



Introducing Sybase® Adaptive Server™ Enterprise
for Windows NT

Adaptive Server™

Document ID: 34977-01-1150

September 1997

Copyright Information

Copyright © 1989–1997 by Sybase, Inc. All rights reserved.

Sybase, Inc., 6475 Christie Avenue, Emeryville, CA 94608.

Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, the Sybase logo, APT-FORMS, Certified SYBASE Professional, Data Workbench, First Impression, InfoMaker, PowerBuilder, Powersoft, Replication Server, S-Designer, SQL Advantage, SQL Debug, SQL SMART, SQL Solutions, Transact-SQL, VisualWriter, and VQL are registered trademarks of Sybase, Inc. Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Monitor, ADA Workbench, AnswerBase, Application Manager, AppModeler, APT-Build, APT-Edit, APT-Execute, APT-Library, APT-Translator, APT Workbench, Backup Server, BayCam, Bit-Wise, ClearConnect, Client-Library, Client Services, CodeBank, Column Design, Connection Manager, DataArchitect, Database Analyzer, DataExpress, Data Pipeline, DataWindow, DB-Library, dbQ, Developers Workbench, DirectConnect, Distribution Agent, Distribution Director, Dynamo, Embedded SQL, EMS, Enterprise Client/Server, Enterprise Connect, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, EWA, Formula One, Gateway Manager, GeoPoint, ImpactNow, InformationConnect, InstaHelp, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, Logical Memory Manager, MainframeConnect, Maintenance Express, MAP, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, MethodSet, Net-Gateway, NetImpact, Net-Library, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, PB-Gen, PC APT-Execute, PC DB-Net, PC Net Library, Power++, Power AMC, PowerBuilt, PowerBuilt with PowerBuilder, PowerDesigner, Power J, PowerScript, PowerSite, PowerSocket, Powersoft Portfolio, Power Through Knowledge, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Quickstart Datamart, Replication Agent, Replication Driver, Replication Server Manager, Report-Execute, Report Workbench, Resource Manager, RW-DisplayLib, RW-Library, SAFE, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILLS, smart.partners, smart.parts, smart.script, SQL Anywhere, SQL Central, SQL Code Checker, SQL Edit, SQL Edit/TPU, SQL Modeler, SQL Remote, SQL Server, SQL Server/CFT, SQL Server/DBM, SQL Server Manager, SQL Server SNMP SubAgent, SQL Station, SQL Toolset, Sybase Client/Server Interfaces, Sybase Development Framework, Sybase Gateways, Sybase IQ, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase Synergy Program, Sybase Virtual Server Architecture, Sybase User Workbench, SybaseWare, SyBooks, System 10, System 11, the System XI logo, SystemTools, Tabular Data Stream, The Architecture for Change, The Enterprise Client/Server Company, The Model for Client/Server Solutions, The Online Information Center, Translation Toolkit, Turning Imagination Into Reality, Unibom, Unilib, Uninull, Unisep, Unistring, Viewer, Visual Components, VisualSpeller, WarehouseArchitect, WarehouseNow, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, and XA-Server are trademarks of Sybase, Inc. 6/97

All other company and product names used herein may be trademarks or registered trademarks of their respective companies.

Table of Contents

About This Book

Audience	ix
How to Use This Guide	ix
Related Documents	ix
Conventions	x
Adaptive Server Enterprise Documents	x
Other Sources of Information	xii
If You Need Help	xii

1. Overview of Adaptive Server Enterprise

Relational Database Management Systems	1-1
Tables	1-1
SQL	1-4
Auxiliary Facilities of Adaptive Server's RDBMS	1-4
Client/Server Architecture	1-4
Division of Tasks	1-6
Distributed Computing	1-6
The Internet or Intranets	1-8
Components of Adaptive Server Enterprise	1-11
Server Components	1-12
Client Components	1-14
Sybase Documentation	1-17

2. Terms and Concepts

Database Object	2-1
Row and Column	2-1
View	2-4
Index	2-4
Rule	2-5
Stored Procedure	2-5
Trigger	2-6
Transact-SQL	2-6
Queries	2-7
Other Commands	2-9
Presenting Transact-SQL to the RDBMS	2-10

System Database	2-10
Optional Sample Databases	2-11
Database Device	2-11
System Table	2-12
Permission	2-12
Utility Programs	2-13
Transaction Log	2-13
System Administration	2-13

3. Using Adaptive Server Enterprise

Installing Adaptive Server	3-1
Configuring Adaptive Server	3-2
Designing Databases	3-3
Creating Databases	3-4
Loading Data	3-5
Creating Indexes	3-6
Backing Up Data	3-7
Assigning Logins, Users, and Permissions	3-7
Querying and Modifying Data	3-8
Modifications	3-9
Transactions	3-9
Stored Procedures	3-9
Building or Buying the Client Applications	3-10
Setting Up a Web Site	3-12
Monitoring and Tuning Performance	3-14

Glossary

Index

List of Figures

Figure 1-1:	Tables in a database	1-2
Figure 1-2:	Retrieving data across tables using key columns	1-3
Figure 1-3:	Sybase client/server environment	1-5
Figure 1-4:	Basic client/server interaction	1-5
Figure 1-5:	Division of tasks.....	1-6
Figure 1-6:	Distributed client/server environment	1-7
Figure 1-7:	Adaptive Server configuration implementing SQL Remote	1-8
Figure 1-8:	NetImpact Dynamo Web site configuration	1-9
Figure 1-9:	Traditional, file-based Web site.....	1-9
Figure 1-10:	How NetImpact Dynamo interacts with Adaptive Server and the Web server.	1-11
Figure 1-11:	Overview of Adaptive Server Enterprise components	1-12
Figure 2-1:	Rows and columns in a table	2-2
Figure 2-2:	Primary and foreign key columns.....	2-3
Figure 2-3:	A view	2-4
Figure 2-4:	A simplified index schematic.....	2-5
Figure 2-5:	Selection	2-7
Figure 2-6:	Projection	2-8
Figure 2-7:	A join	2-9
Figure 2-8:	A Sybase Central wizard	2-14
Figure 3-1:	Configuring Adaptive Server from the Server Config utility	3-2
Figure 3-2:	PowerDesigner PDM of the pubs2 database	3-4
Figure 3-3:	Creating a data pipeline in InfoMaker	3-6
Figure 3-4:	A simple InfoMaker form.....	3-11
Figure 3-5:	PowerBuilder development window	3-12
Figure 3-6:	NetImpact Dynamo Web site displayed in Sybase Central.....	3-13

About This Book

This book, *Introducing Sybase Adaptive Server Enterprise for Windows NT*, provides an overview of Adaptive Server Enterprise™ for Windows NT, a set of Sybase® products for developing and deploying relational database applications on desktop platforms.

Audience

Read this guide before you install and begin using Adaptive Server to establish a foundation for using the product. This guide is for system administrators, managers, or anyone who will be involved in setting up or using Adaptive Server Enterprise or who wants to understand how the product set fits together. This guide assumes no technical knowledge of the Sybase products.

How to Use This Guide

This guide includes the following chapters:

- Chapter 1, “Overview of Adaptive Server Enterprise,” introduces relational database management systems and client/server architecture. It also describes each of the products in the Adaptive Server Enterprise set of products.
- Chapter 2, “Terms and Concepts,” defines some of the basic terminology necessary for understanding discussions about using Adaptive Server.
- Chapter 3, “Using Adaptive Server Enterprise,” provides a brief introduction to some of the most common tasks a user performs with Adaptive Server Enterprise.

Related Documents

The *Read Me First* card in your product package lists the products that are included with Adaptive Server Enterprise. It also provides a “road map” for locating the related Sybase documentation available to you. *Installing Adaptive Server Enterprise for Windows NT* lists the version numbers of the included products.

Conventions

This manual uses the following style conventions:

- Within text, the names of files, directories, and database objects appear in italics:

\data\master.dat

- Examples showing the use of Transact-SQL commands are printed like this:

```
select titles, pub_id
      from titles
```

- The names of utilities, procedures, commands, and scripts appear in the following font:

`sp_revokelogin`

- Terms that you can find in the glossary appear in the following font:

global variable

Adaptive Server Enterprise Documents

The following documents comprise the Sybase Adaptive Server Enterprise documentation:

- The *Release Bulletin* for your platform – contains last-minute information that was too late to be included in the books.

A more recent version of the *Release Bulletin* may be available on the World Wide Web. To check for critical product or document information that was added after the release of the product CD, use SyBooks™-on-the-Web.

- The Adaptive Server Enterprise installation documentation for your platform – describes installation and upgrade procedures for all Adaptive Server and related Sybase products.
- The Adaptive Server Enterprise configuration documentation for your platform – describes configuring a server, creating network connections, configuring for optional functionality, such as auditing, installing most optional system databases, and performing operating system administration tasks.
- *What's New in Adaptive Server Enterprise Enterprise?* – describes the new features in Adaptive Server release 11.5, the system changes

added to support those features, and the changes that may affect your existing applications.

- *Navigating the Documentation for Adaptive Server* – an electronic interface for using Adaptive Server Enterprise. This online document provides links to the concepts and syntax in the documentation that are relevant to each task.
- *Transact-SQL User's Guide* – documents Transact-SQL, Sybase's enhanced version of the relational database language. This manual serves as a textbook for beginning users of the database management system. This manual also contains descriptions of the *pubs2* and *pubs3* sample databases.
- *System Administration Guide* – provides in-depth information about administering servers and databases. This manual includes instructions and guidelines for managing physical resources and user and system databases, and specifying character conversion, international language, and sort order settings.
- *Adaptive Server Reference Manual* – contains detailed information about all Transact-SQL commands, functions, procedures, and datatypes. This manual also contains a list of the Transact-SQL reserved words and definitions of system tables.
- *Performance and Tuning Guide* – explains how to tune Adaptive Server Enterprise for maximum performance. This manual includes information about database design issues that affect performance, query optimization, how to tune Adaptive Server for very large databases, disk and cache issues, and the effects of locking and cursors on performance.
- The *Utility Programs* manual for your platform – documents the Adaptive Server utility programs, such as *isql* and *bcp*, which are executed at the operating system level.
- *Security Features User's Guide* – provides instructions and guidelines for using the security options provided in Adaptive Server from the perspective of the non-administrative user.
- *Error Messages and Troubleshooting Guide* – explains how to resolve frequently occurring error messages and describes solutions to system problems frequently encountered by users.
- *Security Administration Guide* – explains how to use the security features provided by Adaptive Server Enterprise to control user access to data. This manual includes information about how to add users to Adaptive Server Enterprise, administer both system

and user-defined roles, grant database access to users, and manage remote Adaptive Server Enterprises.

- *Component Integration Services User's Guide for Adaptive Server Enterprise and OmniConnect* – explains how to use the Adaptive Server Component Integration Services feature to connect remote Sybase and non-Sybase databases.
- *Adaptive Server Glossary* – defines technical terms used in the Adaptive Server documentation.
- *Master Index for Adaptive Server Enterprise Publications* – combines the indexes of the *Adaptive Server Reference Manual*, *Component Integration Services User's Guide*, *Performance and Tuning Guide*, *Security Administration Guide*, *Security Features User's Guide*, *System Administration Guide*, and *Transact-SQL User's Guide*.

Other Sources of Information

Use the SyBooks™ and SyBooks-on-the-Web online resources to learn more about your product:

- SyBooks documentation is on the CD that comes with your software. The DynaText browser, also included on the CD, allows you to access technical information about your product in an easy-to-use format.

Refer to *Installing SyBooks* in your documentation package for instructions on installing and starting SyBooks.

- SyBooks-on-the-Web is an HTML version of SyBooks that you can access using a standard Web browser.

To use SyBooks-on-the-Web, go to <http://www.sybase.com>, and choose Documentation.

If You Need Help

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you have any questions about this installation or if you need assistance during the installation process, ask the designated person to contact Sybase Technical Support or the Sybase subsidiary in your area.

1

Overview of Adaptive Server Enterprise

Adaptive Server Enterprise is an integrated set of software products for designing, developing and deploying relational database applications. It consists of a high-performance relational database management system (RDBMS), which runs database servers, and a collection of applications and libraries, which run on database clients. This arrangement, consisting of servers that are accessed by multiple clients over a network, forms the basis for Sybase's **client/server architecture**.

The product includes an application called NetImpact Dynamo. With NetImpact Dynamo, you can deploy your database over a corporate intranet or over the Internet.

This chapter provides an overview of the following topics:

- Relational Database Management Systems 1-1
- Client/Server Architecture 1-4
- The Internet or Intranets 1-8
- Components of Adaptive Server Enterprise 1-11
- Sybase Documentation 1-17

Relational Database Management Systems

A **relational database management system (RDBMS)** is a system for storing and retrieving data in which the data is represented in two-dimensional **tables**. In early relational systems, tables were called relations. A relational database consists of a collection of tables that store interrelated data.

