Pervasive.SQL 2000i

Getting Started with Pervasive.SQL 2000i

Workstation/Workgroup Edition

Getting Started with Pervasive.SQL 2000i
March 2001
100-003670-004

Contents

Getting Ready to Install

Installing the Engine

4 Installing Pervasive.SQL 2000i with No Previous Installation . . . 4-1

Instructions for First-time Workstation/Workgroup Installation

5 Upgrading Your Pervasive.SQL Installation 5-1

Instructions for Upgrading an Existing Pervasive.SQL Installation

Customizing the Workgroup Engine

After Installation

13   Troubleshooting After Installation   13-1

How to Proceed When You Encounter Errors During Installation

14   Pervasive.SQL Resources and Contacts 14-1

A Guide to Pervasive.SQL Customer Information Resources

Figures

Tables

About This Manual

This manual contains information about installing the Pervasive.SQL 2000i database system. Pervasive.SQL 2000i is a complete database system, providing the best of both worlds. It combines a transactional system designed for high-performance data handling and improved programming productivity with an embeddable and scalable relational system.

This manual also contains information about common installation pitfalls, general network protocol information, and Pervasive.SQL 2000i components.

For information about configuring the Pervasive.SQL 2000i engine, see the *Pervasive.SQL Advanced Operations Guide.*

Who Should Read This Manual

This manual provides information for users who install and run Pervasive.SQL 2000i Workstation or Workgroup edition.

Pervasive Software would appreciate your comments and suggestions about this manual. As a user of our documentation, you are in a unique position to provide ideas that can have a direct impact on future releases of this and other manuals. If you have comments or suggestions for the product documentation, post your request at http://www.pervasive.com/devtalk or send e-mail to docs@pervasive.com.

Manual Organization

This manual is arranged in the order of the main installation sequence. You complete the installation by following the chapters in order. *Getting Started with Pervasive.SQL* is divided into the following sections:

Getting Ready to Install

- Chapter 1—"Welcome to Pervasive.SQL 2000i"

  This chapter describes the components and documentation that come with Pervasive.SQL 2000i.

- Chapter 2—"Preparing to Install Pervasive.SQL 2000i"

  This chapter discusses important preparations that you should undertake before attempting to install Pervasive.SQL 2000i .

- Chapter 3—"Pervasive.SQL 2000i and Terminal Server Products"

  This chapter describes how to use Pervasive.SQL 2000i with Terminal Server Products.

Installing the Engine

- Chapter 4—"Installing Pervasive.SQL 2000i with No Previous Installation"

  This chapter describes how to install Pervasive.SQL 2000i for the first time on Windows NT, 2000 or Windows 95/98/ME.

- Chapter 5—"Upgrading Your Pervasive.SQL Installation"

  This chapter describes how to upgrade a previous version of Pervasive.SQL on Windows NT, 2000 or Windows 95/98/ME.

- Chapter 6—"Working with User Counts"

  This chapter describes how to apply a user count to a Workgroup engine.

Customizing the Workgroup Engine

- Chapter 7—"Network Settings for Workgroup Engine"

  This chapter describes how to configure your network for use with the Workgroup engine.

- Chapter 8—"Client Network Configuration"

  This chapter describes how to configure your clients network settings for use with the Workgroup engine.

■   Chapter 9—"Application Configuration Scenarios"

This chapter describes different application configurations for special scenarios.

■   Chapter 10—"Using the DOS Requesters"

This chapter describes how to install and configure the DOS requesters.

■   Chapter 11—"Configuring the Workgroup Engine"

This chapter describes how to configure the Workgroup engine using the Gateway Locator utility.

After Installation

■   Chapter 12—"Connecting Clients to a Pervasive.SQL Database"

This chapter steps through a few procedures using the Pervasive Control Center, including connecting to a database and retrieving some data.

■   Chapter 13—"Troubleshooting After Installation"

This chapter outlines some common problems that can occur when installing and how to resolve them. It also provides information on Pervasive.SQL 2000i tools that aid in diagnosing problems. This chapter gives contact information for Pervasive.SQL support for the case that you do not find the answer to your problem.

■   Chapter 14—"Pervasive.SQL Resources and Contacts"

This chapter describes online resources for help and provides contact information for Pervasive Software.

This manual also contains an index.

Conventions

Unless otherwise noted, command syntax, code, and examples use the following conventions:

CASE
: Commands and reserved words typically appear in uppercase letters. Unless the manual states otherwise, you can enter these items using uppercase, lowercase, or both. For example, you can type `MYPROG`, `myprog`, or `MYprog`.

Bold
: Words appearing in bold include the following: menu names, dialog box names, commands, options, buttons, statements, etc.

`Monospaced font`
: Monospaced font is reserved for words you enter, such as command syntax.

[ ]
: Square brackets enclose optional information, as in [*log_name*]. If information is not enclosed in square brackets, it is required.

|
: A vertical bar indicates a choice of information to enter, as in [*file name* | *@ file name*].

< >
: Angle brackets enclose multiple choices for a required item, as in `/D=<5|6|7>`.

variable
: Words appearing in italics are variables that you must replace with appropriate values, as in *file name*.

…
: An ellipsis following information indicates you can repeat the information more than one time, as in [*parameter*…].

::=
: The symbol ::= means one item is defined in terms of another. For example, a::=b means the item *a* is defined in terms of *b*.

PERVASIVE®
SOFTWARE
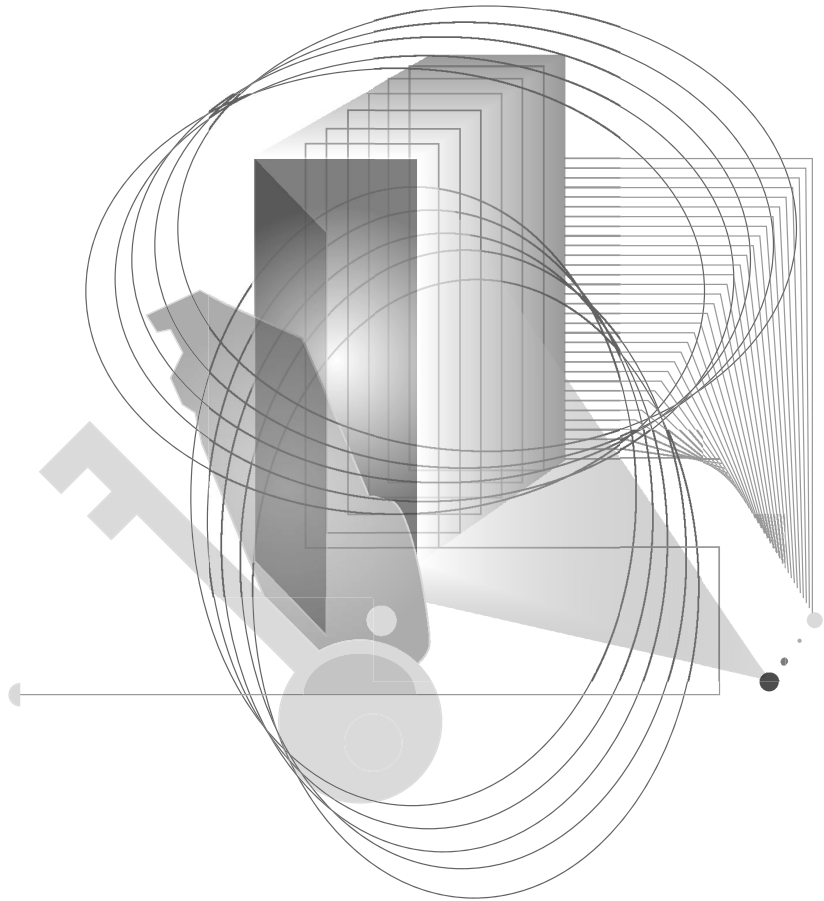
GETTING READY TO INSTALL

Welcome to Pervasive.SQL 2000i

A Basic Introduction to Pervasive.SQL 2000i

Thank you for purchasing Pervasive.SQL 2000i. We are confident that you will find Pervasive.SQL 2000i to be the very best high performance, low maintenance database engine on the market.

This chapter contains the following topics:

- "About Pervasive.SQL 2000i" on page 1-2
- "About the Workstation and Workgroup Engines" on page 1-4

About Pervasive.SQL 2000i

Pervasive.SQL 2000i is a reliable, low-maintenance, high-performance database management system (DBMS). Thousands of companies around the world license Pervasive.SQL and distribute it as the underlying data storage program for their data-intensive software products. These companies see no reason to build their own DBMS or license from a competitor once they experience the ease-of-use, reliability, and value offered by Pervasive.SQL.

No matter whether you received Pervasive.SQL with another product or purchased it yourself, this section explains a little about the product and why it is right for you.

If you would like to read a detailed discussion of Pervasive.SQL architecture and interfaces, please refer to *Pervasive Products and Services.*

Competitive Advantages

Pervasive.SQL provides a number of advantages over other products available on the market. Here are just a few:

- *Lowest total cost of ownership.* An independent study conducted by Aberdeen Group concluded that no major database product can match Pervasive.SQL's low total cost of ownership. How do we do it? See the next bullet.

- *No Database Administrator (DBA) required.* You can look in the newspaper any day of the week and see classified ads for Oracle, Sybase, or SQL Server database administrators, with sky-high salaries. Pervasive.SQL offers the unique Zero Database Administrator, or Z-DBA™, architecture. Its easy-to-use tools, bulletproof installation, and set-it-and-forget-it simplicity make it the perfect workhorse for desktop, workgroup, and departmental applications.

- *Scalable from the desktop to the Web.* Pervasive.SQL is available in three editions: the Ultra-light™ Workstation database engine supports applications running locally on the same computer as the engine; the Workgroup engine comes with a three-user license and scales up to about ten users; and the Server engine comes with a 10-user license and scales to hundreds of

concurrent users, including intranet and extranet applications. Upgrading to another configuration requires no changes to the supported application, just plug-and-play with the new database engine.

■ *Cross-platform support.* Unlike some competitors, Pervasive.SQL does not lock you in to a single platform. Pervasive.SQL databases are binary-compatible and supported across Microsoft Windows, Novell NetWare, Sun Solaris, and several varieties of Linux. No matter where your data is or where it is going to be, Pervasive.SQL is there for you.

■ *Big database features at a small price.* Pervasive.SQL offers full security, encryption, management and monitoring tools, and a host of other features you would expect to see in more expensive DBMS products.

■ *Legendary stability and reliability.* There's no doubt why 70% of the Windows desktop accounting market uses Pervasive.SQL as the underlying database of choice. When you've got to manage important data, you go for the database engine that won't let you down.

