where *sa* is the user ID of the System Administrator, *password* is the System Administrator's password, and *servername* is the Adaptive Server where you plan to install the database.

If you have set the DSQUERY environment variable to the *servername*, you can replace the server name with DSQUERY. For example:

```
isql -Usa -Ppassword -S$DSQUERY <
%SYBASE%\%SYBASE_ASE%\scripts\Sins_syn_sql
```

5 To ensure that you have installed the sybsyntax database and that it is working correctly, use isql to log in to the server on which you installed the database, and execute sp_syntax. For example:

```
isql -Usa -Ppassword -Sservername

1> sp_syntax "select"
2> go
```

Adaptive Server displays a list of commands that contain the word or word fragment “select”.

CHAPTER 12

Troubleshooting Network Connections

Net-Library enables clients and Adaptive Servers to interact with each other over a network. If the Net-Library software is not functioning properly, the client/server environment will not function properly either.

This chapter describes how to use the Server Ping utility in the Directory Services Editor (dsedit) to get information about Adaptive Servers on a network.

The dsedit Server ping utility

Use the Directory Services Editor (dsedit) utility's Server Ping utility to run tests on the Net-Library-to-server connections across your network software. Server Ping reports information about both successful connections and failed connection attempts.

This test is particularly useful when you have multiple server names to identify more than one server in the *sql.ini* file.

You do not need to have a valid user name on Adaptive Server to run Server Ping.

Running server ping

You can test the connections to any server that has a name in the *sql.ini* file on your client, as described in "How a client accesses Adaptive Server" on page 32.

1 Start dsedit.

2 Select the directory service to open from the Select Directory Service dialog box, and click OK.

   The Interfaces Driver dialog box for the server appears.

3 Select the name of the server to test from the list of server names.

   The server information displayed depends upon the specific Net-Library driver that you have installed.

4 Select Server Object/Server Ping.

   The Ping dialog box appears.

5 Click Ping to test the connection.

If Server Ping makes a successful connection to the server, a message indicating the success appears in a dsedit dialog box. A successful connection indicates that you have properly configured your Adaptive Server for network access.

If Server Ping reports an unsuccessful connection to the server, see "Troubleshooting connection failures" on page 176.

Troubleshooting connection failures

When a client application fails to connect to a server, you can test the application for diagnostic purposes. The messages that the Server Ping utility displays may provide you with enough information to solve the problem.

This test, however, cannot diagnose all types of network connection problems. Some problems may result from issues in your Adaptive Server setup, rather than in your Net-Library-to-network-software connection.

For tips on troubleshooting these setup problems, see "Failure of other applications" on page 178.

When a test fails

When Server Ping reports an unsuccessful connection, check to make sure that:

- Adaptive Server is running on the target server.
- A network hardware connection exists between your client machine and the target server.
- The server meets the minimum hardware and software requirements (see the *Installation Guide*).
- The network software is installed and configured on the client and the server.
- The connection information in the *sql.ini* file is correct for the server.
- The connection information in your client's network configuration file is correct. For more information, see the Net-Library documentation for your client.
- The format of the connection information is correct for the network protocol. See"Components in the sql.ini file" on page 33.

If you need to edit *sql.ini*, use dsedit.

Warning! Make sure that no more than one copy of any Net-Library DLL is installed on your computer.

Using returned messages to diagnose a failure

When you are sure that the requirements named in "When a test fails" on page 177 have been met, determine the point at which the Server Ping failed by reviewing the resulting messages.

Failure to connect to Adaptive Server

When Server Ping does not connect to a server, dsedit displays information about what went wrong. For example, if the server is not running, the message shown in this next screen might appear:

Since it loaded the Net-Library DLL, dsedit found connection information in *sql.ini*. When the connection succeeds in finding the information, but notifies you that the server is not responding, you can use that information to discover the problem.

❖ **Troubleshooting an unsuccessful Server Ping**

1 Verify that the server is running.

2 Check that your networking software and hardware are properly configured.

3 Check to see if any network error messages are displayed.

4 Check that the connection information is correct for your network protocol and that format of your entries matches the format shown in Chapter 4, "Network Communications Using sql.ini."

Failure to load Net-Library DLLs

Server Ping displays a message when it cannot load the Net-Library DLL. verify that the directory containing Net-Library DLL is included in the PATH environment variable.

Failure of other applications

When Server Ping reports no errors, but your other applications fail to run, use this information to discover the problem.

❖ **Troubleshooting a falsely successful Server Ping**

1 Verify that the Net-Library driver that you want to use is listed in the *libtcl.cfg* file.

The utility does not look in *libtcl.cfg*, so Server Ping can be successful, even if the *libtcl.cfg* file contains incorrect information. The *libtcl.cfg* file is in the *ini* subdirectory of the Sybase installation directory.

2 Use isql to verify that you can access Adaptive Server locally from your computer.

3 Use isql to verify that the databases and tables used by your client application exist.

4 Verify that you have a valid user login name for Adaptive Server.

5 Verify that you have permissions on databases and tables that are consistent with the permissions required to run your applications.

Occasionally, a Server Ping result might indicate inaccurately a successful connection to Adaptive Server because dsedit found some other application listening at the specified Adaptive Server address. dsedit does not recognize that the non-Sybase application is not an Adaptive Server. To determine if this is the case, try to connect to the server with isql.