Tables

Figure 1-1 shows portions of tables that you might create in a relational database that stores related data about books: The tables are called *titles*, *authors*, *publishers*, and *titleauthor*.

titles				
title_id	title	type	pub_id	price
BU1032	The Busy Executive's Database Guide	business	1389	19.99
PS2106	Life Without Fear	psychology	0736	7.00
MC3021	The Gourmet Microwave	mod_cook	0877	2.99
BU7832	Straight Talk About Computers	business	1389	19.99
BU2075	You Can Combat Computer Stress	business	0736	2.99

publishers			
pub_id	pub_name	city	state
1389	Algodata Infosystems	Berkeley	CA
0736	New Age Books	Boston	MA
0877	Binnet & Hardley	Washington	DC

authors		
au_id	au_lname	au_fname
213-46-8915	Green	Marjorie
9948-72-3567	Ringer	Albert
274-80-9391	Straight	Dick

titleauthor	
au_id	title_id
213-46-8915	BU1032
9948-72-3567	PS2106
274-80-9391	BU7832
213-46-8915	BU2075

Figure 1-1: Tables in a database

Each table in the database holds information about different things—books, publishers, and authors. But some information in each table overlaps with information in another table—by design. Columns that appear in more than one table and that link the tables together are called **keys**. Keys, which are discussed in more detail in “Key” on page 2-2, allow you to retrieve information that is distributed over multiple tables.

For example, you might want information about books published by Marjorie Greene, which spans the *authors* and the *titles* tables. Figure 1-2 shows the key columns for the sample databases and how they allow you to retrieve the correct data from the tables.

authors		
au_id	au_lname	au_fname
213-46-8915	Green	Marjorie
9948-72-3567	Ringer	Albert
274-80-9391	Straight	Dick

titleauthor	
au_id	title_id
213-46-8915	BU1032
9948-72-3567	PS2106
274-80-9391	BU7832
213-46-8915	BU2075

titles				
title_id	title	type	pub_id	price
BU1032	The Busy Executive's Database Guide	business	1389	19.99
PS2106	Life Without Fear	psychology	0736	7.00
MC3021	The Gourmet Microwave	mod_cook	0877	2.99
BU7832	Straight Talk About Computers	business	1389	19.99
BU2075	You Can Combat Computer Stress	business	0736	2.99

Figure 1-2: Retrieving data across tables using key columns

One advantage of relational databases is that they allow you to set up multiple tables, a structure that eliminates redundancy and possible inconsistencies caused by redundancy. For example, both the sales and accounts payable departments might enter and look up information about publishers. In a relational database, the information about publishers is entered only once, in a table that both departments can access.

Defining how tables represent a body of information, determining how the data should be distributed among the tables, and deciding what keys are needed to define the relationships are the subjects of **logical design**. Sound database design is the key factor in realizing Adaptive Server's performance potential and makes a database easy to maintain and update. Many excellent books are available on this subject. *The Practical SQL Handbook* (Bowman, Emerson, and Damovsky, Addison-Wesley, 1993) offers a clear introduction to logical database design, and it specifically relates the issues to an Adaptive Server context. (See "Designing Databases" on page 3-3.)

SQL

SQL (Structured Query Language) is the database language designed for the RDBMS model. It is used to set up, use, and maintain a relational database. This is in contrast to database systems in previous computing models, which required large Information Services departments to be involved in data requests. An excellent reference for learning SQL is the *LAN Times Guide to SQL* (Groff and Weinberg, Osborne McGraw-Hill, 1994).

End users need not learn SQL in order to access data, however—in most systems, purchased or developed GUI client software translates user requests into SQL. Application programming interfaces (APIs) or libraries installed on the clients allow applications to send SQL commands to and retrieve results from the RDBMS.

Sybase's enhanced version of SQL is called **Transact-SQL** (see "Transact-SQL" on page 2-6).

Auxiliary Facilities of Adaptive Server's RDBMS

In addition to storing and retrieving data, Adaptive Server provides a number of auxiliary facilities for managing the relational databases. These include programs for starting and stopping the server, for backing up and restoring data, for copying data between a database and a file system, for authenticating users of the server and the databases, for managing the physical devices on which the data is stored, for displaying system messages in different languages, for storing and preprocessing a reusable procedure or set of procedures, and so forth. Together with the basic storage and retrieval capabilities, these facilities make up Adaptive Server's relational database management system.

Client/Server Architecture

In a Sybase Adaptive Server environment, the databases and the RDBMS software reside on one or more servers, where multiple clients can access them concurrently over a network. Applications that access and manipulate the shared data, as well as the libraries that allow the applications to communicate with the RDBMS, reside

on the clients. Figure 1-3 illustrates the Sybase client/server configuration..




Figure 1-3: Sybase client/server environment

The basic client/server relationship is one in which client applications request services from servers, and servers return results to client applications. A typical request for service in a client/server RDBMS is for the client to request that the server retrieve or modify the data that it is storing. A typical result is confirmation by the server that the request was successfully or unsuccessfully carried out, followed by the requested data. Figure 1-4 illustrates a basic client/server interaction.

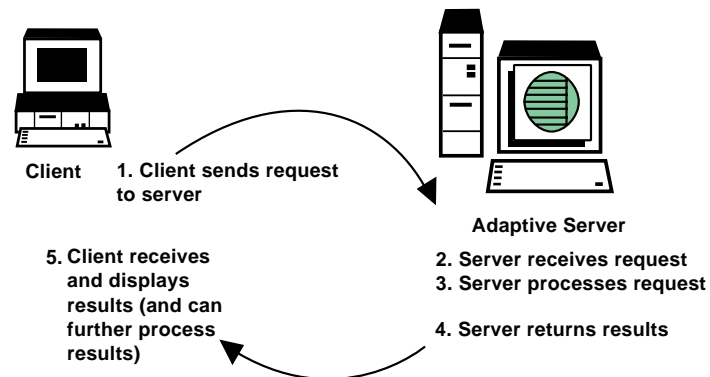


Figure 1-4: Basic client/server interaction

Division of Tasks

A client/server model provides the best performance for data sharing among a large number of PCs and workstations because clients and servers share the work—each performs the work it can do most efficiently. The server handles activities related directly to the maintenance and retrieval of shared data, while other tasks, such as displaying returned data or interpreting requests from users, are off-loaded to the clients. Figure 1-5 represents how the tasks are divided.




Figure 1-5: Division of tasks

Distributed Computing

The ability to distribute data among multiple servers provides a structure for distributed computing environments in which control is not centralized at a single data source or geographic location. Clients can access the data they need from any server.

Figure 1-6 shows a noncentralized, distributed computing configuration in which data is distributed over several sites; each site houses a different subset of the data.




Figure 1-6: Distributed client/server environment

Some organizations want each site to have a complete copy of the entire data set to improve the time it takes to access the data and to eliminate downtime in the case of one server's failure. Data **replication** offers a solution to these organizations.

Replication Server[®], a Sybase product that you can purchase separately, can be integrated with the Adaptive Server Enterprise setup. Replication Server replicates peer-to-peer data replication across a distributed environment, to and from heterogeneous hardware and data sources. All clients have equally fast and reliable access to all data.

SQL Remote

SQL Remote[™], which is part of Adaptive Server Enterprise, provides data replication between a consolidated (centralized) database and a remote database or databases. SQL Remote supports hierarchical configurations of databases; it does not support peer-to-peer replication between more than two databases or other non-hierarchical configurations. It works well for organizations in which many users work on laptop or desktop computers that are connected only occasionally to the company network.

Remote users subscribe to all or portions of a company database; each user keeps his or her own copy of his or her own data in a stand-alone database on their laptop computer. Data is updated in both directions when they make a dial-up or e-mail connection to the company network. Figure 1-7 shows an Adaptive Server configuration that implements SQL Remote.

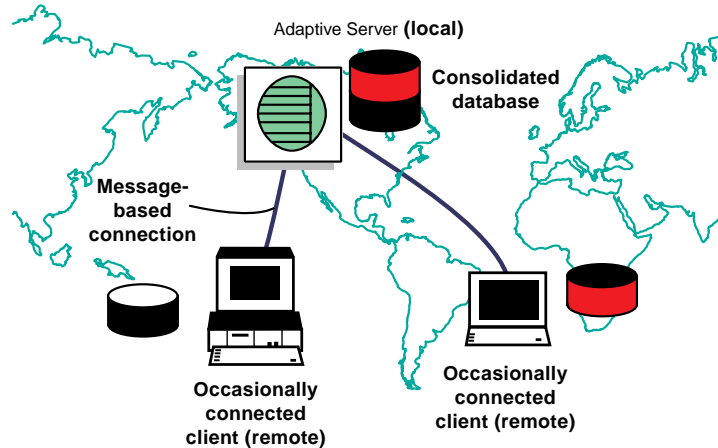


Figure 1-7: Adaptive Server configuration implementing SQL Remote

The Internet or Intranets

Many organizations supplement the traditional client/server configuration by distributing some of their data via the World Wide Web or a corporate **intranet**. Figure 1-8 illustrates how NetImpact™ Dynamo, the powerful tool and Web site application server included with Adaptive Server Enterprise, allows you to connect your

database to the Internet or a corporate intranet and to expand the traditional client/server environment.




Figure 1-8: NetImpact Dynamo Web site configuration

HTML is the language used for formatting Web documents. Noninteractive Web sites rely on static HTML text pages to present information. Figure 1-9 illustrates how static data is traditionally distributed over the Internet by a **web server**.

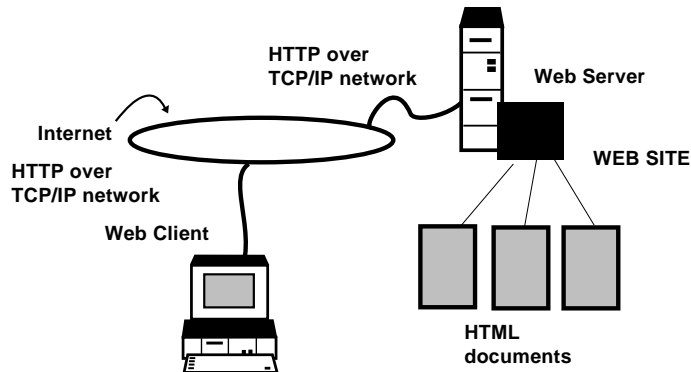


Figure 1-9: Traditional, file-based Web site

The Common Gateway Interface (**CGI**) allows Web servers to provide some interactivity. Web servers call CGI programs that accept and process input from users. CGI programs are usually written in an interpreted language like **PERL** and can also interact in a limited way with external programs like database servers.

NetImpact Dynamo makes your Web site fully interactive and fully connected to the Adaptive Server RDBMS. As shown in Figure 1-10 on page 1-11, it includes an **application server**, or gateway, that acts as an intermediary between the Web server and Adaptive Server.

The NetImpact Dynamo application server receives and processes instructions from a Web server, retrieves data from a database server, and delivers it to the Web server.

Because a NetImpact Dynamo Web site can accept input from the user, **web clients** can act as traditional clients to Adaptive Server through the site by inputting data in forms embedded in the HTML documents. NetImpact Dynamo forwards the data to Adaptive Server. For further details, see “NetImpact Dynamo” on page 1-14 and “Setting Up a Web Site” on page 3-12.

For busy sites, you can offload some of the processing work, such as verifying user input or displaying error messages, to the Web client by embedding client-side JavaScript scripts into HTML templates.

Figure 1-10 illustrates how the NetImpact Dynamo application server interacts with the Web server and with Adaptive Server. The application server itself replaces the CGI program, or the equivalent, in typical interactive Web sites.

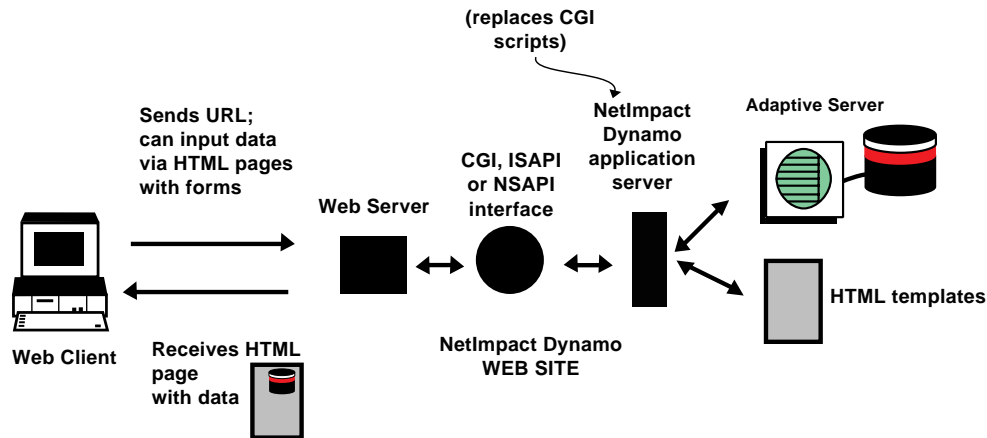


Figure 1-10: How NetImpact Dynamo interacts with Adaptive Server and the Web server

Components of Adaptive Server Enterprise

As described in “Client/Server Architecture” on page 1-4, the basic client/server RDBMS model consists of the database server and third-party or developed client software that communicates with the database server over a network. Adaptive Server Enterprise also includes software for backing up, monitoring, and administering Adaptive Server; client libraries that enable applications to communicate with the server; and developed client applications that can meet many of your computing needs.