Relational or Transactional Access

Pervasive.SQL offers an architecture that is totally unique in the database management market. Our product allows you to access the exact same data through ODBC and OLE-DB, supporting applications like Microsoft ASP, Excel, and Access, or through the lightning-fast transactional interface called Btrieve. ODBC allows you to do complex reporting and data mining, while Btrieve provides massive throughput when you need the ability to view, update, or create millions of records a day.

It is the choice of each application vendor as to which interfaces they use. If you want to know which access methods are used in your application, ask your application vendor.

About the Workstation and Workgroup Engines

This section provides some basic information about the Pervasive.SQL Workstation and Workgroup engines. For a detailed discussion of Pervasive.SQL architecture, please refer to *Pervasive Products and Services.*

If you would like more information about what comes in your product and what gets installed, please see Chapter 2, "Preparing to Install Pervasive.SQL 2000i."

Purpose   You have purchased the Pervasive.SQL Workgroup or Workstation engine.

The Workstation engine is designed for single-user, desktop use where data access over a network is not necessary. The Workstation engine offers the same level of performance, reliability, and features as the Workgroup and Server engines. The only difference is the multi-user and networking capabilities.

The Workgroup engine is designed to support multi-user workgroups up to about ten users. At ten users and above, Server licenses are less expensive than Workgroup licenses.

The Workgroup engine offers a flexible approach to accessing data on remote servers, allowing a variety of small network configurations. If you have data files on a remote file system with no database engine present, you can configure your Workgroup engines so that a particular engine is always used to access the remote data, or you can set it up so that the first engine to access the files then "serves" those files until there are no more requests for data. After this point, again the first engine to access the files then owns the files while requests are coming in.

Where to Install   The Workstation engine must be installed and used on the same computer where the database files are located.

The Workgroup engine can be installed on the same computer where the database files are located, or it can be installed on other computers to access the data over the network. All client components are installed with the Workgroup engine, so even if you access remote files through another Workgroup engine, you do not need to install the client.

In the case of web applications, the client must be installed on the same computer as the web server. Multiple web server platforms require a client on each platform.

Additional User Licenses

If you receive Status Code 7063 or other status codes indicating you have exceeded your licensed user count, you may wish to buy additional user licenses. You can purchase additional licenses from your application reseller or directly from Pervasive. The additional license can be applied easily using a 1-step process.

Features

All Pervasive database engines offer the same powerful feature set and full-functioned support for programming interfaces. All engine editions are plug n' play compatible, requiring no application changes to switch from one engine to another. Data files are binary-compatible across all supported platforms.

The chart below shows the major differences between the three different editions of the product.

Table 1-1   Comparison of Server, Workgroup, Workstation Features

Feature	Server	Workgroup	Workstation
Supports Btrieve, ODBC, OLE-DB, JDBC, PDAC and ActiveX interfaces	✔	✔	✔
Full-featured relational support (online backup, security, referential integrity, management tools, and so on)	✔	✔	✔
Binary compatible data files across all platforms and engine editions	✔	✔	✔
Easy plug n' play upgrading, no application changes required	✔	✔	✔
Comes with online documentation	✔	✔	✔
Can access data on a file server where no database engine is installed		✔	✔
Engine runs on Windows	✔	✔	✔
Engine runs on NetWare	✔		

Table 1-1   Comparison of Server, Workgroup, Workstation Features

Feature	Server	Workgroup	Workstation
Engine runs on Linux	✔		
Engine runs on Solaris	✔		
Multi-user for small groups	✔	✔	
Scales to thousands of users	✔		
Extranet license available	✔		

Preparing to Install Pervasive.SQL 2000i

2

Preparation Needed for Pervasive.SQL 2000i Workstation/Workgroup Installation

This chapter prepares you to install Pervasive.SQL 2000i by providing an overview of major components, a list of what items are installed during the installation procedure, and providing a detailed checklist to gauge your readiness to proceed with the Pervasive.SQL 2000i installation.

This chapter contains the following sections:

- "Installation Overview" on page 2-2
- "Installation Checklist" on page 2-6

Installation Overview

This section provides an overview to the components that make up Pervasive.SQL 2000i. Also included in this section is information on the files included in the different types of installation.

Major Components of Pervasive.SQL 2000i

Pervasive.SQL 2000i consists of the following components:

Pervasive.SQL 2000i Database Engine

Pervasive.SQL 2000i consists of two database sub-engines:

- MicroKernel Database Engine (MKDE), which provides Btrieve/ MicroKernel API support for Pervasive.SQL applications
- SQL Relational Database Engine (SRDE), which provides ODBC/SQL API support for Pervasive.SQL applications

> **Note** The Pervasive.SQL 2000i stand-alone client install is not included in the Workstation/Workgroup edition of Pervasive.SQL 2000i. Each installation of Workstation/Workgroup edition includes all client components, so a separate install is not needed.

Client

Requesters that allow you to access the server-based MicroKernel from a DOS or Windows 3.x/9X/NT/2000 client.

- The 16-bit requesters used in DOS and Windows 3.x clients are Btrieve requesters only and have no relational access.

Pervasive.SQL 2000i Utilities

Pervasive.SQL 2000i includes a suite of useful utilities for every phase of database management from engine installation and configuration to data maintenance. They include the following:

- Pervasive System Analyzer (PSA)
- Pervasive Control Center (PCC)
- Configuration
- Monitor
- Function Executor

- Maintenance
- SQL Data Manager
- Table Designer
- Rebuild
- User Count Administrator
- ODBC Administrator
- Gateway Locator

Wizards

As part of the tools available in PCC, a suite of wizards provide an interactive way to assist you in database development, maintenance, and manipulation. Wizards available in Pervasive.SQL 2000i include the following:

- Create Database
- Drop Database
- Check Database
- Import Data
- Export Data
- Null Conversion
- Create Table
- Drop Table

Online Documentation

In addition to context-sensitive help for the utilities, you can install the complete Pervasive.SQL 2000i documentation set in either HTML or HLP formats, depending on your operating system and your available browser. Your product package contains a printed copy of *Getting Started with Pervasive.SQL 2000i (Workstation/ Workgroup Edition)* and the *Status Codes and Messages Quick Reference* card. You can order a complete hardcopy set of the manuals through the Pervasive Software Sales Team.

- *Getting Started with Pervasive.SQL 2000i (Server Edition)* - helps you to get Pervasive.SQL running with installation, setup, and basic configuration information based on your operating environment.

- *Getting Started with Pervasive.SQL 2000i (Workstation/
  Workgroup Edition)* - helps you to get Pervasive.SQL 2000i
  Workstation or Workgroup engine running with installation,
  setup, and basic configuration information based on your
  operating environment.

- *Pervasive Products and Services* - provides an overview of
  Pervasive Software products and services, including support
  options and information regarding licensing and registration.
  This manual also describes the database and product
  architectures.

- *Pervasive.SQL User's Guide* - gives basic information on the
  utilities provided with the Pervasive Control Center, including
  several interactive wizards designed to help you accomplish
  basic database tasks. This manual also contains information on
  the Unix version and the utilities available for that particular
  platform.

- *SQL Engine Reference* - provides database programmers a
  complete reference guide to the SQL relational database
  language supported in Pervasive.SQL 2000i.

- *Advanced Operations Guide* - details tasks typically performed by
  Network or System Administrators, using various tools such as
  the Rebuild, Maintenance, Monitor, and Function Executor
  utilities.

- *Status Codes and Messages* - lists the error and informational
  codes and messages that can be received when using Pervasive
  software.

Installation Options

Pervasive.SQL 2000i offers Typical and Custom installation options.
You can accept default installation parameters by selecting the
Typical install. On the other hand, you can specify certain
installation parameters by selecting the Custom install option. The
following outlines the major differences in these two installation
options.

Typical Setup Installation

The Typical installation is recommended for most users. This option
takes default actions for most operations performed during the
installation.

The Typical setup installs the following components:

- Pervasive.SQL 2000i Engine (including ODBC interface)
- All possible requesters (MS-DOS, Win16 and Win32 requesters)
- Utilities
- Transactional and Relational interfaces
- Client installation copies
- Online documentation

Custom Setup Installation

Selecting Custom setup type allows you to specify the program folder name and select the components to install.

During a Custom setup, the following component is optional:

- Online documentation

The default installation location for Pervasive.SQL 2000i is C:\PVSW, assuming the Windows system drive is C.

Installation Checklist

This section provides you with a checklist to prepare you for installation. Please use this checklist as a guide for a successful installation.

Quick Checklist Each of the following bullet points is described in more detail in the sections that follow.

❑ Your system hardware meets the minimum requirements (listed later in this chapter) to install Pervasive.SQL 2000i Workstation or Workgroup.

❑ Your operating system is supported by Pervasive.SQL 2000i.

❑ You are a system administrator or have administrative rights on the Windows system where you plan to install the product.

❑ Your application vendor supports the Pervasive.SQL 2000i engine.

❑ Schedule the upgrade appropriately and perform the installation with care.

❑ Licenses for older versions of the product cannot be migrated up.

❑ Read the README file for important, late-breaking warnings and information.

Check the Hardware Requirements Pervasive.SQL 2000i Workstation/Workgroup requires the following specifications:

❑ Windows 95/98/ME/2000 or Windows NT v4, SP5 or later.

❑ At least 16 MB of free memory (32 MB is recommended).

❑ Up to 8 MB of free space for system components.

❑ Up to 22 MB of free space in the installation location, depending upon the installation options chosen.

Check the Software Requirements

Your operating system must be supported by Pervasive.SQL 2000i. The following operating systems are supported:

- Windows 95/98
- Windows ME
- Windows 2000 (SP1 or later)
- Windows NT 4 (SP5 or later)
- If you are installing onto a terminal server, see Chapter 3, "Pervasive.SQL 2000i and Terminal Server Products."

Support for Previously Released Pervasive Server Engines

Pervasive.SQL 2000i client components support the following previously released Pervasive server engines:

- Btrieve Server v6.15 (NetWare, Windows NT)
- Pervasive.SQL v7.0 (NetWare, Windows NT)
- Pervasive.SQL 2000 v7.5 (NetWare, Windows NT)

See the *Application Migration Guide*, available with the Pervasive.SQL 2000i SDK, for more information.

Ensure Adequate Permissions

To install Pervasive.SQL 2000i on Windows NT or Windows 2000, you must be a system administrator or have administrative rights on your system. For information on granting administrative rights, please refer to *Pervasive.SQL User's Guide*.