Before calling Sybase Technical Support

For problems with your Net-Library application, have the following information available when you call Sybase Technical Support:

- The text of the diagnostic utility error
- A listing of your *sql.ini* file
- The name and version number of your network software
- The name and version number of the operating system on which your client and server networking software is running
- The version number of the server to which you are connected
- The date and size of your Net-Library DLL

  To locate this library information, execute the dir command to display a file list that includes the DLL.

APPENDIX A

Adaptive Server Registry Keys

The Windows operating system stores configuration information in a tree-structured file called the Registry.

When you install Adaptive Server for Windows, the installation program and Server Config utility write configuration information to several branches, called *keys*, in the Windows Registry.

This appendix presents the Registry values in a series of tables, one table for each key that appears under HKEY_LOCAL_MACHINE in the Registry:

- \SOFTWARE\SYBASE\Server\server_name – Table A-1
- \SOFTWARE\SYBASE\SQLServer\server_name\parameter – Table A-2
- \SOFTWARE\SYBASE\SQLServer – Table A-3
- \SYSTEM\CurrentControlSet\Services\SYBSQL_server_name – Table A-4

In some cases, you can use the information in this appendix to configure features of Adaptive Server. However, you can seriously impair your Windows system if you make incorrect changes to the Registry.

Warning! Do not modify key values in the Registry unless you are an experienced Windows administrator, and you are familiar with the regedt32 utility. See your system Windows documentation for information about using regedt32.

Table A-1: \SOFTWARE\SYBASE\Server\server_name

HKEY_LOCAL_MACHINE\SOFTWARE\SYBASE\Server\server_name			
Key name	**Type**	**Default**	**Description**
DefaultDomain	REG_SZ	None	The default domain for mapping Windows user names to Adaptive Server logins
DefaultLogin	REG_SZ	None	The login ID to use for access to Adaptive Server when an authorized user does not have an Adaptive Server login defined in syslogins

HKEY_LOCAL_MACHINE\SOFTWARE\SYBASE\Server\server_name			
Key name	**Type**	**Default**	**Description**
LoginMode	REG_DWORD	0	The login security mode: • 0 indicates Standard • 1 indicates Integrated • 2 indicates Mixed
Map#	REG_SZ	Dash (-)	The special character mapped to the valid Adaptive Server pound sign (#) character
Map$	REG_SZ	Space ( )	The special character mapped to the valid Adaptive Server dollar sign ($) character
Map@	REG_SZ	Space ( )	The special character mapped to the valid Adaptive Server at sign (@) character
Map_	REG_SZ	Domain Separator (\)	The special character mapped to the valid Adaptive Server underscore (_) character
ServerType	REG_SZ	SQLServer	The type of server
SetHostName	REG_DWORD	0	Replacement status of the host name from the client login by the network user name under integrated security; • 1 = yes • 0 = no

Table A-2: \SOFTWARE\SYBASE\SQLServer\server_name\parameter

HKEY_LOCAL_MACHINE\SOFTWARE\SYBASE\SQLServer\server_name\parameters			
Key name	**Type**	**Default**	**Description**
Arg0	REG_SZ	-d*D:\sybase\ASE-15_0\data\master.dat*	The location of the master device file
Arg1	REG_SZ	-s*server_name*	The name of the Adaptive Server
Arg2	REG_SZ	-e*d:\sybase\ASE-15_0\install\errorlog*	The location and name of the error log file
Arg3	REG_SZ	-i*d:\sybase\ini*	The location of the *sql.ini* file
Arg4	REG_SZ	-M*d:\sybase*	The directory that stores shared memory files
Arg5	REG-SZ	-N*d:\sybase\ASE-15_0\sysam\<srv_name>.properties*	Location and name of license cache file.

Table A-3: \SOFTWARE\SYBASE\SQLServer

HKEY_LOCAL_MACHINE\SOFTWARE\SYBASE\SQLServer			
Key name	**Type**	**Default**	**Description**
CurrentVersion	REG_SZ	Windows 15.0	The version number for the Adaptive Server software installed on the computer
DefaultBackupServer	REG_SZ	*server_nam*e_BS	The name of the default Backup Server

HKEY_LOCAL_MACHINE\SOFTWARE\SYBASE\SQLServer			
Key name	**Type**	**Default**	**Description**
DSEVNTLOG	REG_SZ	LocalSystem	The destination machine for logging messages to the Windows event log
DSLISTEN	REG_SZ	*server_name*	The name Adaptive Server uses to listen for client connections when no name is given during Adaptive Server start-up
RootDir	REG_SZ	*D:\sybase*	The location of the Sybase installation directory for client applications to look for. Lists the SYBASE environment variable.
Version	REG_SZ	15.0	The version number of the Adaptive Server

Table A-4: \SYSTEM\CurrentControlSet\Services\SYBSQL_server_name

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\SYBSQL_server_name			
Key Name	**Type**	**Default**	**Description**
DisplayName	REG_SZ	Sybase SQL Server_*server_name*	The Adaptive Server name used in the Services list under Control Panel
ErrorControl	REG_DWORD	0x1	For system use only
ImagePath	REG_EXPAND_SZ	*D:\Sybase\ASE-15_0\bin\sqlsrvr.exe -s<server_name>* -C	The path for the Adaptive Server executable file
ObjectName	REG_SZ	LocalSystem	For system use only
Start	REG_DWORD	0x2	For system use only
Type	REG_DWORD	0x10	For system use only

Index

Symbols

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z