You install some Adaptive Server Enterprise software on the server and some on the client. It is also convenient to run some of the client administrative tools and utilities from the server desktop.

On the server, you create files for use as **database devices**, files dedicated to storing data. Use the Transact-SQL `disk init` command to initialize the devices.

On the client, you can install (in addition to the client components of Adaptive Server) **development tools** such as PowerBuilder® that help you build applications that access Adaptive Server, the applications themselves, and/or third-party software.

Figure 1-11 shows an overview of the client and server components of an Adaptive Server system. You can install any or all client

components on the server instead of on a client. The Open Client™, and, in some cases, the ODBC drivers, must be installed on all machines that will run client software.

The sections “Server Components” on page 1-12 and “Client Components” on page 1-14 describe each component in more detail.

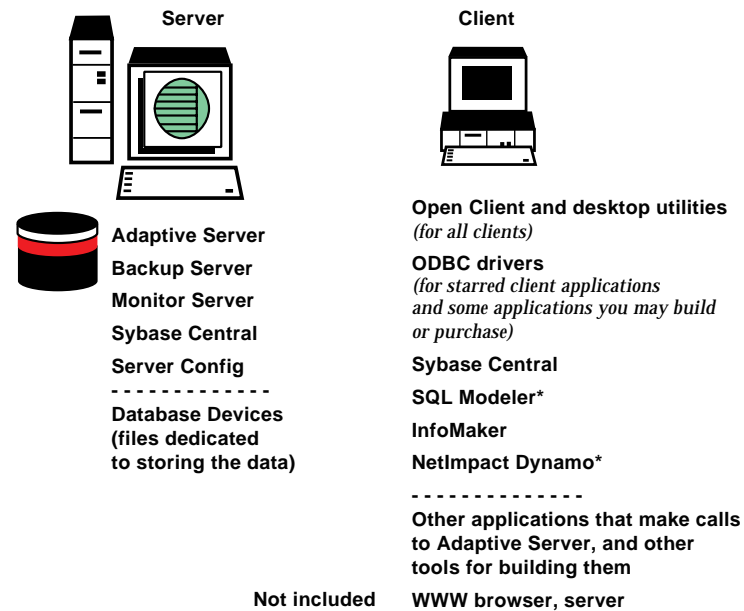


Figure 1-11: Overview of Adaptive Server Enterprise components

Server Components

The following Adaptive Server Enterprise components reside on the server:

- **Adaptive Server**

Adaptive Server is Sybase’s high-performance RDBMS. You can get a technical overview of Adaptive Server release 11.5.x from the product description and the white papers accessible from the Sybase home page on the World Wide Web at the following address:

<http://www.sybase.com>

Also you can ask your sales representative for more information. Adaptive Server also includes several utilities, which are described in *Adaptive Server Utility Programs for Windows NT*.

- **Backup Server™**

Backup Server is a server application that runs concurrently with Adaptive Server to perform high-speed on-line database **dumps** and **loads**. Backup Server is automatically installed when you install Adaptive Server.

- **Server Components of Adaptive Server Monitor™**

Monitor Server and **Historical Server** are the server components of a client/server application called Adaptive Server Monitor, which allows you to capture, display, and evaluate Adaptive Server performance data and to tune Adaptive Server performance. Monitor Server captures performance data from Adaptive Server's shared memory; Monitor Historical Server, writes the data to files for offline analysis. You must install Monitor Server on the same machine as the Adaptive Server that you want to monitor. Monitor Historical Server performs best when installed on a different machine from the Adaptive Server being monitored.

- **Component Integration Services**

Component Integration Services extends Adaptive Server and provides enhanced interoperability. It allows Adaptive Server to present a uniform view of enterprise data to client applications and provides location transparency to enterprise-wide data sources.

- **Server Config**

Server Config facilitates the configuration of Adaptive Server immediately after installation. See "Configuring Adaptive Server" on page 3-2 for a discussion of these post-installation tasks.

- **SQL Remote**

SQL Remote provides message-based data replication between a central consolidated database and a set of databases on laptop or desktop computers at the same or different sites. The remote sites need be connected only occasionally and indirectly to the consolidated database site. The remote computers must be running their own database servers, such as SQL Anywhere.

Client Components

When you install client products on a PC, the installation procedure creates a Sybase program group. This program group contains entries for many of the client components.

The following Adaptive Server Enterprise components are client components:

- **Sybase Central™**

Sybase Central is a Windows application for managing Sybase databases and NetImpact™ Dynamo. It helps you manage database objects and perform common administrative tasks. It makes complex tasks (such as backing up a database) easy.

Sybase Central helps with the following typical system or database administrator's tasks:

- Managing Adaptive Servers, including connecting to, disconnecting from and stopping servers, and troubleshooting Adaptive Server problems
- Managing **data caches**
- Managing Adaptive Server physical resources
- Managing databases, including creating, deleting, backing up, and restoring databases
- Managing access, including creating and deleting Adaptive Server **logins**, creating and deleting database users and user **groups**, administering Sybase **roles**, and managing object and command **permissions**
- Monitoring Adaptive Server performance data and tuning Adaptive Server performance parameters

- **NetImpact Dynamo**

NetImpact Dynamo is a suite of tools for building and managing Web sites linked to Adaptive Server databases. It allows you to store, manage, and access both static HTML and dynamic database data locally or over a corporate **intranet** or the Internet.

NetImpact Dynamo consists of:

- HTML templates that you create with wizards, store in a database, and can embed with SQL queries, Dynamo scripts, and JavaScript scripts. The Dynamo script language is a **server-side script** language that uses the same syntax as JavaScript. It provides programmable control of the content returned to a

Web client. You can use JavaScript to write **client-side scripts** to offload some of the processing to a Web client or to call Java applets.

- The NetImpact Dynamo **application server**, which processes SQL queries and Dynamo scripts. The application server provides an interface between common Web server programs and Adaptive Server. NetImpact Dynamo includes a **CGI** application server, for integration with CGI-compatible servers; an **ISAPI** application server, for integration with Microsoft Internet Information Server or other ISAPI-compatible servers; and an **NSAPI** application server, for integration with the Netscape Web Server or other NSAPI-compatible servers.
- The NetImpact Dynamo personal Web server, which allows you to access your Web site from the same computer you have installed NetImpact Dynamo before you officially deploy it on the Internet.
- The Web site itself, which is stored in a database and maintained through the Sybase Central interface.

You should install NetImpact Dynamo on the same machine on which you install Sybase Central.

- **PowerDesigner™ SQL Modeler™**

PowerDesigner SQL Modeler allows you to design database models and to generate, document, and maintain databases. You can also use PowerDesigner SQL Modeler to reverse-engineer existing databases and re-create them in another RDBMS. For example, you can reverse-engineer a Microsoft Access database and rebuild it in Adaptive Server.

- **InfoMaker™**

InfoMaker is a reporting and datamaintenance tool that helps you make forms, reports, graphs, crosstabs, and tables, as well as simple applications that use these as building blocks. You can also move data between databases using the InfoMaker data pipeline. InfoMaker includes native drivers for connecting to Adaptive Server only. You can connect to desktop databases through the ODBC interface.

- **Open Client**

Open Client contains APIs that enable client applications to interface with Adaptive Server. These APIs are Client-Library™, DB-Library™, and CS-Library. Client applications that you

purchase or develop, as well as application development tools such as PowerBuilder, and InfoMaker, require that these Open Client libraries be installed so that they can communicate with Adaptive Server.

CS-Library contains a collection of utility routines used by all client applications. Client-Library and DB-Library contain a collection of routines that are specific to the programming language being used in an application. Client-Library is an extension of the older DB-Library that accommodates new features supported by Sybase SQL Server releases 10.0 and later; DB-Library supports pre-release 10.0 features. For more information, the documentation in the Open Client/Server™ collection of SyBooks™ (see “Sybase Documentation” on page 1-17 for more information).

Open Client includes Net-Library™, a library containing network protocol services that support connections between client applications and Adaptive Server.

Open Client also includes the following utilities:

- **isql** – an interactive query processor that lets you send commands to the RDBMS from the command line.
- **bcp** – a program that copies data from a database to an **operating system file**, and vice versa.
- **defncopy** – a program that copies definitions of **database objects** that you create from a database to an operating system file, and vice versa.

You install Open Client on all machines that will run client software that accesses Adaptive Server.

- **Desktop Utilities**

Icons for frequently used client utilities appear on the desktop. These include:

- **dsedit** – an editor for creating and modifying the network configuration files. You can also call **dsedit** from Sybase Central. **dsedit** includes a component called **sybping** that tests network connections between client PCs and a server.
- **wdllvers** (Windows only) – a utility for examining Sybase and Windows Dynamic Link Libraries (**DLLs**) loaded into memory.
- **SQL Advantage®** – a utility for executing Transact-SQL® commands and system procedures using a GUI interface.

- **Open Database Connectivity (ODBC) Drivers**

ODBC is an API developed by Microsoft to allow clients to connect to heterogeneous RDBMSs.

PowerDesigner and NetImpact Dynamo require access to Adaptive Server through the ODBC interface. The install program automatically installs ODBC drivers on clients where you install these products. You should also install them on clients where you will be running any third-party products that require the ODBC interface. You need to install Open Client software in addition to the ODBC drivers in order to connect properly.

After you install the ODBC drivers, you create an **ODBC data source** for Adaptive Server in the ODBC Administrator. The ODBC Administrator is a utility that appears on the desktop after you install the drivers. Then you connect through the ODBC interface from the client application.

Sybase Documentation

Except for the books and online help files included in the Adaptive Server Enterprise package, all the product documentation is contained in an online library called SyBooks, which is included on the product media. You can install SyBooks on a single server, where users can access it over the network, or you can install it on individual clients, if they have enough disk space.

If you have Internet access and a Web browser, you can also access SyBooks from the World Wide Web. See “Other Sources of Information” on page xii for more information about SyBooks documentation.

You can also order the documentation in the SyBooks collection in printed form. Refer to the product ID numbers listed on the SyBooks content list included in the product package.

2

Terms and Concepts

Before you delve into Adaptive Server and Adaptive Server documentation, familiarize yourself with some of the terminology that you will frequently encounter as you explore the Sybase family of relational database products.

Chapter 1 defined some terms, including **client/server architecture**, **relational database**, **relational database management system (RDBMS)**, and **table**.

This chapter describes some additional terms and concepts that relate to Adaptive Server, including:

- Database Object 2-1
- Transact-SQL 2-6
- System Database 2-10
- Database Device 2-11
- Permission 2-12
- System Table 2-12
- Utility Programs 2-13
- Transaction Log 2-13
- System Administration 2-13

Database Object

The term **database object** refers to a component of a relational database. The primary database object is the **table**. This section describes other database objects: **rows** and **columns**, **defaults**, **views**, **indexes**, **rules**, **triggers**, and **stored procedures**.

Row and Column

A table contains information about an entity. For example, a table might contain information about authors. A **row**, or record, describes one instance of that entity—a person, a company, a sale, or some other entity. The following example describes a particular author. A **column**, or field, describes an attribute of that entity, such as name, size, color, price, and so on. The example in Figure 2-1 shows

columns that represent the author's first and last names. In early relational systems, columns were called attributes.

The figure illustrates two perspectives of a table named 'authors'. The table has three columns: 'au_id', 'au_lname', and 'au_fname'. The data rows are:

authors		
au_id	au_lname	au_fname
213-46-8915	Green	Marjorie
9948-72-3567	Ringer	Albert
274-80-9391	Straight	Dick

In the left view, the word 'rows' is written vertically to the left of the table, and the entire table is highlighted with a blue border. In the right view, the word 'columns' is written below the table, and the header row and the three data columns are highlighted with a blue background.

Figure 2-1: Rows and columns in a table

Key

In some cases, a column does not describe a real-world attribute but, instead, is a **key** column that serves a logical function, (as introduced in “Tables” on page 1-1). A **primary key** uniquely identifies a row in a table. A **foreign key** relates a row in one table to a row in another table by matching its value to the value of the other table's primary key.

The example shown in Chapter 1 shows three key columns. The *pub_id* column in the *publishers* table is a primary key that provides a unique identifier for every row in the *publishers* table—there are no duplicated *pub_id* values. In the *titles* table, however, the *pub_id* column is a foreign key that relates the information about a book in the *titles* table to information about that book's publisher in the *publishers* table. In the *titles* table, there may be *pub_ids* that appear more than once, because a publisher probably publishes more than one book.

The diagram illustrates two tables: **publishers** and **titles**.