Verify Compatibility of Vendor Applications

Contact your application vendor or review the documentation provided by your vendor to ensure that they support the Pervasive.SQL 2000i engine version and mode that you want to install. Please contact your application vendor if upgrading from Workstation to Workgroup for configuration information.

Schedule the Upgrade and Perform with Care

If installing the Pervasive.SQL 2000i Workgroup edition, the installation and upgrade should be performed during a period of low use.

As with any significant software installation, be sure to back up any important files on the target hard drive, including data files, before you begin the installation.

If you are performing an upgrade, keep the installation media and instructions from the old installation, in the unlikely event that you need to fall back to the previous version of the product.

Read the README file

Pervasive Software strongly recommends that you read the information contained in the README file. This file contains important product news that could not be included in the product documentation but is essential to your installation and use of the product release.

Pervasive.SQL 2000i and Terminal Server Products

How to Use Pervasive.SQL with Terminal Servers

This product is supported for use with Microsoft Terminal Server and Citrix MetaFrame. This chapter explains how to use Pervasive.SQL 2000i with either product, and the limitations of doing so. In general terms, you can use Pervasive.SQL 2000i on a terminal server in almost the same way as on a normal computer.

This chapter contains the following sections:

- "Supported Terminal Server Products" on page 3-2
- "Installation Notes" on page 3-3
- "Licensing Notes" on page 3-4
- "Configuration and Runtime Notes" on page 3-5

Supported Terminal Server Products

You may install Pervasive.SQL 2000i database engine and/or client in any of the following terminal server environments:

- Windows NT 4 Terminal Server SP6
- Windows 2000 Terminal Server SP1
- Windows NT 4 Terminal Server SP6 with Citrix MetaFrame 1.8
- Windows 2000 Terminal Server SP1 with Citrix MetaFrame 1.8

Supported Database Engines

Pervasive.SQL 2000i Server and Workgroup engines are supported on the platforms listed above.

Only one instance of the database engine may run on any terminal server platform. You cannot run separate copies of the database engine within two or more terminal sessions.

Installation Notes

Pervasive.SQL has in the past supported installation of only the client software on systems running a terminal server. The server engine components had to be installed on a separate system. With Pervasive.SQL 2000i, access from terminal sessions to server engine components located on the same system is supported.

The following environments are supported:

■ Windows NT 4.0 Terminal Server Edition (Service Pack 6)

■ Windows 2000 (Service Pack 1)

■ Citrix MetaFrame

Permissions

To install Pervasive.SQL 2000i on a terminal server, you must be logged into the terminal server as a user with system administrator rights.

Installing on Microsoft Terminal Server

Prior to installation on a computer with Microsoft terminal server installed, you must switch the terminal server to install mode using the `change user /install` command. This ensures that the software is available to all users after installation. When the installation is complete, you should set the terminal server back to runtime mode using the command `change user /execute`.

For more information about this command, see the Microsoft Knowledge base article (especially the "Additional Notes" at the bottom of the article):

http://support.microsoft.com/support/kb/articles/Q186/5/04.ASP

Workgroup Engine

After installing the Workgroup engine, you must remove it from the **Programs | Startup** folder for all terminal server users. Terminal server users cannot run their own copy of the database engine.

The engine must be started at the server console by a user with administrative rights.

Licensing Notes

You must purchase sufficient user count licenses to cover the number of users accessing each database engine. For example, if you have a 10-user server engine installed, you may not have more than 10 users accessing the database at once, regardless of whether they are connected through terminal sessions or remote database connections.

You may not use a terminal server as a method of bypassing the user count mechanism of the database engine.

Configuration and Runtime Notes

This section provides information about several configuration and runtime issues.

Remote Configuration with PCC

You cannot use PCC through a terminal session to configure a remote database engine. You must use PCC at the terminal server console or through a standard remote ODBC connection to configure the database engine on the terminal server machine.

Multiple Engines

Only one copy of any database engine may run at one time on the terminal server; and the Workgroup/Workstation engine cannot run in each terminal session.

Microsoft Terminal Server as Network Server

There are no particular problems with using your terminal server as your main network server and database server, however, if you have high usage of the server as a file server and also many terminal sessions running at once, you may find the performance less than satisfactory.

Another concern is having all of your mission critical services on the same machine. If it goes down, all your services go down at once.

For these reasons, you may wish to consider distributing your mission critical services on two or more computers.

PERVASIVE®
SOFTWARE
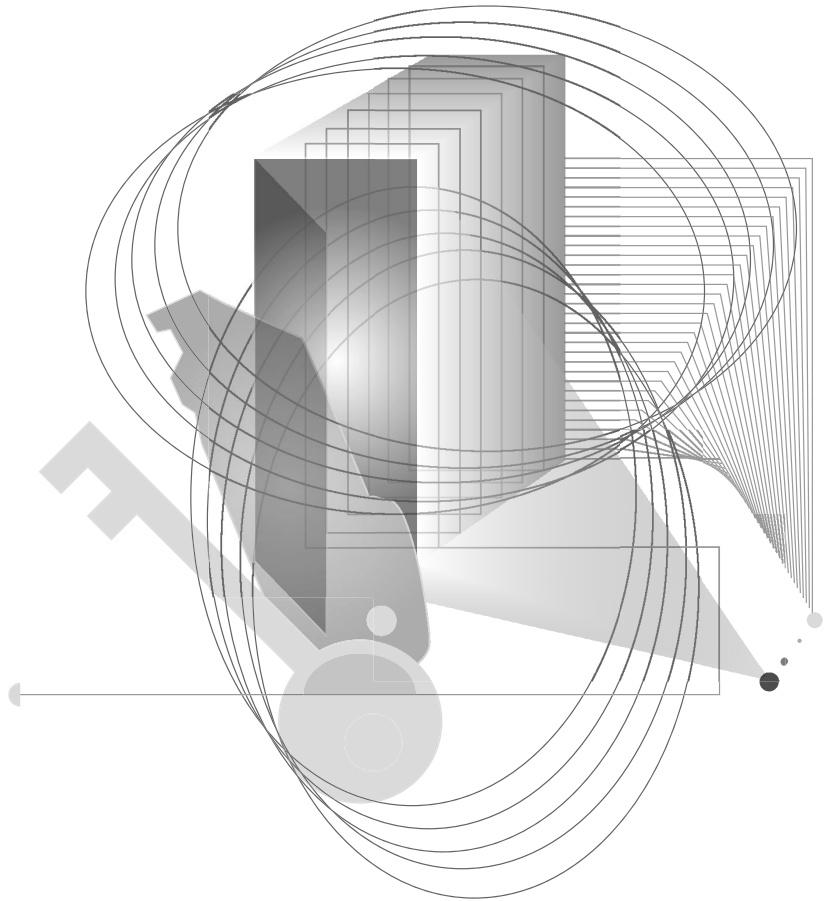
INSTALLING THE ENGINE

Installing Pervasive.SQL 2000i with No Previous Installation

4

Instructions for First-time Workstation/Workgroup Installation

This chapter contains procedures for installing the Pervasive.SQL 2000i Workstation/Workgroup engine when there is no previous installation of the product on your system. The chapter contains the following sections:

- "Before You Begin" on page 4-2
- "Installing the Pervasive.SQL 2000i Workstation/Workgroup Engine" on page 4-4
- "Custom Installation Path" on page 4-16
- "Common Questions After Installing Pervasive.SQL 2000i" on page 4-22

Before You Begin

Before installing Pervasive.SQL 2000i, begin by reviewing the following documents for important information:

❏ *Chapter 2, "Preparing to Install Pervasive.SQL 2000i"* - This chapter provides important information including system requirements and platform specific notes that are relevant to your operation.

❏ *"What's New in Pervasive.SQL 2000i"* - This document is included in your product package and contains information regarding the latest features and improvements in this release.

❏ *README file* - This file is located on the distribution media and contains late-breaking product news that could not be included in the product documentation.

Installing Over Existing Pervasive Products

Note If you are upgrading from Btrieve 6.15 or 6.30 and wish to make use of all the new version features, you must rebuild your data files so they use the v7 file format. See the *Pervasive.SQL User's Guide* section for detailed information on how to use the Rebuild Utilities to convert your data files.

If you install Pervasive.SQL 2000i over an existing Btrieve 6.15 or 6.30 installation, there will be no component conflicts because of the Smart Components system. Your Btrieve 6.x and Scalable SQL 4.x licenses, however, will not be migrated to Pervasive.SQL 2000i.

Installation Tips

■ When installing Pervasive.SQL 2000i for the first time on a system, Setup will check if all of the needed system files meet the minimum requirements. In some cases, these files are locked by the operating system and a reboot is required before Setup can continue. The following dialog box is displayed if this is the case.

Figure 4-1 Pervasive.SQL 2000i Setup Reboot Required
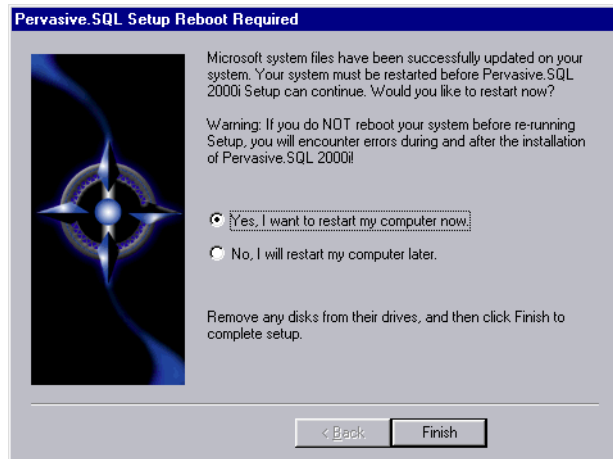
Click **Yes** to reboot the system. Setup is then automatically be restarted.

> **Note** It is strongly recommended that you reboot your system if you encounter this message. If you do not reboot your system, Setup will encounter failures during engine and utilities configuration.

- If you have any trouble with the following installation, see Chapter 13, "Troubleshooting After Installation."

Accessing README file information

It is highly recommended that you go through the README file to find out the latest changes and additions to Pervasive.SQL 2000i.

Installing the Pervasive.SQL 2000i Workstation/Workgroup Engine

> ⬗
>
> **Note** If you simply want to upgrade an existing Workstation engine to a Workgroup engine, go to "Upgrading from the Workstation Engine to the Workgroup Engine" on page 5-31 for those instructions.

➤ To install Pervasive.SQL 2000i on a Windows NT, 2000 or 9x workstation

1 Launch the installation program from your Windows workstation

    a. Insert the Pervasive.SQL 2000i Workstation/Workgroup CD in the CD-ROM drive of your Windows workstation.

    b. If the installation does not start automatically, click **Start**, select **Run**, and type `drive:setup` where *drive* is the drive letter of your CD-ROM device.

> ⬗
>
> **Note** Please see "Installation Tips" on page 4-2 regarding the Pervasive.SQL 2000i Setup Reboot Required dialog box.

The following dialog box displays:

Figure 4-2 Pervasive.SQL 2000i Setup Welcome Dialog Box
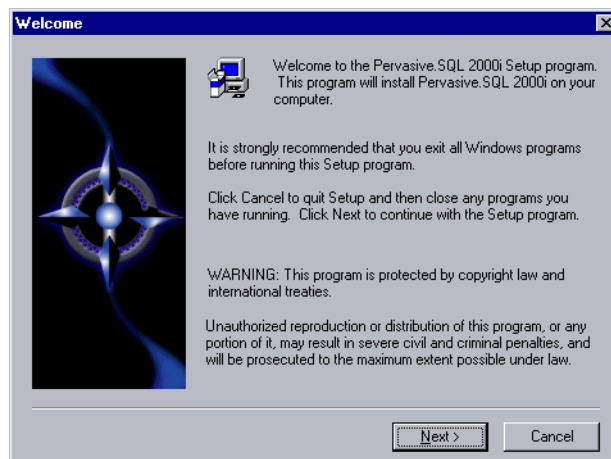
Click **Next** to proceed with the installation.

2   Read and accept the Software License Agreement

Figure 4-3 Pervasive Software's Software License Agreement
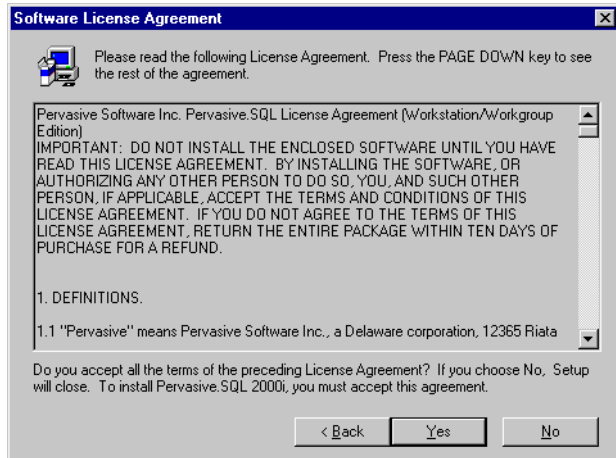
Read the Pervasive Software License Agreement. Click **Yes** to accept the terms of the agreement and continue installation.

3   Select the Setup Type - Typical or Custom

Figure 4-4 Setup Type Dialog Box
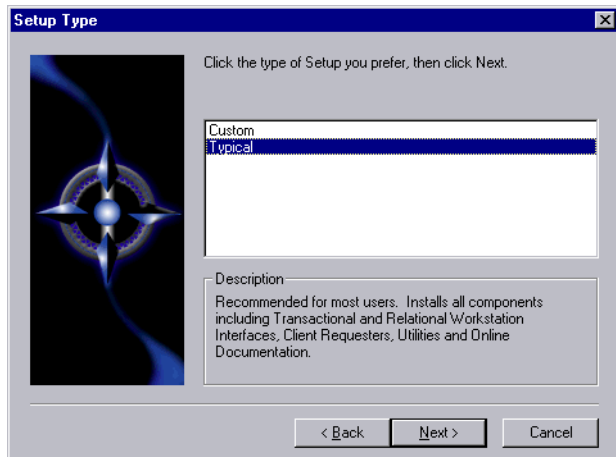
Select the **Typical** setup option and click **Next** to continue with installation.

Notes About Typical Setup

The Typical installation, which is recommended for most first time users, takes default actions for most operations performed during the installation.

The Typical workstation installation installs the following components:

- Pervasive.SQL 2000i Engine (including ODBC interface)
- Transactional and Relational Interfaces
- Utilities
- Online Documentation
- Client Requesters

Notes About Custom Setup

The Custom installation is typically for advanced users; it lets you choose the components you want to install and specify the program folder name.

Selecting Custom setup type allows you to specify the program folder name and select the components to install.

During a Custom installation you have the option to install the online documentation. The Transactional and Relational Workstation Interfaces, Client Requesters and Utilities are always installed.

> If you choose a Custom install, skip now to "Custom Installation Path" on page 4-16.

4 Select the workstation installation directory

The default installation location for Pervasive.SQL 2000i is C:\PVSW (assuming the Windows system drive is C:\). You can specify a different location for the installation as shown in the following dialog box.

Figure 4-5 Select Your Installation Directory
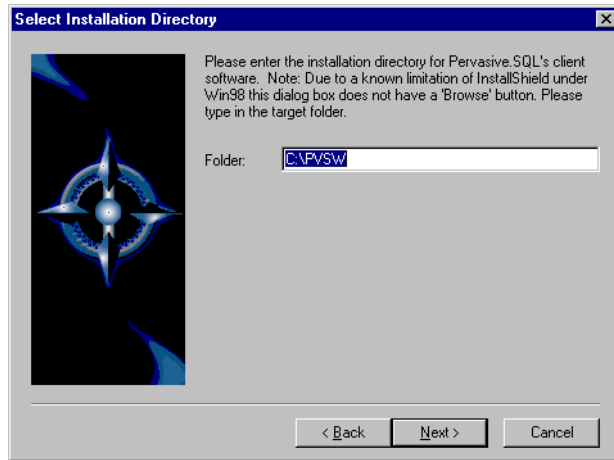
Click **Next** to accept the default installation directory and
proceed with the installation.

Note To select a different directory, click **Browse** and select or enter
the new installation directory. If the folder you have chosen doesn't
exist, you'll need to answer **Yes** to have setup create the folder for you.

Next the Pervasive System Analyzer (PSA) launches. PSA scans
your system for previous versions of Pervasive Software
products.

Figure 4-6 Pervasive System Analyzer Welcome Screen
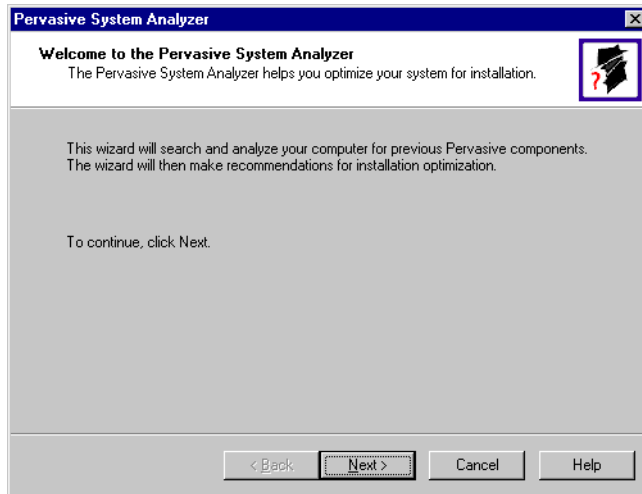
5   Analyze your system

Click **Next** to continue with the installation and begin system analysis. As PSA analyzes your system the following dialog box displays.

Note It is recommended that you perform this step, as it will intercept any potential problems after installation. You may be unaware of older Pervasive.SQL components on your system.

Figure 4-7 System Analysis Using PSA
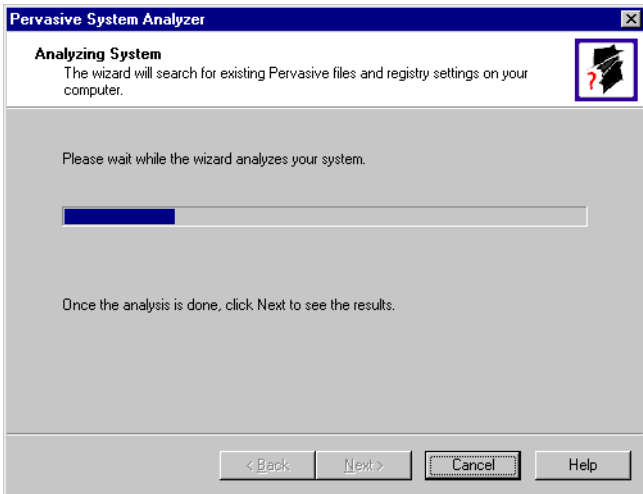
If you have previous versions of Pervasive Software products, PSA displays those products by the software version. (See Chapter 5, "Upgrading Your Pervasive.SQL Installation.")

The progress bar completes as PSA analyzes your system. When PSA has completed the initial system analysis, click the **Next** button to continue and see the results.

If you do not have any previous versions of Pervasive Software products, the following dialog box displays.

Figure 4-8 Pervasive System Analyzer Results
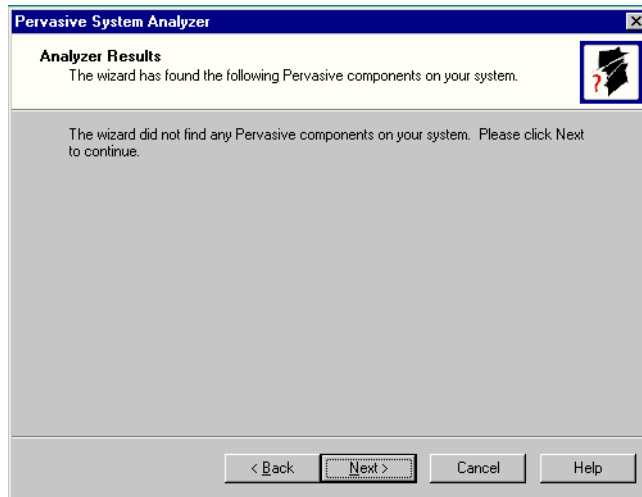
Click **Next** to display your current installation settings.

6   Verify your installation settings

Figure 4-9 Your Selected Installation Settings
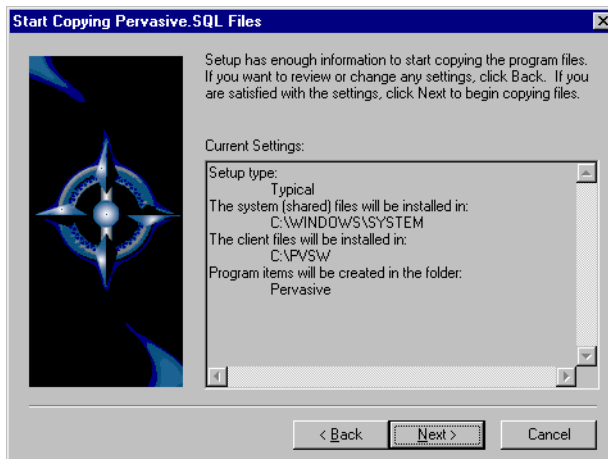
If the displayed selections are correct, click **Next** to continue.