The **publishers** table has a primary key on the **pub_id** column. Its data is as follows:

pub_id	pub_name	city	state
1389	Algodata Infosystems	Berkeley	CA
0736	New Age Books	Boston	MA
0877	Binnet & Hardley	Washington	DC

The **titles** table has a primary key on the **title_id** column and a foreign key on the **pub_id** column. Its data is as follows:

title_id	title	type	pub_id	price
BU1032	The Busy Executive's Database Guide	business	1389	19.99
PS2106	Life Without Fear	psychology	0736	7.00
MC3021	The Gourmet Microwave	mod_cook	0877	2.99
BU7832	Straight Talk About Computers	business	1389	19.99
BU2075	You Can Combat Computer Stress	business	0736	2.99

Figure 2-2: Primary and foreign key columns

You must enforce **referential integrity** for foreign keys. In other words, for every foreign key that does not have a NULL value, there must be a corresponding primary key with the same value. In the example in Figure 2-2, this means that if a *pub_id* in a row in the *titles* table (the “referencing” table) has a value of 1389, there must be a row with a *pub_id* primary key value of 1389 in the *publishers* table (the referenced table). You can enforce referential integrity by specifying **constraints** when you create the table or you can enforce it with **triggers** (see “Trigger” on page 2-6).

A **default** is the value that Adaptive Server assigns to a column when a row is added to a table specifying no particular value for that column. You can specify a default for a column when you create a table. For example, for a column that identifies the price of a published book, you can define “0.00” as the default value until a price is assigned to the book.

View

A **view** is a virtual table, composed of specific columns and/or rows from one or more tables or other views. A view is virtual because it has no physical existence in the database.

You can use a view to simplify and customize a user's perception of a database. For example, sales clerks might need to access some data about authors, but they should not have access to confidential information about an author's royalty advances. You can create a view that customizes the way sales clerks can access the database. Figure 2-3 gives an idea of how a view can select information from a table or tables..

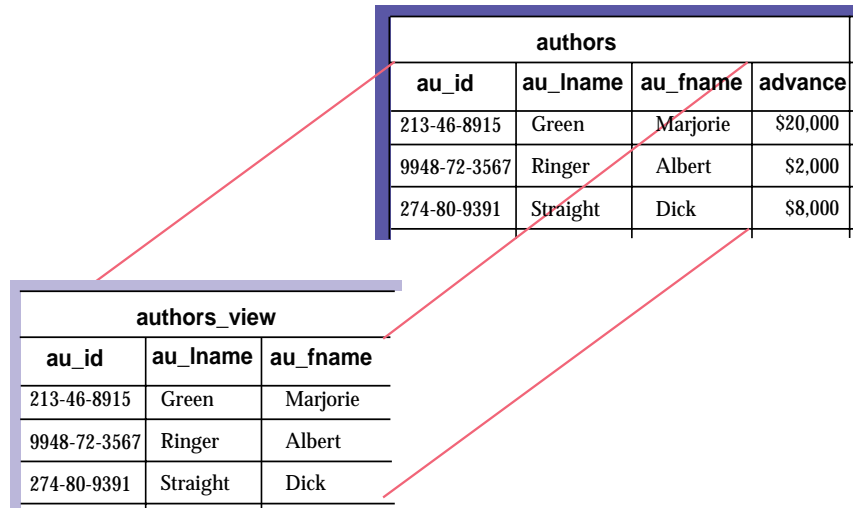


Figure 2-3: A view

Different views can present combinations of data from various tables in different ways, independent of the table structure in which the data is stored. In some cases, data can be directly updated through a view; in other cases, a view is read-only.

Index

An **index** is a database object that helps Adaptive Server locate data. Indexes speed up data retrieval by pointing Adaptive Server to the location of a table column's data on disk.

Conceptually, an index in a database is like an index in a book. In a book, the index relates each indexed term to the page or pages on which that word appears. In a database, the index relates each indexed column value to the **page**, or physical location, at which the row of data containing the indexed value is stored. See Figure 2-4.

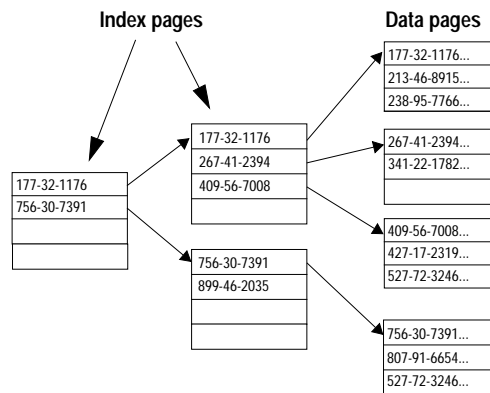


Figure 2-4: A simplified index schematic

Indexes are an important physical design element for high performance. The *Performance and Tuning Guide* explains in detail how indexes work and how to create indexes that improve performance.

Rule

A **rule** defines and enforces which data can be entered for a particular table column or **user-defined datatype**. For example, you can define a rule that requires that all the values in the *title_id* column in the *titles* table begin with two uppercase letters followed by four digits. Or you can define a rule that requires all values of the *money datatype* to be less than \$1,000,000.

Stored Procedure

A **stored procedure** is a set of one or more commands stored in a database. A stored procedure is partially processed before it is stored, so it executes faster than it would if you executed its constituent commands individually.

A stored procedure is invoked by its name. The caller can pass **parameters** to and receive results from a stored procedure. You can create and name your own stored procedures to execute specific database queries and to perform other database tasks. For example, you might create a stored procedure that returns the names of all authors whose books have sold more than the number of books that you specify as a parameter at the time you call the procedure.

You can also use a set of stored procedures supplied by Sybase to help you accomplish numerous administrative tasks. These Sybase-supplied stored procedures are called **system procedures** and are automatically installed when you install Adaptive Server. The system procedures are described in the *Adaptive Server Reference Manual*.

An **extended stored procedure (ESP)** is a stored procedure that executes procedural language code instead of database commands. You invoke the extended stored procedure in the same way as a regular stored procedure. A Sybase-supplied extended stored procedure is called a system extended stored procedure or a system ESP. The system ESPs are described in the *Adaptive Server Reference Manual*.

Trigger

A **trigger** is a special stored procedure that automatically fires whenever a user updates, deletes, or inserts data, depending on how you define the trigger. You associate a trigger with a table or columns within a table.

For example, some organizations keep duplicate data to improve retrieval speed for **decision support** queries. They can define triggers to duplicate data automatically after an update, deletion, and insertion. Managers can query duplicated decision support data without slowing down the data entry system. Another common use of a trigger is to implement some aspects of **referential integrity**.

Transact-SQL

Transact-SQL is Sybase's enhanced version of SQL. It is the language that you, or the client application you are running, uses to communicate requests to the RDBMS. Some of the enhancements that Transact-SQL offers, such as **stored procedures**, **control-of-flow language**, and **error handling** allow Adaptive Server to be a truly programmable database server.

Queries

The “Q” in “SQL” stands for **query**. You query or retrieve data from a database with Transact-SQL’s select command. The basic query operations in a relational system are **selection**, **projection**, and **join**. The select command implements all of them.

The following sections present some simplified scenarios of retrieving data from a database using the select command and show the actual SQL statements you use to accomplish the queries. The purpose of these statements is to familiarize you with what Transact-SQL looks like, not to teach you SQL. For detailed instructions in how to use Transact-SQL, see the *Transact-SQL User’s Guide*.

Selection

A **selection** (also called restriction) is a subset of the rows in a table, based on some conditions. Figure 2-5 shows an example of selecting the rows for books of the type “business” from the *titles* table.

titles				
title_id	title	type	pub_id	price
BU1032	The Busy Executive’s Database Guide	business	1389	19.99
PS2106	Life Without Fear	psychology	0736	7.00
MC3021	The Gourmet Microwave	mod_cook	0877	2.99
BU7832	Straight Talk About Computers	business	1389	19.99
BU2075	You Can Combat Computer Stress	business	0736	2.99

Figure 2-5: Selection

The Transact-SQL select command that retrieves the highlighted rows is:

```
select *
  from titles
 where type like "business"
```

Projection

A **projection** is a subset of the columns in a table. Figure 2-6 shows a sample projection in which you want look only at the title and price of all books in the *titles* table.

titles				
title_id	title	type	pub_id	price
BU1032	The Busy Executive's Database Guide	business	1389	19.99
PS2106	Life Without Fear	psychology	0736	7.00
MC3021	The Gourmet Microwave	mod_cook	0877	2.99
BU7832	Straight Talk About Computers	business	1389	19.99
BU2075	You Can Combat Computer Stress	business	0736	2.99

Figure 2-6: Projection

The Transact-SQL select command that retrieves the highlighted columns is:

```
select title, pub_id
from titles
```

Join

A **join** links the rows in two or more tables by comparing the values in key columns and concatenating rows that have matching values. For example, you might want to select data from the relevant tables showing information about Berkeley publishers and the types and titles of books they published. Figure 2-7 shows a join.

Primary

publishers			
pub_id	pub_name	city	state
1389	Algodata Infosystems	Berkeley	CA
0736	New Age Books	Boston	MA
0877	Binnet & Hardley	Washington	DC

titles				
title_id	title	type	pub_id	price
BU1032	The Busy Executive's Database Guide	business	1389	19.99
PS2106	Life Without Fear	psychology	0736	7.00
MC3021	The Gourmet Microwave	mod_cook	0877	2.99
BU7832	Straight Talk About Computers	business	1389	19.99
BU2075	You Can Combat Computer Stress	business	0736	2.99

Foreign

Figure 2-7: A join

The Transact-SQL select statement that returns the highlighted data from the two tables is:

```
select *
  from publishers, titles
 where publishers.pub_id = titles.pub_id
    and city = "Berkeley"
```

Other Commands

You can do more with Transact-SQL than just query. Transact-SQL includes other commands that create tables and views, such as `create table` and `create view`. It also includes commands that modify tables (the `insert`, `update`, and `delete` commands), and commands that perform many other database tasks discussed in this manual.

Transact-SQL offers extensions to the standard library of SQL commands. We have already looked at one powerful extension, in "Stored Procedure" on page 2-5. Another is **control-of-flow language**, Transact-SQL's programming-like constructs that control the flow of execution of the Transact-SQL statements. When

combined with Transact-SQL's sophisticated error handling techniques, including the ability to capture a **return status** and get reports from **global variables**, they make Adaptive Server a fully programmable database server.

The Transact-SQL commands are described in the *Adaptive Server Reference Manual*. For in-depth discussions of how to use many of the Transact-SQL commands, see the *Transact-SQL User's Guide*.

Presenting Transact-SQL to the RDBMS

There are several ways to present Transact-SQL commands to the RDBMS. One way is to use Sybase's interactive query parser, *isql*. You can invoke *isql* at the operating system command line and then enter Transact-SQL commands at the *isql* prompt. You can also use SQL Advantage, which is a GUI-based query parser, accessible from the Sybase program group or the Utilities section of Sybase Central on Windows or Windows NT clients.

End users typically do not use SQL itself but access the databases using specialized client applications. The client applications send Transact-SQL commands to the RDBMS by passing them as parameters to Open Client library routines (see "Building or Buying the Client Applications" on page 3-10).

System Database

There are two kinds of databases: **system databases** and **user databases**. User databases are databases that you create to store and manage your enterprise data. System databases are databases that the system uses to manage the system.

The *master* database is a system database that is installed with every Adaptive Server. It stores information about all the user databases and other system databases and their associated devices.

The *sybssystemprocs* database is the database that stores all the system procedures provided by Sybase.

The *model* database is a system database that Adaptive Server uses as a template for creating user databases. You can customize the default *model* database that Sybase provides to include your own database objects. For example, if you have created some special stored procedures that you want included in all the user databases that are subsequently created, you can accomplish this by creating the procedures in the *model* database.

The *tempdb* database is a system database that Adaptive Server uses to provide temporary working storage for the objects it creates in the process of executing queries.

The *sybsecurity* database is an optional system database that you can install if you want to use Adaptive Server's audit system. For information about the audit system and the *sybsecurity* database, see the *Security Administration Guide*.

The *sybsyntax* system database is an optional system database that stores the syntax for all the Transact-SQL commands, Sybase system procedures, Adaptive Server utilities, and some Open Client routines. Although the syntax for these facilities is available in the Sybase documentation, if you want to be able to access the syntax quickly online, using the *sp_syntax* system procedure, you must install the *sybsyntax* database.

For more information about the system databases, see the *System Administration Guide*.

Optional Sample Databases

The *pubs2* and *pubs3* databases are optional sample user databases that are used in most of the examples appearing in the Adaptive Server documentation. If you want to try some of the examples that appear in the Adaptive Server documentation, install the *pubs2* or *pubs3* database as described in the *Installing Adaptive Server and OmniConnect on UNIX Platforms*. These databases are described in the *Transact SQL User's Guide*.

The *psdemodb* sample database is used in the examples appearing in the NetImpact Dynamo and InfoMaker documentation. If you want to try the examples for those client products, install the sample database. See *Installing Adaptive Server Enterprise for Windows NT* and the NetImpact Dynamo documentation for instructions on installing the sample database.