Pervasive.SQL 2000i setup begins to copy the necessary files to complete your specified installation requirements.

Figure 4-10 Setup Begins Copying Files
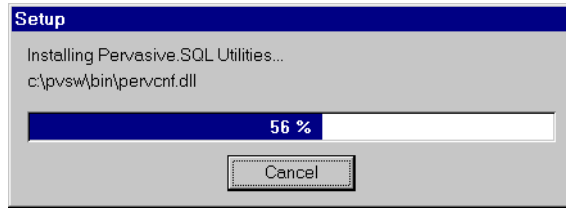
Note Only click **Cancel** if you want to stop installation and exit the Setup process.

7   Test the Transactional Engine

PSA now performs a series of tasks to ensure that the transactional engine is working properly.

Figure 4-11 Transactional Engine Test
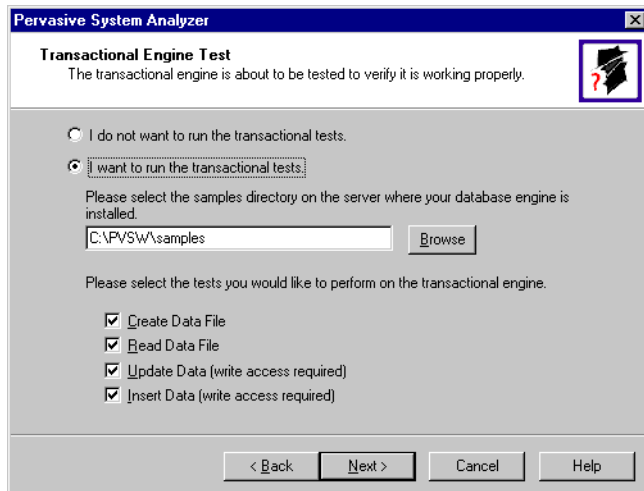
Select the samples subdirectory of the installation location. Then once you have selected to run the transactional tests and selected the tasks associated with that test, click **Next** to run the tests and see the results.

> **Note** It is recommended that you perform this test to ensure that the transactional engine is communicating properly and the transactional databases are now functional.

8   View the Transactional Test Results

PSA displays a check mark for each test that passes and an X for each task that fails during the transaction engine test. A summary report is also provided in the information window, as shown in the following figure.

Figure 4-12   Transactional Engine Test Results from PSA
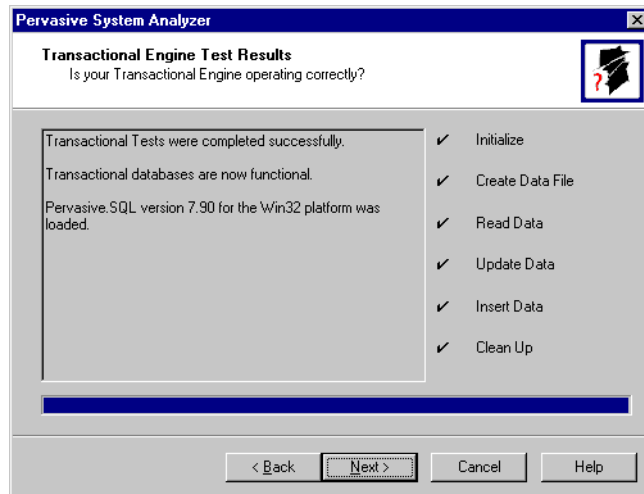
Once your transactional engine tests are complete, click **Next** to continue with installation and test the relational engine.

9   Test the Relational Engine

PSA now performs a series of tasks to ensure that the relational engine is working properly.

Figure 4-13 Relational Engine Test
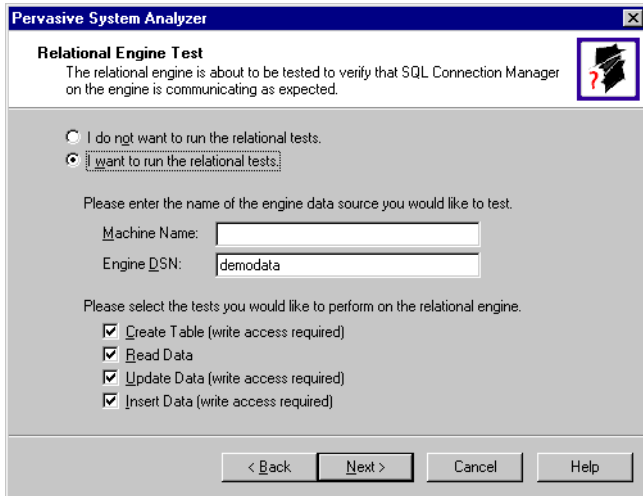
Enter the Machine Name and Engine DSN for the engine data you want to test. Then once you have selected to run the relational tests and selected the tasks associated with that test, click **Next** to run the tests and see the results. The "demodata" DSN is automatically setup during installation and should be used unless you have previously set up another database for testing purposes.

Note It is recommended that you perform this test to ensure that the relational engine is communicating properly and the relational databases are now functional.

10 View the Relational Test Results

PSA displays a check mark for each test that passes and an X for each task that fails during the relational engine test. A summary report is also provided in the information window, as shown in the following figure.

Figure 4-14 Relational Engine Test Results from PSA
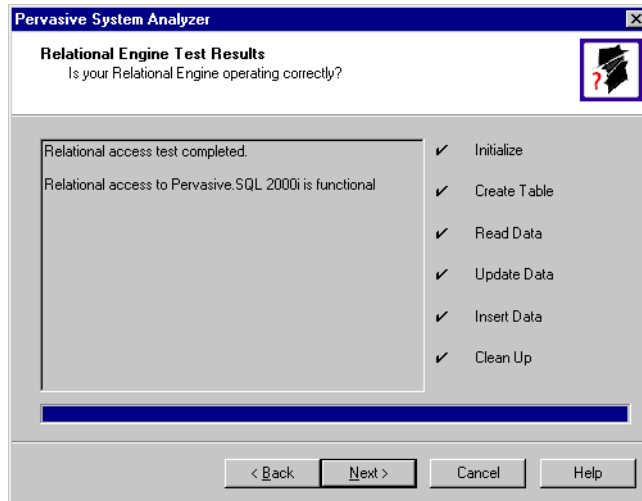
Once your relational engine tests are complete, click **Next** to finish the installation.

11 Read the README file

If you did not read the readme file prior to installation as described in "Before You Begin" on page 4-2, please do so now.

12 Your Pervasive.SQL 2000i workstation/workgroup installation is complete.

A dialog box will appear asking if you want to read the README file. It is highly recommended that you leave the **View readme file** check box selected and click **OK**. The following dialog box informs you that the installation setup is complete.

Figure 4-15 Installation Completion Dialog Box
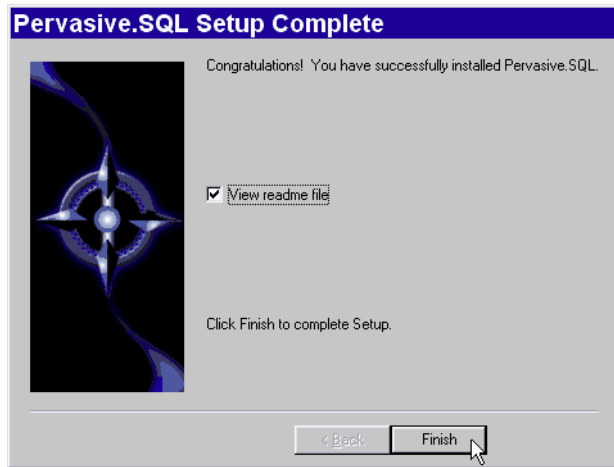
Click **Finish**.

Note If you have any trouble with the installation, see Chapter 13, "Troubleshooting After Installation."

If installed files could not be copied because they were locked in memory, a reboot may be necessary at this point. Setup only prompts for a reboot if a locked file or some other event was detected that requires a reboot. Please reboot your system if prompted to do so in order to ensure proper operation of your Pervasive.SQL 2000i product.

The install program modifies the PATH and CLASSPATH environment variables at the end of the installation process, if needed. These settings control how your Windows operating system finds Pervasive components. In Windows NT, these environment variables are stored in the **Control Panel** under the **Environment** tab on the **System** icon. In Windows 2000, these variables are stored in the **Control Panel** under the **Advanced tab** on the **System** icon.

What to do Next Skip over the next section for custom installations and see "Common Questions After Installing Pervasive.SQL 2000i" on page 4-22.

Custom Installation Path

This section describes how to customize your installation of Pervasive.SQL 2000i.

1 Choose the installation directory

The default installation location for Pervasive.SQL 2000i is C:\PVSW, assuming the Windows system drive is C:\. You can specify a different location for the installation as shown in the following dialog box.

Figure 4-16 Choosing Installation Directory of Pervasive.SQL 2000i
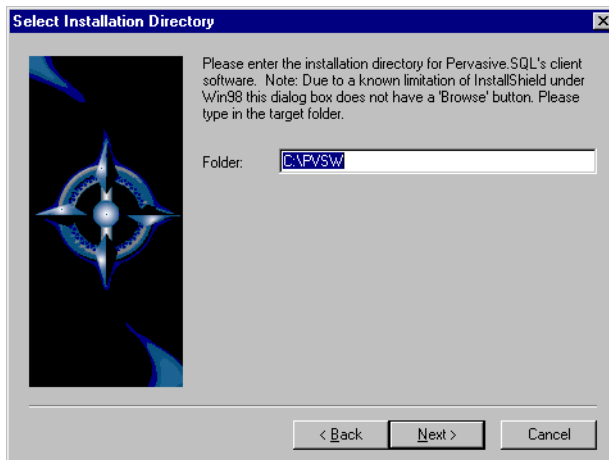
If you wish to change the default directory, click **Browse** and select the desired directory.

Click **Next** to proceed with the installation.

2 Select components to install

Pervasive.SQL 2000i Workstation/Workgroup edition currently contains only one optional component, which is online documentation.