Database Device

A **database device** stores Adaptive Server databases. It can be a disk or a portion of a disk or, sometimes, an operating system file.

The system and user databases reside on database devices connected to the server. The *master* database must be installed on a device called the master device, which is automatically created at install time. The

other system databases and the user databases should be created on other devices.

Part of the process of installing Adaptive Server consists of preparing the database devices. This procedure is described in *Installing Adaptive Server and OmniConnect on UNIX Platforms*.

System Table

Every user database contains a set of **system tables**, which are special tables used by the system to manage data and the system. The *master* database contains some special system tables, which do not appear in other databases, that keep track of Adaptive Server as a whole. In other relational database products, what Sybase calls the system tables may be called the data dictionary or the system catalogs.

Adaptive Server Enterprise includes a poster illustrating the contents of and relationships among the system tables. The *Adaptive Server Reference Manual* also describes each system table in detail.

Although you can query the system tables directly to retrieve information about the system, such as the names of all the user databases or all the user accounts on the server, it is usually simpler to use the appropriate system procedure. The system procedures are designed to retrieve the information that a user or System Administrator is likely to need, and they save you the trouble of having to formulate complex queries and become familiar with the details of the system tables on your own.

Permission

A **permission**, sometimes called a privilege, is the authority to perform an action on a certain database object. Adaptive Server permissions are similar to permissions at the operating system level, but Adaptive Server maintains its own permissions for each user, apart from those conferred by the operating system. Examples of Adaptive Server permissions are permission to select data from a table, to modify data in a table, to execute a particular stored procedure, or to create databases or tables.

An Adaptive Server System Administrator or the owner of a particular database is responsible for setting up and maintaining permissions for various users, or groups of users, for Adaptive Server or for particular database.

Utility Programs

Adaptive Server utility programs perform RDBMS tasks. You invoke them from the operating system rather than through `isql` or a similar program.

`isql` is a utility. Another utility is `bcp`, which copies bulk data from an operating system file to a database table and vice versa. See *Adaptive Server Utility Programs* for your platform for information on the Sybase utilities.

Transaction Log

The **transaction log** is a special system table that Adaptive Server maintains for each database to record modifications to the data in the database tables. If a system failure occurs while data is being modified, Adaptive Server uses the transaction log to recover the data and restore the database to the state it was in before the failure occurred. Although the log is part of the database, it should be stored on a separate device to reduce the likelihood of a physical disk error damaging both a database and its associated transaction log.

System Administration

The term **system administration** refers to an assortment of tasks including, but not limited to, managing Adaptive Server's physical storage, creating and backing up databases, creating user accounts, granting permissions, and running diagnostic and repair functions.

The user responsible for system administration is the Adaptive Server **System Administrator**, who may or may not be the same individual or individuals who are system administrators for the server machine's operating system. For Adaptive Server, the System Administrator can log into Adaptive Server as the special, predefined user "**sa**", using the System Administrator password. The "**sa**" account has special privileges not held by normal users. The System Administrator can grant other users permission to perform certain system administration tasks.

Sybase Central allows you to execute the most common system administration tasks using wizards. A wizard is a special kind of dialog box or linked dialog boxes that guide you through a task. Figure 2-8 shows a Sybase Central wizard.



Figure 2-8: A Sybase Central wizard

Some tasks must be executed using Transact-SQL commands or Adaptive Server utilities directly. See Sybase Central's online help for information about administration tasks you can perform from Sybase Central. See the *System Administration Guide* and the *Performance and Tuning Guide* for complete information about Adaptive Server system administration.

3

Using Adaptive Server Enterprise

This chapter briefly describes common tasks that a user performs with a new Adaptive Server. It gives general information about how to begin using this product.

This chapter covers the following topics:

- Installing Adaptive Server 3-1
- Configuring Adaptive Server 3-2
- Designing Databases 3-3
- Creating Databases 3-4
- Loading Data 3-5
- Creating Indexes 3-6
- Backing Up Data 3-7
- Assigning Logins, Users, and Permissions 3-7
- Querying and Modifying Data 3-8
- Building or Buying the Client Applications 3-10
- Setting Up a Web Site 3-12
- Monitoring and Tuning Performance 3-14

For complete technical information about Adaptive Server, see the online documentation in SyBooks. You can also consult some of the excellent books available from Sybase Press, such as *Understanding Sybase System 11: A Hands-On Approach*.

Installing Adaptive Server

You need to install Adaptive Server and the other server components on the server machine, and the client components on one or more client machines. You can install some of the client components on the server, if you find that it more convenient. On both servers and clients, you install the software into a Sybase installation directory that you can designate. You should also install SyBooks in order to have access to the online documentation.

If you have an earlier release of Adaptive Server already installed, you may need to upgrade your existing databases before you can access them with the newly installed release.

To install the software, follow the instructions in *Installing Adaptive Server Enterprise for Windows NT* and the *Release Bulletin*, which are packaged with the software media.

Configuring Adaptive Server

Configuration of Adaptive Server falls into two categories: post-installation configuration and maintenance configuration.

When you install Adaptive Server, the installation program sets configuration parameters to default values. After installing, you need to configure Adaptive Server for your setup. This involves:

- Setting up client/server network connections
- Specifying the location of the **error log**
- Localizing the language, **character set**, and **sort order**
- Setting certain operating system values
- Setting the **System Administrator** (“sa”) password
- Setting up **auditing**, if desired

You can easily configure Adaptive Server with the Server Config utility that appears in the Sybase program group after you install Adaptive Server Enterprise. Figure 3-1 shows the main menu of Server Config and the tasks it helps you accomplish.



Figure 3-1: Configuring Adaptive Server from the Server Config utility

See *Configuring Sybase SQL Server Professional for Windows NT* for detailed information about performing these post-installation configuration tasks.

Maintenance configuration is concerned more with fine-tuning memory and disk storage. See the *System Administration Guide* for information about ongoing Adaptive Server configuration.

To configure Adaptive Server from a Windows client with a graphical interface, use the Sybase Central configuration utility to set configuration parameters and configure network connections.

Choose Configure from the Server pop-up menu.

Those who prefer a command line interface can set configuration parameters by invoking the `sp_configure` system procedure from `isql`. `sp_configure` is fully documented in the *System Administration Guide* and the *Adaptive Server Reference Manual*.

Designing Databases

There are two stages of database design: **logical design** and **physical design**. In logical design, you define the entities that will be represented by tables in the database, choose the attributes of those entities (which will be represented by columns in the table), and determine how the entities are interrelated through the designation of primary and foreign keys.

Logical design can be accomplished independently of the particular relational database vendor that your organization chooses to supply its software. Logical design should rely heavily on input from the intended users of the databases, and not just on database and software experts. Ideally, you will have completed the logical design of your user databases before you purchase and install Adaptive Server.

In the physical design stage, you map the logical design to the Transact-SQL data definition commands that will create the databases on the server. These are commands like `create database`, `create table`, `create view`, `create rule`, and so on.

You can produce a physical design manually or with Sybase Central. If the databases are large and complex, you can use PowerDesigner SQL Modeler, a powerful database design tool included with Adaptive Server Enterprise. PowerDesigner guides you through creating a physical data model from scratch (**PDM**), and it provides a visual representation of the model as you make it. Figure 3-2 shows Power-Designer's visual representation of the PDM it created for *pubs2*, the sample database included with Adaptive Server Enterprise (see "Optional Sample Databases" on page 2-11).

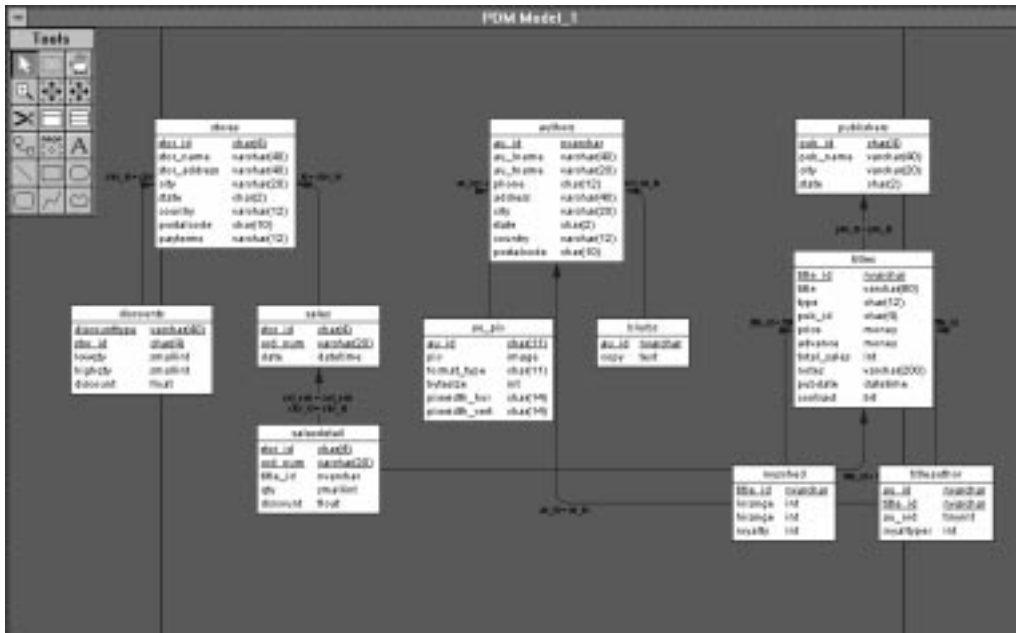


Figure 3-2: PowerDesigner PDM of the pubs2 database

You can also reverse-engineer a PDM from an existing database. PowerDesigner can connect to an existing database and create a PDM for it that you can use as is or modify to create an Adaptive Server database.

You can also define what are called **extended attributes** for the database. Extended attributes provide additional information to external client applications such as PowerBuilder. For more information on extended attributes, see the PowerDesigner documentation in SyBooks.

It is a good idea to archive the PDM you create for your database within PowerDesigner. You can use archived PDMs to modify existing databases.

Creating Databases

The output of the physical design effort is a collection of Transact-SQL commands that create databases and, create database objects within the database. To create the databases, you execute these create commands on the server.

If you use Sybase Central to design your databases, Sybase Central itself sends the create commands to the server. If you use PowerDesigner, after you create the PDM, you can generate the database directly from Power-Designer or from a database creation script that you can run at a later time from Power-Designer, Sybase Central, or the command line.

It is also possible, of course, to create databases at the `isql` command line using the Transact-SQL create commands. Since most databases consist of several tables, if you use the command line method, it is a good idea to save the original data definition commands in a script file that can be run to re-create consistent images of the database as many times as necessary.

Loading Data

After the database tables have been created, you can fill them with data. You may have to insert the data manually, using either the Transact-SQL insert command or a client application that has been written to input the data. In many cases, however, the data already exists either in another vendor's database system or in **operating system files**.

If the data is in another database, such as SQL Anywhere™ or Microsoft Access, you can use InfoMaker's Data Pipeline Painter to access the source database through the ODBC interface and move the data into the new database. The Data Pipeline Painter can reproduce data within a database, across databases, and across RDBMSs. When combined with Power-Designer's reverse-engineering capability, you have all the tools you need to convert existing databases for Adaptive Server as quickly as possible. Figure 3-3 illustrates how data pipelines are created in InfoMaker.



Figure 3-3: Creating a data pipeline in InfoMaker

If the data is in a file, or can be converted to a file, you can copy it into a database table using the `bcp` bulk copy utility. See the *Adaptive Server Utility Programs* for your platform for information on how to use `bcp`. You can also import data from a file in InfoMaker's Data Manipulation painter.

If the data is in a Sybase database from an earlier release, the original tables may be accessible after they have been upgraded. See the description of the upgrade process in the installation guide to determine whether the databases can be upgraded.

Creating Indexes

If you are going to load a large amount of initial data into a table after creating it, it is best to create the table indexes after you load the data, because indexes slow performance in commands that insert data. Also, for every row added to the table, a row must also be added to the index. However, well-chosen indexes greatly improve performance on commands that retrieve data, so it is advisable to create indexes on the tables that are frequently queried.

When you create an index on a table, you have to decide on which table column or combination of columns the index will be based. You also have to decide what type of index to create and whether it should be unique. See `create index` in the *Adaptive Server Reference Manual* and the chapter on indexes in the *Performance and Tuning Guide* to learn about selecting useful indexes.

You can create indexes using the Sybase Central Index Creation wizard or at the `isql` command line using the Transact-SQL `create index` command. You can also add an index to an existing database by generating a database alter script with Power-Designer. You can also modify the database directly from Power-Designer.

Backing Up Data

After you input data and create indexes on the tables, you should back up the databases so that you can restore them in the event of a system failure. You should continue to back up data at regular intervals. How often you back up a database depends on how frequently the data is updated and how costly it would be for you to lose it.

Databases that are constantly being modified need to be backed up frequently; databases containing relatively static data can be backed up less frequently. You should regularly back up production databases, but you can usually be less diligent about backing up databases used for test and development purposes, if you can afford to lose their most recent data.

If you experience a **media failure**, you can restore the data from the most recent backups, using Adaptive Server's `restore` command.

You do not have to shut down Adaptive Server to back up or restore a database. Backup Server makes it possible to perform online backup and restore operations while Adaptive Server is running.

You can back up and restore databases and their transaction logs by using the Sybase Central Backup and Restore wizards or by using the Transact-SQL `dump` and `load` commands or at the `isql` command line.

See the *System Administration Guide* for information about developing a backup plan and about backing up and restoring user and system databases.

Assigning Logins, Users, and Permissions

Adaptive Server security is based on a three-tiered system that separates server access, database access, and data access.

To connect to a server, a user must have a **login** on the server. A login does not automatically give the person permission to access data or create database objects, just to connect. It consists of a login name and a password.

To give others access to a server, the Sybase System Administrator creates logins using either the Sybase Central Add Login wizard or the `sp_addlogin` system procedure from `isql`.

To enable a user to connect to a database, the System Administrator gives the user a database **user ID**. That database username may or may not be the same as the user's login name. The System Administrator or the **Database Owner** (the user who created the database or was given ownership by the database creator) create database users. You can create users with either the Sybase Central Add User wizard or the `sp_adduser` system procedure from `isql`.

In addition to individual users, the System Administrator or Database Owner can create a **group** of database users in Adaptive Server. The group is identified by a group name and provides a convenient way to grant and revoke permissions to more than one user in a single statement. For example, you can define a "managers" group or a "payroll" group, assign users to the group, and assign the same privileges to all the users in the group by assigning the privileges to the group as a whole. There is also a system-defined group, named "all," to which all database users belong.

You can create a group using the Sybase Central Create Group wizard or the `sp_addgroup` system procedure from `isql`. When you add a user to the database, you can designate the groups that the user is a member of as one of the options in the `sp_adduser` system procedure or Sybase Central Add User wizard.

Having a login and username gives an individual access to a server and a database, but it does not give automatic permission to select data from tables or views, modify data in tables, or execute stored procedures. To give a user permission to access a database object, the System Administrator or the Database Owner must explicitly grant the user, or a group to which the user belongs, permission to access that object. The System Administrator or Database Owner can grant and deny various permissions to other users to use specific database objects by either using Sybase Central or the Transact-SQL `grant` and `revoke` commands from `isql`.

Querying and Modifying Data

Most activity between a user and Adaptive Server involves either querying data or modifying data. Querying data is discussed in "Queries" on page 2-7.

Modifications

Modifying data means:

- Adding a new row to a table (the `insert` command or copying in data with the `bcp` utility),
- Deleting a row from a table (the `delete` command), or
- Changing the value of one or more columns in a row (the `update` command).

Transactions

If you are executing multiple data modification commands on the same or related data, you can execute them together in a single **transaction**. A transaction consists of multiple Transact-SQL commands, enclosed between begin and end delimiters, that are executed as an **atomic** command rather than as separate commands. The tables affected by the transaction are locked so that other users cannot make changes that affect the work being accomplished during the transaction. The transaction as a whole is either entirely committed or not committed, ensuring that if the system fails in the middle of an operation, the database is not left in an inconsistent state.

The classic example of a transaction is one that removes money from your savings account (`delete`) and adds it to your checking account (`insert`). You would want to perform the `delete` and the `insert` in a single transaction, so that if the system failed after the `delete` but before the `insert`, the record of the funds would not be lost. If such a failure occurred, the transaction would be rolled back by Adaptive Server's recovery program and the database would appear as though the `delete` had never been executed.

Stored Procedures

If you find that there are commands, or sets of commands, that users will perform frequently, you should consider creating a **stored procedure**. Even for a single command, a stored procedure usually executes faster because the stored procedure stores the command in a partially processed form. Also, you can control access more securely by giving users permission to execute the stored procedure without giving them permission to access all the objects referenced

by the commands in the stored procedure. See the *Transact-SQL User's Guide* for more information on creating stored procedures.

Building or Buying the Client Applications

You have a number of options for establishing a client interface through which Adaptive Server users can access data. You can purchase one of many third-party applications, you can use InfoMaker as a simple kind of client interface, or you can build your own.

If your organization is very small, and your information needs are simple, you may be able to manage by using a set of Transact-SQL statements and creating some stored procedures. Users could perform queries through an isql or SQL Advantage client and direct the results of a query to a file.

If your needs are more complex, and you want to provide a customized interface to users who may not know anything about relational databases or Transact-SQL, you can purchase a client application that has been designed for a particular purpose and runs against an Adaptive Server database from a third-party software company or a VAR (Value Added Reseller). Look for the "Industry Solutions" page in the "Partners and Solutions" section of the Sybase Web site for up-to-date news and reference information about industry-specific applications that run against Adaptive Server.

InfoMaker itself is a GUI client application that may suit most of your needs. You can use InfoMaker to access your Adaptive Server data and to make forms, sophisticated reports, graphs, cross-tabs, and labels. Figure 3-4 shows a very simple InfoMaker form set up against the *pubs2* sample database. First, you design the form and use the built-in tools to define the queries that will populate it. After you run the data through the form, you can scroll through the records and you can enter and update data in the InfoMaker form window.



The screenshot shows a window titled "pubs? Authors" with a form containing the following fields and buttons:

Field Label	Value
Au Id	12345678
Au Name	White
Au Phone	Johnson
Phone	408-896-7223
Address	10000 Biggs Rd
City	Menlo Park
State	CA
Country	USA
Postcode	94025

Buttons: Add an entry, Delete an entry, Update Entry

Figure 3-4: A simple InfoMaker form

You can build a customized, user-friendly client application with a **development tool**. A development tool helps you create GUI applications that send Transact-SQL commands to access and manipulate Adaptive Server data. InfoMaker itself is a development tool. You construct an InfoMaker application, complete with toolbars and menus, from a collection of reports and forms that you design. . You can then run the standalone application from any client desktop, outside the InfoMaker environment.

You can purchase (separately) a development tool, such as PowerBuilder, which creates robust standalone GUI client applications. PowerBuilder takes advantage of the extra functionality of any extended attributes you may have defined in Power-Designer.

After building an application, you deploy it on the client, where it is linked with the Open Client libraries. Typically, application-building tools have their own simplified programming APIs and a graphical interface for program development. Figure 3-5 shows the PowerBuilder development window.

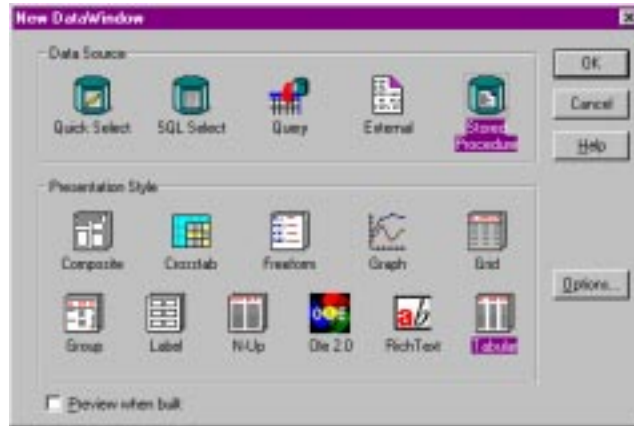


Figure 3-5: PowerBuilder development window

You can also design client applications from scratch using a traditional procedural programming language, such as C, C++, or COBOL. This method usually requires more time and more skilled software engineers than are needed to develop an application using a development tool.

The advantage of the second approach is that the applications are generally smaller, and sometimes faster, because the executable does not have the overhead of the tool. These programs make direct calls to the Open Client library routines to communicate with Adaptive Server. See the Open Client documentation in the Open Client/Server collection in SyBooks for complete information about coding with the Open Client libraries.

Setting Up a Web Site

If you decide you want to connect your data to the Internet or to a corporate intranet, you can use NetImpact Dynamo to create an interactive Web site. You should install the Web server software and NetImpact Dynamo software on the same machine on which you are running Sybase Central. Adaptive Server itself can be on a different machine and can be accessed through ODBC drivers installed on the client.

You operate NetImpact Dynamo and create your Web site through the Sybase Central interface. When you first connect to your database, NetImpact Dynamo creates tables in the database that

allow it to store the Web site itself in the database. As you work, NetImpact Dynamo stores all the site-related documents in the database. It also creates a visual representation of the site that makes it easy to work with, displaying the scripts, templates, static HTML files, and images in a familiar folder-based hierarchical structure (see Figure 3-6). Dynamo scripts have a *.ssc* extension and templates have a *.stm* extension, instead of a *.html* extension.

Name	Connection	Size	Comment
Add template			
Add folder			
Add script			
Add linked folder			
app	<default>		
doc	<default>		
images	<default>		
system	<default>		
contents.html	<default>	867	NetImpact Dynamo
mine.stm	<default>	1,453	
newquery.stm	<default>	667	
Products.stm	<default>	543	sample
query.stm	<default>	1,359	

Figure 3-6: NetImpact Dynamo Web site displayed in Sybase Central

The visual representation of the site also includes some of the wizards available for creating new templates and scripts, which contain programmatic instructions for merging static HTML with the results of embedded SQL queries. You can also import existing files and images into your site through the Sybase Central interface.

While you are creating the site, you can connect to it through the NetImpact Dynamo personal Web server and view it locally in a Web browser, such as Netscape. When you are ready to deploy the site, you connect the site to the Web server. Then, the site can be accessed anywhere on the Internet or from a corporate intranet.

Monitoring and Tuning Performance

After building a production database, you may want to examine some performance statistics and adjust some performance parameters to tune the system for high **performance**.

The easiest way to monitor Adaptive Server performance is with Adaptive Server Monitor Server. Monitor Server gives specific details about performance with numeric measurements of cache usage, network traffic, **device I/O**, and **locking** activity, all of which are discussed in the *Performance and Tuning Guide*. You view statistics collected by Monitor Server using Monitor Viewer, which is available from Sybase Central.

Most of your tuning efforts ought to address the amount of time it takes for Adaptive Server to respond to queries. You can achieve high performance by starting out with a well-designed database that includes strategic indexes and by learning to work with the Adaptive Server query **optimizer**, an Adaptive Server feature that analyzes queries and database objects and selects the appropriate query plan, based on how much time it will take to execute. For more information, see the *Performance and Tuning Guide*.

Glossary

Adaptive Server Enterprise

The server in Sybase's client/server architecture. Adaptive Server Enterprise manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

APIs

Application Program Interfaces, subroutines that allow client applications to interface with Adaptive Server Enterprise. Also known as **libraries**.

aggregate function

A function that works on a set of cells to produce a single answer or a set of answers, one for each subset of cells. The aggregate functions available in Transact-SQL are: average (*avg*), maximum (*max*), minimum (*min*), sum (*sum*), and count of the number of items (*count*).

application-building tool

See **development tool**.

application server

The programmable part of NetImpact Dynamo that functions like a CGI program and that serves as an interface between a **web server** and Adaptive Server Enterprise.

argument

A value supplied to a function or procedure that is required to evaluate the function.

auditing

Recording security-related system activity that can be used to detect penetration of the system and misuse of system resources.

automatic recovery

A process that runs every time Adaptive Server is stopped and restarted. The process ensures that all transactions that completed before the server went down are brought forward and that all incomplete transactions are rolled back.

backup

A copy of a database or transaction log used to recover from a **media failure**.

business rules

The rules to which the data in a database must conform.

cache

See **data cache**.

CGI

Common Gateway Interface. A standard that allows Web servers to interface with and pass information on to external applications such as search engines and database servers.

character set

A set of specific (usually standardized) characters with an encoding scheme that uniquely defines each character. ASCII and ISO 8859-1 (Latin 1) are two common character sets.

client

The user's side of a client/server arrangement; can refer to the software that is making the calls to the server or to the machine that is running the client software.

client-side script

In a World Wide Web application, a script interpreted by a web client and embedded directly in an HTML page. Compare to **server-side script**.

client/server architecture

A computer system architecture in which clients request a service and a server provides that service. Each machine can specialize in the tasks it is best suited for.

clustered index

An **index** in which the physical order and the logical (indexed) order is the same. A table can have only one clustered index.

column

A data value that describes one characteristic of an **entity**. Also called a **field**.

command

An instruction that specifies an operation to be performed by the computer. Each command or SQL statement begins with a keyword, such as **insert**, that names the basic operation to be performed.

consolidated database

A database that serves as the “master” database in a hierarchical data **replication** setup. The consolidated database contains all the data to be replicated; while its remote databases can contain only their own subsets of the data. In case of conflict or discrepancy, the consolidated database is considered to have the primary copy of all data.

constraint

A rule applied to a database object that ensures that all entries in the database object to which it applies satisfy a particular condition. For example, a column may have a constraint requiring that all values in the column be unique.

control-of-flow language

Transact-SQL’s programming-like constructs (such as **if**, **else**, **while**, **goto label**) that control the flow of execution of Transact SQL statements.

data cache

Also called named cache or cache. A cache is an area of memory within Adaptive Server that contains the in-memory images of database **pages** and the data structures required to manage the pages.

data definition

The process of setting up databases and creating database objects such as tables, indexes, rules, defaults, procedures, triggers, and views.

data dictionary

The system tables that contain descriptions of the **database objects** and how they are structured.

data integrity

The correctness and completeness of data within a database.

data modification

Adding, deleting, or changing information in the database with the **insert**, **delete**, and **update** commands.

data retrieval

Requesting data from the database and receiving the results. Also called a **query**.

database

A set of related data tables and other database objects that are organized and presented to serve a specific purpose.

database administration

The tasks involved in maintaining, designing, implementing changes to, tuning, and expanding a database. See also **system administration**.

database device

See **device**.

database object

One of the components of a database: **table, view, index, stored procedure, trigger, column, default, or rule**.