Figure 4-17 Choose Components To Install
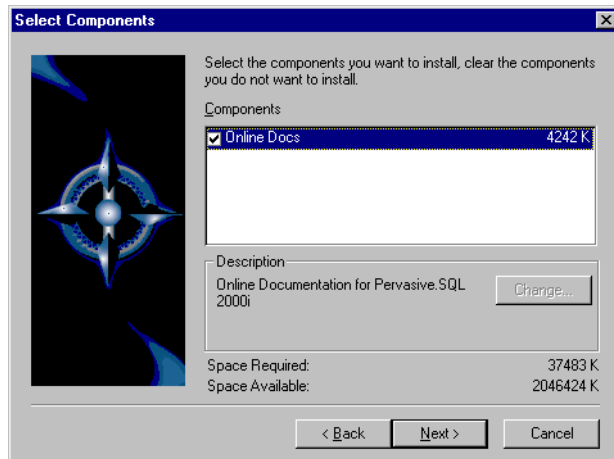
The **Online Docs** option is selected by default. If you choose not to install the online documentation, clear this checkbox. When you are finished, click **Next**.

Next the Pervasive System Analyzer (PSA) launches. PSA scans your system for previous versions of Pervasive Software products.

Figure 4-18 Pervasive System Analyzer Welcome Screen
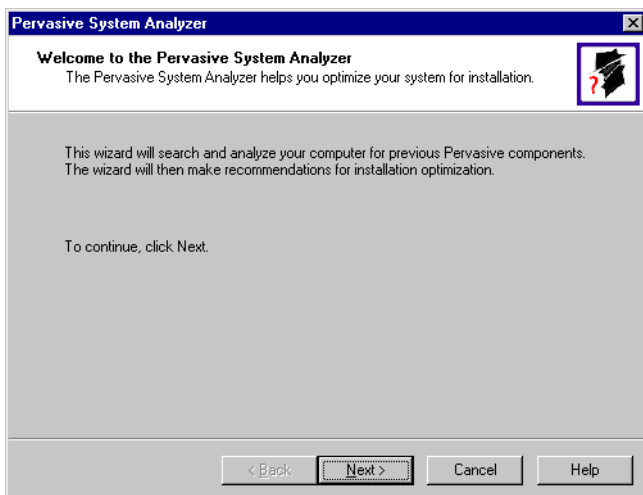
3 Analyze your system

Click **Next** to continue with the installation and begin system analysis. As PSA analyzes your system the following dialog box displays.

Figure 4-19 System Analysis Using PSA
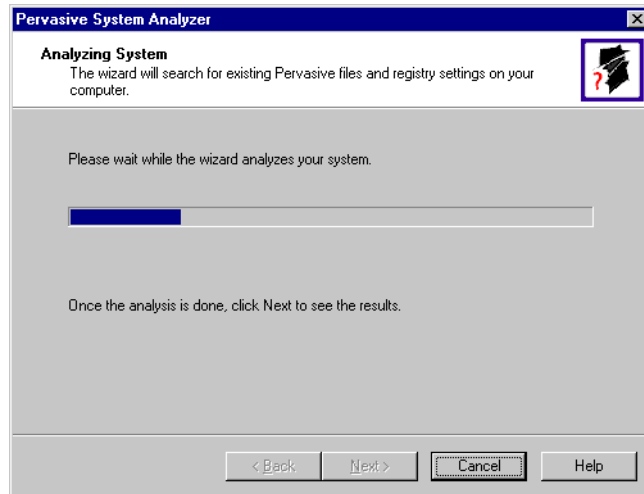
If you have previous versions of Pervasive Software products, PSA displays those products by the software version. (See Chapter 5, "Upgrading Your Pervasive.SQL Installation.")

The progress bar completes as PSA analyzes your system. When PSA has completed the initial system analysis, click the **Next** button to continue and see the results.

If you do not have any previous versions of Pervasive Software products, the following dialog box displays.

Figure 4-20 Pervasive System Analyzer Results
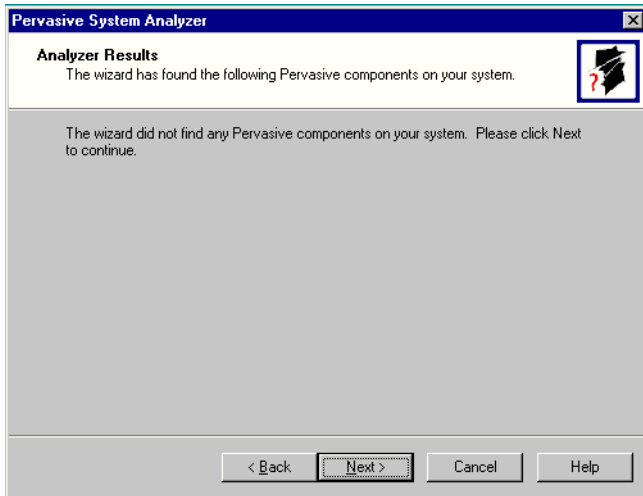
4   Choose the Program folder for menu items

Specify the folder in which the Start Menu items will be placed. Pervasive installation suggests a name for the folder, which you can accept by clicking **Next**.

If you have previously done an install on this machine, Install recommends that you do not use the same folder name so that your ability to cleanly uninstall the software in the future is preserved. (See Chapter 5, "Upgrading Your Pervasive.SQL Installation.")

Figure 4-21 Choose your Program Folder
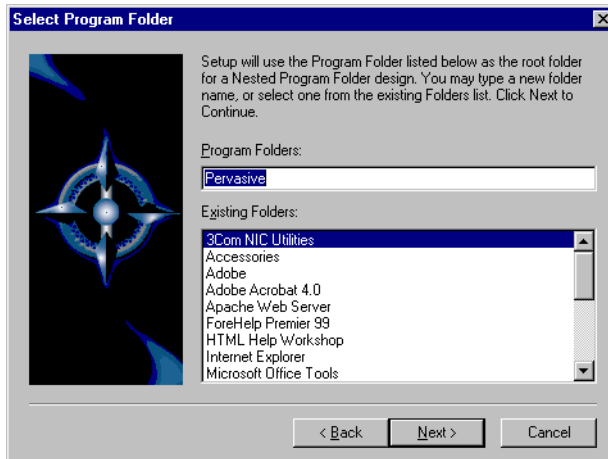
Click **Next** to accept the Program Folder name and display your current installation settings.

5 Verify your installation settings

Install now has enough information to start the installation process. A dialog box displays the list of options you have selected. If you want to change any of these selections, click **Back**.

If the displayed selections are correct, click **Next** to continue.

Figure 4-22 Your Custom Installation Selections
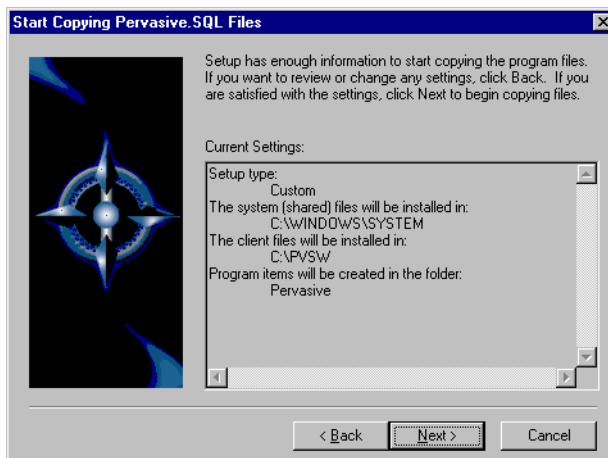
Pervasive.SQL 2000i setup begins to copy the necessary files to complete your specified installation requirements.

Figure 4-23 Setup Begins Copying Files
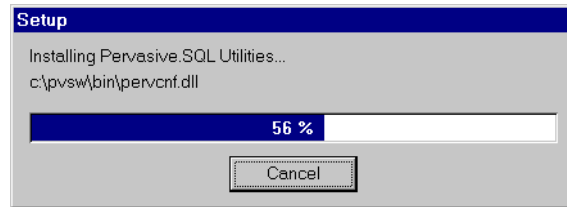
Note Only click **Cancel** if you want to stop installation and exit the Setup process.

The custom path rejoins the Typical installation path. See the step "Test the Transactional Engine," on page 4-11.

# Common Questions After Installing Pervasive.SQL 2000i

This section contains information that you should read after running the installation program. If you are having problems with your installation, go to Chapter 13, "Troubleshooting After Installation" or get help online from our Knowledge Base website at http://support.pervasive.com/eSupport.

If you are installing the Workgroup engine, you must install it on every computer that is expected to share data within your workgroup. Because every computer must have the Workgroup engine installed, there is no need to install the client software separately as there is for the Server engine.

How to Handle Data Source Names (DSNs)

After you have successfully installed the product, you have several choices depending on your situation, as shown in the table below.

If your situation is like this....	... then you should do this next:
You have existing Pervasive.SQL 7 DSNs already created.	You must delete and re-create all existing Pervasive.SQL 7 DSNs. Follow the instructions provided in "Deleting DSNs" on page 2-53 of *Pervasive.SQL User's Guide*.
You have existing Pervasive.SQL 2000 DSNs already created.	You should be able to access your data by connecting to the existing DSNs. Follow the instructions provided in "Accessing Data via ODBC From Other Applications" on page 2-46 of *Pervasive.SQL User's Guide*.
You are installing this product for the first time.	You should be able to connect to the sample DEMODATA database now. To create your own database or set up access to a database associated with another application, see *Pervasive.SQL User's Guide*.

How Do I Convert My Files From Previous Pervasive Products?

Converting your data files to 7.x format is not required, but you must convert them if you wish to take advantage of new features offered by the Pervasive.SQL 2000i engine.

Users of previous Btrieve versions: Use the Rebuild utility to convert your existing pre-6.0 or 6.x files to 6.x or 7.x format. For more information, refer to *Advanced Operations Guide*.

How Do I Read the Online Documentation?

Pervasive.SQL 2000i documentation is in the form of WinHelp files for Win32 platforms. To read the documentation, select the desired book name from the **Documentation** subfolder of the **Pervasive** program folder on the **Start** menu.

Documentation for Win16 platforms is in the form of Adobe Acrobat (PDF) files. Shortcuts to these files are found in the **Pervasive.SQL** documentation file group. The installation program for an Acrobat reader is available on the Pervasive.SQL 2000i CD-ROM.

There is also a README file on the distribution media that has the latest changes and additions to Pervasive.SQL 2000i.

What Files Were Installed As Part of Pervasive.SQL 2000i (Workstation/Workgroup Edition)?

Once installed, Pervasive.SQL 2000i files reside in the following default directory structure.

\WINNT\SYSTEM32	Microsoft system files Pervasive glue dlls (Win 9x only)
\PVSW	
\BIN	Pervasive.SQL 2000 engines, requesters, and utilities
\DEMODATA	Contains sample database
\DOC	Online documentation
\SAMPLES	Contains the sample database SAMPLE.BTR and UPPER.ALT, a sample Alternate Collating Sequence

Upgrading Your Pervasive.SQL Installation

Instructions for Upgrading an Existing Pervasive.SQL Installation

This chapter contains information specific for upgrading your current Pervasive product installation to Pervasive.SQL 2000i.