Database Owner

The user who creates a database. A Database Owner has control over all the database objects in that database. The login name for the Database Owner.

datatype

Specifies what kind of information each column will hold, and how the data will be stored. Datatypes include *char*, *int*, *money*, and so on. Users can construct their own datatypes, based on the Adaptive Server system datatypes.

dbo

In a user's own database, Adaptive Server recognizes the user as "dbo." A **Database Owner** logs into Adaptive Server using his or her assigned login name and password.

decision support

Data used to help managers make business decisions. Usually involves a lot of queries, as opposed to updates.

default

The option chosen by the system when no other option is specified.

default database

The database that a user connects with when he or she logs in.

development tool

Software such as PowerBuilder that helps you build specialized GUI applications for accessing Adaptive Server Enterprise databases.

device

Any piece of disk (such as a partition) or a file in the file system that you specially prepare for storing databases and their objects.

device I/O

Reads or writes to or from a database device.

disk initialization

The process of preparing a database device or a file for Adaptive Server use. Once the device is initialized, it can be used for storing databases and database objects. The command used to initialize a database device is **disk init**.

DLLs

Dynamic Link Libraries. Software used by Microsoft Windows and IBM OS/2 to provide services to applications.

dump

To make a backup of an entire database, including both the data and the **transaction log**, which you accomplish with the **dump database** command. Also, the data that results from this action.

entity

The basic unit described by tables in a relational database. Identifying them is the first step in the **logical design**.

error handling

Techniques available to Transact-SQL programmers for basing code on and displaying errors and error messages.

error log

A file that stores severe error messages and results about the start-up and recovery of databases on a server being started up.

error message

A message issued by Adaptive Server, usually to the user's terminal, when it detects an error condition.

field

A data value that describes one characteristic of an **entity**. Also called a **column**.

foreign key

A key column in a table that logically depends on a **primary key** column in another table. Also, a column (or combination of columns) whose values are required to match a primary key in some other table.

global variable

System-defined variables that are updated by Adaptive Server on an ongoing basis. For example, @@error contains the last error number generated by the system.

group

Uniquely named set of users assigned a set of permissions for the objects and operations within a database.

guest

Default user in the *model* database that allows any user with a valid Adaptive Server login to use that database, with limited privileges.

HTML

HyperText Markup Language. The language with which World Wide Web documents are formatted. It defines basic details of the document, such as paragraph tags, graphics, and hypertext links.

HTTP

HyperText Transfer Protocol. The protocol that negotiates document delivery to a **web client** from a **web server**. When a web client requests an **HTML** document from a web server, the server obeys the request, transfers the document, and closes the connection.

index

A database object that consists of key values from the data tables and pointers to the pages that contain those values. Indexes speed up access to data rows by pointing Adaptive Server Enterprise to the location of a table column's data on disk.

integrity rules

Rules that describe how data will be kept accurately and consistently in the relational model.

intranet

A private network that uses Internet software and standards. It allows you to use a web server to publish information within the organization served by the private network.

ISAPI

Internet Server API. A web server programming interface developed by Microsoft.

JavaScript

A cross-platform World Wide Web scripting language. You can embed client-side JavaScript scripts into NetImpact Dynamo templates to allow the web client to handle some processing tasks or to call Java applets.

join

A basic operation in a relational system that links the rows in two or more tables by comparing the values in specified columns.

key

A field used to identify a record, often used as the **index** field for a table.

key value

Any value that is indexed.

keyword

A word or phrase that is reserved for exclusive use by Transact-SQL. Also known as a reserved word.

libraries

See **APIs**.

load

To restore data stored in a backup created during a **dump**.

local variables

User-defined variables defined with a **declare** statement.

locking

The process of restricting access to resources in a multiuser environment to maintain security and prevent concurrent access problems. Adaptive Server automatically applies locks to tables or **pages**.

logical design

Defining the tables, relations, and keys of a relational database. Distinguished from **physical design**.

logical key

The primary, foreign, or common **key** definitions in a database design that define the relationship between tables in the database. Logical keys are not necessarily the same as **physical keys** (the keys used to create indexes) on the table.

logical read

The process of accessing a data page or an index page that is already in memory, to satisfy a query. Compare to **physical read**.

login

The name a user uses to log into Adaptive Server. A login is valid if Adaptive Server has an entry for the user in the system table *syslogins*.

MAPI

Microsoft's Message Application Programming Interface, a message system used in several popular e-mail systems such as Microsoft Mail.

Master database

Controls the user databases and the operation of Adaptive Server Enterprise as a whole. Known as *master*, it keeps track of such things as user accounts, ongoing processes, and system error messages.

master table

A table that contains data on which data in another table logically depends. The detail table typically has a foreign key that joins to the primary key of the master table.

media failure

A media failure occurs when the information on a medium (typically a hard disk drive) becomes unusable.

messages

Message-based communication between applications or computers does not require a direct connection. Instead, a message agent sent at one time by an application can be received at another time by another application.

message system

In SQL Remote™ replication, a protocol for exchanging messages between the consolidated database and a remote databases. Usually, a consolidated database and a remote database send and receive messages using the same messaging system.

message number

The number that uniquely identifies an error message.

model database

A template for new user databases. The installation process creates *model* when Adaptive Server is installed. Each time the `create database` command is issued, Adaptive Server makes a copy of *model* and extends it to the size requested, if necessary.

nested queries

select statements that contain one or more subqueries.

nonclustered index

An **index** that stores key values and pointers to data.

normalization rules

The standard rules of database design in a relational database management system.

NSAPI

Netscape Server API. A programming specification for Netscape's web servers to access back-end applications.

null

Having no explicitly assigned value. NULL is not equivalent to zero or to blank. A value of NULL is not considered to be greater than, less than, or equivalent to any other value, including another NULL value.

object permissions

Permissions that regulate the use of certain commands (data modification commands, plus `select`, `truncate table` and `execute`) to specific tables, views or columns.

objects

See **database object**.

ODBC

The Open Database Connectivity (ODBC) interface, a standard interface to database management systems in the Windows and Windows NT environments.

ODBC data source

A database that can be accessed through the **ODBC** interface and for which the appropriate drivers have been installed on the local client. In order to connect to the database, you configure the connection information in the desktop application called the ODBC Administrator.

operating system

A group of programs that translates your commands to the computer, so that you can perform such tasks as creating files, running programs, and printing documents.

operating system file

Collection of data named and recognized by the operating system. Adaptive Server Enterprise data is not stored in so-called operating system files, but it can be exported to operating system files with the Sybase utility, **bcp**.

optimizer

Adaptive Server Enterprise code that analyzes queries and database objects and selects the appropriate query plan. It estimates the cost of each permutation of table accesses in terms of how much processing and disk writing time it will take.

page

A 2K block of data that is the minimal unit that can be read from or written to disk.

parameter

An **argument** to a **stored procedure**.

PDM

Physical Data Model. The PowerDesigner model that specifies a physical implementation of a database design and that is represented in a graphic format. You can create a PDM from scratch or by reverse-engineering an existing database.

PERL

Practical Extraction and Reporting Language. An interpreted scripting language, typically used in writing CGI scripts.

performance

The speed with which Adaptive Server processes queries and returns results. Performance is affected by several factors.

permission

The authority to perform certain actions on certain database objects or to run certain commands.

physical design

Mapping the **logical design** to the Transact-SQL data definition commands that create databases on the server.

physical key

A column name, or set of column names, used in a create **index** statement to define an index on a table. Physical keys on a table are not necessarily the same as the **logical keys**.

physical read

A disk I/O to access a data, index, or log page. Adaptive Server Enterprise estimates physical reads and logical reads when optimizing queries. See **logical read**.

primary key

The column or columns whose values uniquely identify a row in a table.

privilege

The authority to perform certain actions on certain database objects or to run certain commands. Synonymous with **permission**.

projection

One of the basic query operations in a relational system. A projection is a subset of the columns in a table.

query

1. A request for the retrieval of data with a **select** statement.
2. Any SQL statement that manipulates data.

query plan

The ordered set of steps required to carry out a query, complete with the access methods chosen for each table.

relational database management system

A system for storing and retrieving data from two-dimensional tables in which **SQL** use is standard.

recovery

The process of rebuilding one or more databases from database dumps and log dumps.

referential integrity

The rules governing data consistency, specifically the relationships between the primary keys and foreign keys of different tables. Adaptive Server addresses referential integrity with user-defined **triggers** and with referential integrity **constraints**.

relational database

A collection of tables that stored interrelated data.

relationship

Describes how entities are related. A basic step in the logical design of a database is to identify the relationships between the entities you have identified.

remote database

In SQL Remote replication, a database that exchanges replication messages with a consolidated database. Remote databases may contain all or some of the data in the consolidated database.

remote procedure calls

A **stored procedure** executed on a different Adaptive Server Enterprise from the server the user is logged into.

replication

For databases, a process by which the changes to data in one database (including creation, updating, and deletion of records) are also applied to the corresponding records in other database.

restriction

A subset of the rows in a table. Also called a **selection**, it is one of the basic query operations in a relational system.

return status

A value that indicates that the procedure completed successfully or indicates the reason for failure.

roles

Provide individual accountability for users performing system administration and security-related tasks in Adaptive Server. The **System Administrator**, System Security Officer, and Operator roles can be granted to individual server login accounts.

rollback transaction

A Transact-SQL statement used with a user-defined **transaction** (before a commit transaction has been received) that cancels the transaction and undoes any changes that were made to the database.

row

A set of related **columns** that describes a specific entity. Also called a record.

rule

A specification that controls what data may be entered in a particular column, or in a column of a particular user-defined datatype.

sa

The login name for an Adaptive Server **System Administrator**.

savepoint

A marker that the user puts inside a user-defined **transaction**. The user can later use the **rollback transaction** command with the savepoint name to cancel any commands up to the savepoint, or **commit transaction** to actually complete the commands.

selection

A subset of the rows in a table. Also called a **restriction**, it is one of the basic query operations in a relational system.

server-side script

In a World Wide Web application, a script interpreted on the WWW server before being sent to the web client. Compare to **client-side script**.

server user ID

The ID number by which a user is known to Adaptive Server Enterprise.

sort order

Used by Adaptive Server to determine the order in which to sort character data.

SQL

Structured Query Language (SQL), the language used to communicate with a relational database and that is the subject of standards set by several standards bodies.

SQL Remote

An asynchronous message-based replication system for two-way server-to-laptop, server-to-desktop, and server-to-server replication.

statement

A statement begins with a **keyword** that names the basic operation or command to be performed.

store-and-forward

An exchange of information, typical of message-based systems, that allows information to be exchanged without a direct connection between applications.

stored procedure

A collection of SQL statements and optional control-of-flow statements stored under a name. Adaptive Server-supplied stored procedures are called **system procedures**.

subquery

A select statement that is nested inside another select, insert, update or delete statement, or inside another subquery.

system administration

An assortment of tasks including but not limited to managing Adaptive Server's physical storage, creating and backing up databases, creating user accounts, granting permissions, and running diagnostic and repair functions. See also **database administration**.

System Administrator

A user who is authorized to handle Adaptive Server system administration, including creating user accounts, assigning permissions, and creating new databases.

system databases

The databases on a newly installed Adaptive Server: *master*, which controls user databases and the operation of the Adaptive Server; *tempdb*, used for temporary tables; *model*, used as a template to create new user databases; and *sybsystemprocs*, which stores the system procedures. See also **user databases**.

system procedures

Stored procedures that Adaptive Server supplies for use in system administration. These procedures are provided as shortcuts for retrieving information from the system tables or mechanisms for accomplishing database administration and other tasks that involve updating system tables.

system table

One of the data dictionary tables. The system tables keep track of information about Adaptive Server Enterprise as a whole and about each user database. The *master* database contains some system tables that are not in user databases.

table

A collection of **rows** (records) that have associated **columns** (fields). The logical equivalent of a database file.

temporary database

The database in Adaptive Server Enterprise called *tempdb* that provides a storage area for temporary tables and other temporary working storage needs.

throughput

The volume of work completed in a given time period. It is usually measured in transactions per second.

Transact-SQL

The SQL dialect used in Adaptive Server Enterprise.

transaction

A grouping of a series of Transact-SQL statements so that they are treated as a single unit of work. Either all statements in the group are executed or no statements are executed. The tables queried during the transaction are locked until the transaction is complete.

transaction log

A **system table** (*syslogs*) in which all changes to the database are recorded.

trigger

A special form of **stored procedure** that goes into effect when a user gives a change command such as **insert**, **delete**, or **update** to a specified table or column. Triggers are often used to enforce referential integrity.

unique indexes

Indexes that do not permit any two rows in the specified columns to have the same value. Adaptive Server checks for duplicate values when you create the index (if data already exists) and each time data is added.

URL

Uniform Resource Locator. A standardized character string that identifies the location of a World Wide Web document.

update

An addition, deletion, or change to data, involving the insert, delete, truncate table, or update statements.

user database

A database created by a user. See also **system databases**.

user-defined datatype

A definition of the type of data a column can contain that is created by the user. These **datatypes** are defined in terms of the existing system datatypes. Rules and defaults can be bound to user-defined datatypes (but not to system datatypes).

user ID

The ID number by which a user is known in a specific database. Distinct from **server user ID**.

variable

An entity that is assigned a value. Adaptive Server has two kinds of variables, **local variables** and **global variables**.

view

A named select statement that is stored in the database as an object. It allows users to “view” a subset of rows or columns from one or more tables.

web client

Also known as a web browser. Client software that requests and displays **HTML** documents and other Internet or **intranet** resources.

web server

A computer that serves data over the World Wide Web. Also, an application that receives and serves data in the **HTTP** protocol over the Internet or a corporate **intranet** and turns a computer into a web server. Sometimes called an HTTP server or an HTTP daemon.

Index

A

- Accessing databases
 - with client applications 2-10
- Adaptive Server System
 - Administrator 2-12
- Administrative tasks
 - accomplishing with stored procedures 2-6
- APIs (application programming interfaces) 1-4, 1-15
- Application-building tools. *See* Development tools
- Application server (Net Impact Dynamo) 1-10, 1-15
- Assigning logins 3-7
- Assigning permissions 3-7
- Assigning users 3-7
- Attributes
 - old name for columns 2-2
- Auxiliary facilities 1-4
- Avoiding redundancy 1-3

B

- Backing up data 3-7
 - based on how often data changes 3-7
- Backup Server 3-7
 - functionality 1-13
 - restoring data 3-7
- bcp utility 1-16, 2-13, 3-6

C

- CGI application server
 - NetImpact Dynamo as 1-15
- CGI programs 1-10
- Client
 - typical request for service 1-5
- Client/server
 - advantages of 1-6

- architecture 1-4
- basic interaction 1-5
- basic relationship 1-5
- division of tasks 1-6
- supplemented with Internet data distribution 1-8
- Sybase's configuration 1-5
- Client applications
 - building vs. buying 3-11
 - designing 3-12
 - interfacing with Adaptive Server 1-15
 - isql as 3-10
 - purchasing 3-10
- Client-Library 1-16
- Columns 2-1
- Component Integration Services 1-13
- Components of product 1-11
 - client 1-14
 - overview 1-12
 - server 1-12
- Configuration parameters
 - default 3-2
- Configuring Adaptive Server 1-13, 3-2
 - maintenance tasks 3-3
 - post-installation tasks 3-2
 - with command-line interface 3-3
 - with Server Config utility 3-2
 - with Sybase Central 3-3
- Connecting to Adaptive Server through ODBC interface 1-17
- Control-of-flow language 2-9
- Copying data 1-16
 - create group command 3-8
 - create index command 3-6
 - create table command 2-9
 - create view command 2-9
- Creating databases 3-4
- Creating forms with InfoMaker 3-10
- Creating groups
 - with Sybase Central 3-8
- Creating indexes 3-6

- after importing data 3-6
- at the command-line 3-7
- deciding on the columns to index 3-6
- with Sybase Central 3-7

Creating users
with Sybase Central 3-8

CS-Library 1-16

D

Database design

- logical 1-3, 3-3
- physical 3-3

Database devices

- defined 1-11, 2-11
- preparing 2-11

Database objects

- columns 2-1
- defaults 2-3
- defined 2-1
- indexes 2-4
- rows 2-1
- rules 2-5
- stored procedures 2-5
- tables 2-1
- triggers 2-6
- types of 2-1
- views 2-4

Database Owners

- creating database users 3-8

Databases

- creating 3-4
- creating from PowerDesigner 3-5

Data definition commands

- saving 3-5

Data dictionary. *See* System tables

Data pages 2-5

Data replication 1-7

DB-Library 1-16

Defaults 2-3

- example 2-3

defncopy utility 1-16

delete command 3-9

Desktop utilities 1-16

Development tools 1-11

- InfoMaker as 3-11
- PowerBuilder 3-11

Distributed computing 1-6

DLLs 1-16

Documentation 1-17

- accessing on the World Wide Web 1-17
- online 1-17
- ordering printed 1-17

dump command 3-7

Dynamo script

- defined 1-14

E

Error-handling techniques 2-10

- capturing a return status 2-10
- getting reports from global variables 2-10

Exporting data 1-16

Extended attributes for external client application

- defined by PowerDesigner 3-4

F

Foreign keys 2-2

- function of 2-2
- keeping consistent with primary keys 2-3

Forms

- creating with InfoMaker 3-10

G

Global variables 2-10

grant command 3-8

Granting permissions

- with Sybase Central 3-8

Graphs

- creating with InfoMaker 3-10

Groups

- designating 3-8

H

- Historical Server
 - writing performance data to files for off-line analysis 1-13
- HTML templates
 - part of NetImpact Dynamo 1-14

I

- Importing data 1-16, 3-5
- Indexes 2-4
 - as element for high performance 2-5
 - creating 3-6
 - description 2-5
 - speeding up data retrieval with 2-4
- InfoMaker
 - as GUI client application 3-10
 - Data Manipulation painter for importing data from a file 3-6
 - data pipeline 1-15
 - functionality 1-15
- insert command 3-5, 3-9
- Internet 1-8
- Intranets 1-8
- ISAPI application server
 - Net Impact Dynamo as 1-15
- isql utility 1-16, 2-10

J

- JavaScript scripting language
 - embedded in NetImpact Dynamo templates 1-14
- Join 2-8
 - example 2-8

K

- Key columns 1-2, 2-2
- Keys 1-2, 2-2
 - foreign keys 2-2
 - primary keys 2-2

L

- Labels
 - creating with InfoMaker 3-10
- Libraries 1-4, 1-16, 3-12
 - See also* APIs (application programming interfaces)
- load command 3-7
- Loading data 3-5
- Loading data from an existing database with InfoMaker's data pipeline 3-5
- Logical design 1-3
- Logins
 - assigning 3-7
 - necessary for connecting to a server 3-7

M

- Master database 2-10
- Media failure 3-7
- model* database 2-10
- Modifying data
 - defined 3-9

N

- NetImpact Dynamo
 - application server 1-10, 1-15
 - as client and server 1-10
 - creating a web site with 3-12
 - functionality 1-8, 1-14
 - personal Web server 1-15
 - scripting language 1-14
- Net Library 1-16
- Network configuration files
 - creating and modifying with dsedit 1-16
- Non-centralized computing 1-6
- NSAPI application server
 - NetImpact Dynamo as 1-15

O

- ODBC (Open Database Connectivity)
 - drivers 1-17
 - installed with Open Client software 1-17
- ODBC Administrator 1-17
- ODBC data sources 1-17
- On-line documentation 1-17
 - installing 3-1
- Open Client 3-12
 - functionality 1-15
 - passing parameters to 2-10
- Open Client Libraries 3-12
 - Client-Library 1-16
 - CS-Library 1-16
 - DB-Library 1-16
- Open Database Connectivity Drivers.
 - See* ODBC (Open Database Connectivity) drivers

P

- Pages 2-5
- Performance
 - realizing potential 1-3
- Performance parameters
 - displayed by Adaptive Server Monitor 3-14
 - tuning 1-14, 3-14
- Performance statistics
 - analyzing offline 1-13
 - viewing 1-14, 3-14
- Permissions
 - defined 2-12
 - examples 2-12
 - granting to users 3-7
 - granting with Sybase Central 3-8
- Personal Web server (NetImpact Dynamo's) 1-15
- Physical database design
 - with PowerDesigner 3-3
- Physical Data Model (PDM)
 - created by PowerDesigner 3-3

- reverse-engineered from an existing database 3-4
- PowerBuilder 3-11
- PowerDesigner SQL Modeler
 - as design tool 3-3
 - defining extended attributes with 3-4
 - functionality 1-15
 - using to create a database 3-5
 - using to reverse-engineer a database 3-4
- Primary keys 2-2
 - function of 2-2
- Printed documentation 1-17
- Privileges. *See* Permissions
- Programming Adaptive Server 2-10
- Program window
 - showing client components 1-14
- Projection 2-7
 - example 2-7
- psdemodb* sample database 2-11
- pubs2* database 2-11

Q

- Querying
 - with Transact-SQL 2-7
- Query operations
 - basic in relational system 2-7
- Query optimizer 3-14

R

- RDBMS (relational database management system)
 - advantages of 1-3, 1-4
 - defined 1-1
 - facilities of Adaptive Server 1-4
 - technical overview of Sybase 1-12
- Referential integrity
 - enforcing with triggers or constraints 2-3
- Relational database management system. *See* RDBMS (relational database management system)

- Relations 1-1
 - See also* Tables
- Replication. *See* Data replication
- Replication Server 1-7
- Reports
 - creating with InfoMaker 3-10
- Restoring data from backups 3-7
- Restriction. *See* Selection
- Retrieving
 - information distributed over multiple tables 1-2
 - system administration information 2-12
- Return status 2-10
- revoke command 3-8
- Rows 2-1
- Rules 2-5
 - examples of 2-5

- S**
- Sample databases
 - psdemodb* 2-11
 - pubs2* 2-11
- Security
 - three-tiered system 3-7
- select command 2-7
- Selection 2-7
- Sending commands to the RDBMS 1-16
- Server
 - typical result sent by 1-5
- Server components 1-12
- Server Config
 - functionality 1-13
- Server-side utilities 1-13
- sp_addgroup system procedure 3-8
- sp_addlogin system procedure 3-8
- sp_adduser system procedure 3-8
- sp_configure system procedure 3-3
- sp_syntax system procedures 2-11
- SQL
 - defined 1-4
- SQL Advantage 2-10
- SQL Remote 1-7
- SQL statements
 - examples 2-7
- Stand-alone client applications
 - created by InfoMaker 3-11
- Stored procedures 2-5
 - creating 3-9
 - example of 2-6
 - executing faster than constituent commands 2-6
 - invoking 2-6
 - passing parameters to 2-6
 - system procedures supplied by Sybase 2-6
- Structured Query Language. *See* SQL
- Sybase Adaptive Server
 - environment 1-4
- Sybase Central
 - functionality 1-14
 - using to administer NetImpact Dynamo 3-12
 - using to execute system administration tasks 2-13
- Sybase client/server configuration 1-5
- Sybase System Administrator
 - creating user ids 3-8
 - designating groups 3-8
 - login for 2-13
 - special privileges of 2-13
- sybsyntax* database 2-11
- sybssystemprocs* database 2-10
- System catalogs. *See* System tables
- System databases 2-10
 - master* 2-10
 - model* 2-10
 - sybsecurity* 2-11
 - sybsyntax* 2-11
 - sybssystemprocs* 2-10
 - tempdb* 2-11
- System procedures 2-6, 2-12
- System tables
 - defined 2-12
 - poster illustrating 2-12

T

- Tables 2-1
 - defined 1-1
- tempdb* database 2-11
- Transaction logs 2-13
 - storing on a separate device 2-13
 - used to recover data 2-13
- Transactions 3-9
 - classic example of 3-9
 - defined 3-9
- Transact-SQL 1-4
 - commands 2-9
 - defined 2-6
 - enhancements in 2-6
 - extensions in 2-9
 - presenting commands to the RDBMS 2-10
- Triggers 2-6
 - examples of 2-6
- Tuning performance 3-14

U

- update command 3-9
- Upgrading Adaptive Server 3-1
- Upgrading tables 3-6
- User databases
 - as opposed to system databases 2-10
- User id
 - necessary for connecting to a Server 3-8
- Users
 - assigning logins to 3-7
 - creating 3-8
 - maintaining permissions for 2-12
- Utilities 2-13
 - command-line 1-13
 - examples of 2-13

V

- Views 2-4

- customizing a user's perception of a database with 2-4
- updatable vs. read-only 2-4

W

- wallvers desktop utility 1-16
- Web application scripts
 - client-side 1-14
 - server-side 1-14
- Web servers
 - distributing static data 1-9
 - interactive 1-10
 - with CGI API 1-10
- Web sites
 - non-interactive 1-9
 - stored in database and created by Sybase Central 1-15
 - visual representation of in Sybase Central 3-13