This chapter contains the following sections:

■ "Before You Begin" on page 5-2

■ "Patching an Existing Pervasive.SQL 2000 Installation" on page 5-5

■ "Upgrading to Pervasive.SQL 2000i Workgroup/Workstation Engine from Pervasive.SQL 7 or Btrieve 6" on page 5-12

■ "Custom Installation Path" on page 5-25

■ "Upgrading from the Workstation Engine to the Workgroup Engine" on page 5-31

■ "Common Questions After Installing Pervasive.SQL 2000i" on page 5-33

Throughout this document, when an explicit version number is not specified (for example: Pervasive.SQL 7, Pervasive.SQL 2000, or Pervasive.SQL 2000i), all versions are included.

Before You Begin

Before upgrading to Pervasive.SQL 2000i, begin by reviewing the following documents for important information:

❏ *Chapter 2, "Preparing to Install Pervasive.SQL 2000i"* - This chapter provides important information including system requirements and platform specific notes that are relevant to your operation.

❏ *"What's New in Pervasive.SQL 2000i"* - This document is included in your product package and contains information regarding the latest features and improvements in this release.

❏ *README file* - This file is located on the distribution media and contains late-breaking product news that could not be included in the product documentation.

Considerations Once you have reviewed the latest product information, review this list of considerations to complete your upgrade installation preparation.

❏ *Pervasive.SQL Applications* - Be aware of what applications you have currently using previous versions of Btrieve or Pervasive.SQL in your environment. Don't forget to include both client and server-based applications, such as ArcServe.

❏ *Vendor-Specific Information* - Be sure and check with your application vendors for any specific information regarding their product with Pervasive.SQL.

❏ *TCP/IP Protocol* - DOS applications require different requesters, preferably BTRBOX95, BTRDRVR, or BREQTCP.

❏ *New Features and File Rebuilding* - In order to make use of all the new version features, such as larger file sizes or system key generation, you must rebuild your data files so that they use the newest version file format. This chapter includes a section that details using the Rebuild Utility to rebuild your data files, if you would like to take advantage of newer features.

If you are upgrading from Btrieve 6.15 or 6.30 and wish to make use of all the new version features, you must rebuild your data files so they use the version 7 file format.

> **Note** If you install Pervasive.SQL 2000i over an existing Btrieve 6.15 or 6.30 installation, there will be no component conflicts because of the Smart Components system (unless your previous installation was installed to the Windows system directory). Your Btrieve 6.x and Scalable SQL 4.x licenses, however, will not be migrated to Pervasive.SQL 2000i.

❏ *Back Up Data Files* - Make sure you have a current backup of all your data files prior to beginning upgrade installation

Platform Notes for Windows

This section contains installation information specific to the Windows NT and Windows 2000 platforms.

■ If you run with Advanced Power Management, disable it while installing Pervasive.SQL 2000i. The power management can cause a Pervasive.SQL installation to fail if the server is suspended during the install. In Windows NT, the advanced power management setting is located on the **Power** icon on the **Control Panel**. In Windows 2000, this setting is located in the **Control Panel** under the **APM** tab of the **Power Options** icon.

Installing Over Existing Pervasive Products

If you install Pervasive.SQL 2000i over an existing Btrieve 6.15 or 6.30 installation, there will be no component conflicts because of the Smart Components system (unless your previous installation was installed to the Windows system directory). Your Btrieve 6.x and Scalable SQL 4.x licenses, however, will not be migrated to Pervasive.SQL 2000i.

> **Note** If you are upgrading from Btrieve 6.15 or 6.30 and wish to make use of all the new version features, you must rebuild your data files so they use the version 7 file format. See the *Pervasive.SQL Advanced Operations Guide* for detailed information on how to use the Rebuild Utilities to convert your data files.

Installation Tips

■ When installing Pervasive.SQL 2000i for the first time on a system, Setup will check if all of the needed system files meet the minimum requirements. In some cases, these files are locked by the operating system and a reboot is required before Setup can continue. The following dialog box is displayed if this is the case.

Figure 5-1 Pervasive.SQL 2000i Setup Reboot Required
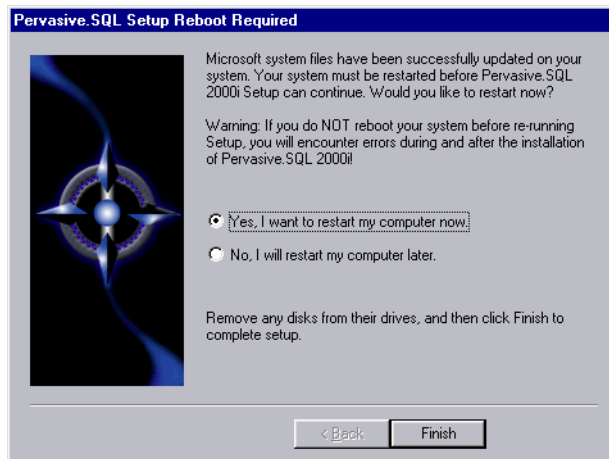
Click **Yes** to reboot the system. Setup is then automatically be restarted.

Note It is strongly recommended that you reboot your system if you encounter this message. If you do not reboot your system, Setup will encounter failures during engine and utilities configuration.

■ If you have any trouble with the following installation, see Chapter 13, "Troubleshooting After Installation."

Accessing README file information

It is highly recommended that you go through the README file to find out the latest changes and additions to Pervasive.SQL 2000i.

Patching an Existing Pervasive.SQL 2000 Installation

If you have an existing Pervasive.SQL 2000 installation on your machine, the installation program for Pervasive.SQL 2000i behaves differently than on a machine with no Pervasive.SQL 2000 components.

This section details the differences you will see with the patch installation.

- "Situation Where Install Runs in Patch Mode" on page 5-5
- "How Patch Mode Differs from Regular Installation" on page 5-5
- "Patching Procedure" on page 5-6
- "How to Run Standard Installation Instead of Patching" on page 5-11

Situation Where Install Runs in Patch Mode

The Pervasive.SQL 2000i installation enters patch mode if both of the following are true:

- You have existing Pervasive.SQL 2000 components on the installation machine.
- You did not run the Uninstall program for your previous Pervasive.SQL 2000 installation.

Note Install can only patch Pervasive.SQL 2000 versions prior to Pervasive.SQL 2000i (SP3). If you need to rerun install at a later date, uninstall Pervasive.SQL 2000i before reinstalling. Install does not support patching a machine that is already running Pervasive.SQL 2000i.

How Patch Mode Differs from Regular Installation

The patch installation differs from a regular installation in these ways:

No Prompting

- The patch install does not prompt you for installation choices such as installation directory, program folders, user count licenses, or component choices.

■ Pervasive System Analyzer runs in a silent mode. It automatically archives all your existing components before installing the new Pervasive.SQL 2000i components.

Existing Files and Licenses Updated

After installation:

■ Your existing Pervasive.SQL directory contains the updated files.

■ Your Windows start menu contains updated links.

■ Your existing user licenses are migrated to Pervasive.SQL 2000i.

Older Scalable SQL Components Preserved During Patch

If you have a version of Scalable SQL installed on the same machine as your previous version of Pervasive.SQL 2000, they are not affected by the patch install. The patch installation only replaces and upgrades Pervasive.SQL 2000 components. In this situation, the install behaves as documented in the section "Upgrading to Pervasive.SQL 2000i Workgroup/Workstation Engine from Pervasive.SQL 7 or Btrieve 6" on page 5-12.

Patching Procedure

Follow these steps to patch your Pervasive.SQL 2000 installation to Pervasive.SQL 2000i:

1 Launch the installation program from your Windows machine.

   a. Insert the Pervasive.SQL Workgroup/Workstation for Windows CD in the CD-ROM drive of your Windows machine.

   b. If the installation does not start automatically, click **Start**, select **Run**, and type drive:setup where *drive* is the drive letter of your CD-ROM device.

Figure 5-2 Pervasive.SQL 2000i Setup Welcome Dialog Box
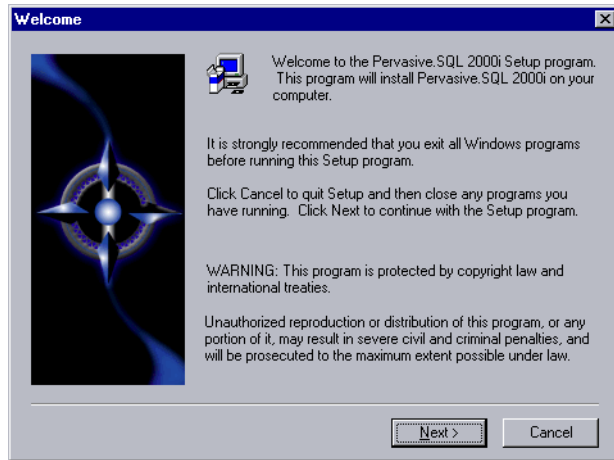
Click **Next** to proceed with the installation.

2 Read and accept the Software License Agreement

Figure 5-3 Pervasive Software's Software License Agreement
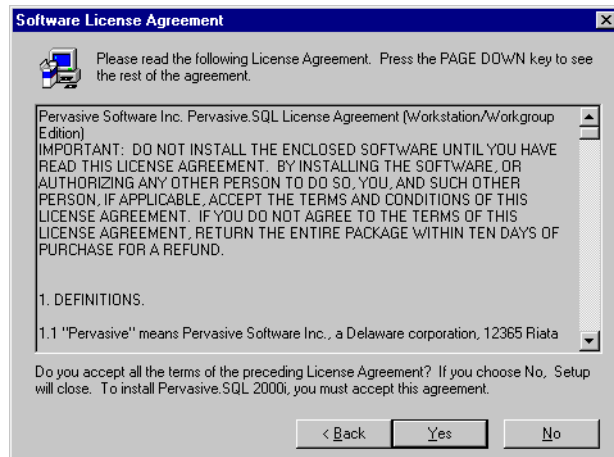
Read the Pervasive Software License Agreement. Click **Yes** to accept the terms of the agreement and continue the patch installation.

3 Pervasive System Analyzer will now archive previous Pervasive.SQL 2000 components. If need be, you can restore these components later using PSA.

Figure 5-4 Pervasive System Analyzer Preparing your System for Patching
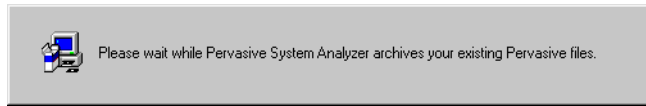
4   Verify your installation settings.

Figure 5-5 Your Selected Patch Installation Settings
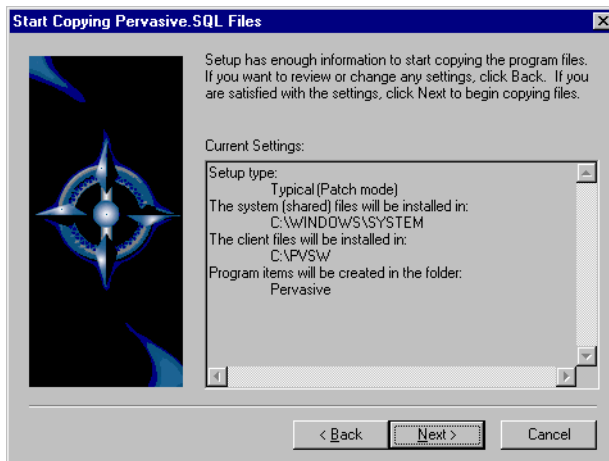
If the displayed selections are correct, click **Next** to continue.

Pervasive.SQL 2000i setup begins to copy the necessary files to complete your specified installation requirements.

Figure 5-6 Setup Begins Copying Files
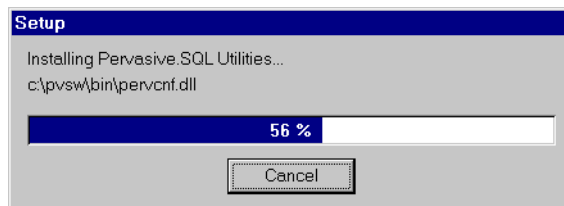
Note Only click **Cancel** if you want to stop installation and exit the Setup process.

5   After the files are copied, Install automatically performs the following functions:

- Sets up your new database engine
- Migrates your existing user count
- Configures program icons and start menu
- Ensures environment variables are set correctly

6   Pervasive System Analyzer checks the functionality of the new components. The following dialog is displayed during the test. Please look in the installation log file (install.log) and the PSA log file (PSAlog.txt) for information regarding the results of these tests.

Figure 5-7 Pervasive System Analyzer Tests Your New Engines

Please wait while Pervasive System Analyzer tests your transactional and relational engines.
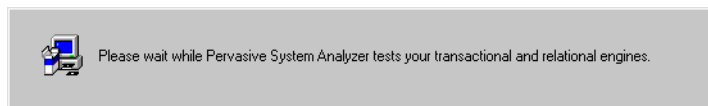
7   Read the README file.

If you did not read the readme file prior to installation as described in "Before You Begin" on page 5-2, please do so now.

Figure 5-8 Option to view the Readme file.

Question

Would you like to view the readme file? If you choose Yes to view the readme you might need to press Alt-Tab to toggle back to your browser.

Yes   No
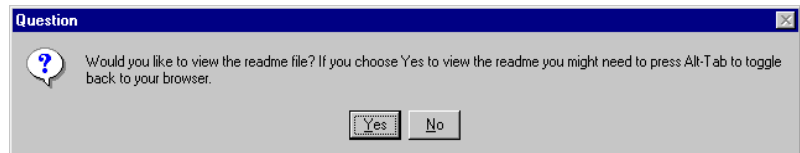
It is highly recommended that you click **Yes** and view the Readme file, if you have not yet done so.

Once you have selected to view the readme file, it is automatically displayed using your current browser or text editor application.

8   Your Pervasive.SQL 2000i workgroup/workstation patch is complete.

The following dialog box informs you that the installation is complete.

Figure 5-9 Installation Complete Dialog Box
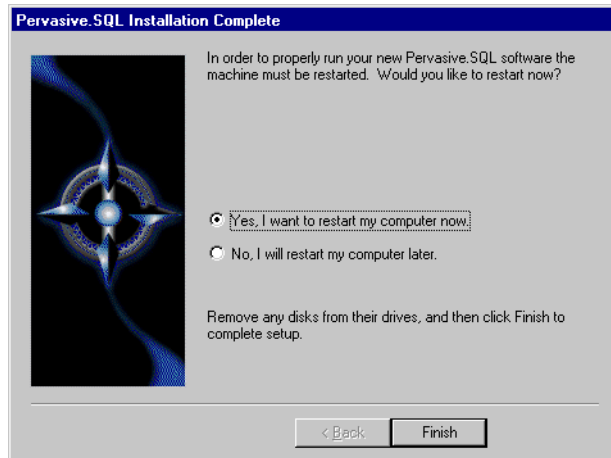
You should always leave the **Yes, I want to restart my computer now** check box selected.

Click **Finish**.

Note If you have any trouble with the installation, see Chapter 13, "Troubleshooting After Installation."

If installed files could not be copied because they were locked in memory, a reboot may be necessary at this point.

The installation program modifies the PATH, and CLASSPATH environment variables at the end of the installation process. These settings control how your Windows operating system finds Pervasive components. These environment variables are stored in the Control Panel under the Environment tab on the System icon. In Windows 2000, this option is under the Advanced tab of the System icon in the Control Panel.

How to Run Standard Installation Instead of Patching

You may want to perform a regular upgrade installation, as opposed to the patch installation, if any of the following are true:

- You already installed Pervasive.SQL 2000i but wish to reinstall it.
- You want to specify a different installation directory than the one used by your current Pervasive.SQL components.
- You want to be aware of which Pervasive.SQL 2000 components Pervasive System Analyzer is archiving, or you want to archive other versions of Pervasive software during the installation process.
- Your existing Pervasive.SQL 2000 installation did not contain features that you now need, such as Networked client support.

If any of these items apply to you, follow this uninstall procedure for your existing Pervasive.SQL 2000 components before proceeding with the Pervasive.SQL 2000i installation:

➤ **To uninstall existing Pervasive.SQL 2000 components:**

1 Stop your existing Pervasive.SQL 2000 engines if they are running.

2 Click **Start | Settings | Control Panel**

3 Click **Add/Remove Programs**

4 Find the entry for your Pervasive.SQL 2000 product (one of the following)

- Pervasive.SQL 2000 NT Server
- Pervasive.SQL 2000 Workgroup
- Pervasive.SQL 2000 Workstation

5 Click **Add/Remove** to remove the installation. You may be prompted to restart your system if any files were in use.

6 You can now install following the procedure in "Upgrading to Pervasive.SQL 2000i Workgroup/Workstation Engine from Pervasive.SQL 7 or Btrieve 6" on page 5-12.

Upgrading to Pervasive.SQL 2000i Workgroup/Workstation Engine from Pervasive.SQL 7 or Btrieve 6

➤ **To install Pervasive.SQL 2000i (Workgroup/ Workstation Edition) on a Windows machine**

1 Stop the Engine currently running

If your Pervasive.SQL Workgroup/Workstation engine is currently running, you must stop the engine and exit before you begin installation. To do this, right-click the engine icon in the task bar at the bottom of your screen and select the **Stop Engines and Exit** option.

2 Launch the installation program from your Windows machine

a. Insert the Pervasive.SQL 2000i Workgroup/Workstation for Windows CD in the CD-ROM drive of your Windows NT workstation.

b. If the installation does not start automatically, click **Start**, select **Run**, and type drive:setup where *drive* is the drive letter of your CD-ROM device.

Note Please see "Installation Tips" on page 5-4 regarding the Pervasive.SQL 2000i Setup Reboot Required dialog box.

Figure 5-10 Pervasive.SQL 2000i Setup Welcome Dialog Box
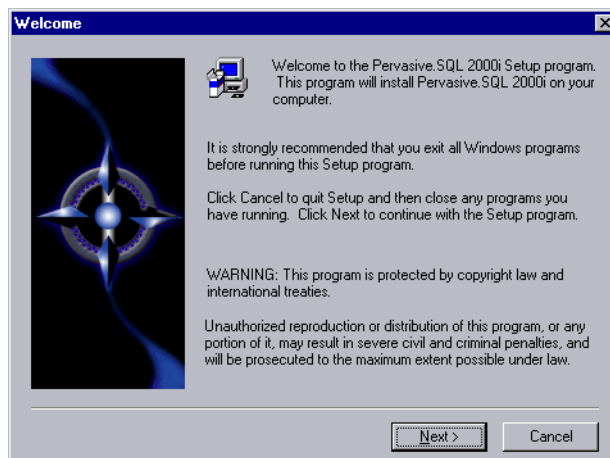
Click **Next** to proceed with the installation.

**3** Read and accept the Software License Agreement

Figure 5-11 Pervasive Software's Software License Agreement
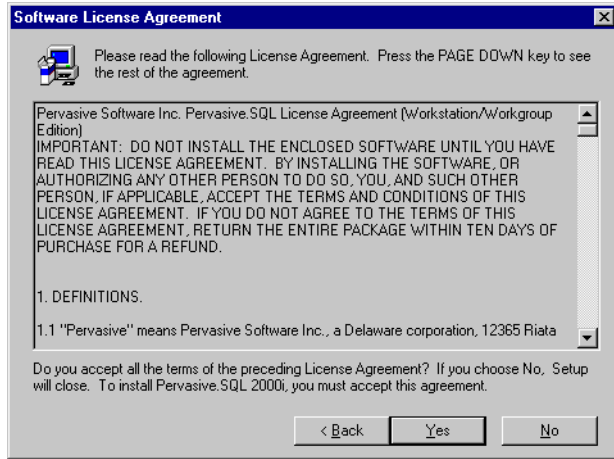
Read the Pervasive Software License Agreement. Click **Yes** to accept the terms of the agreement and continue installation.

4 Select typical or custom installation

Figure 5-12 Setup Type Dialog Box

Notes About Typical Setup

The Typical installation, which is recommended for most first time users, takes default actions for most operations performed during the installation.

The Typical workstation installation installs the following components:

- Pervasive.SQL 2000i Engine (including ODBC interface)
- Transactional and Relational Workstation Interfaces
- Utilities
- Online Documentation
- Client Requesters

Notes About Custom Setup

The Custom installation is typically for advanced users; it lets you choose the components you want to install and specify the program folder name.

Selecting Custom setup type allows you to specify the program folder name and select the components to install.

During a Custom installation you have the option to install the online documentation. The Transactional and Relational Workstation Interfaces, Client Requesters and Utilities are always installed.

If you choose a custom install, skip now to "Custom Installation Path" on page 5-25.

5 Select the workstation installation directory

The default installation location for Pervasive.SQL 2000i is C:\PVSW (assuming the Windows system drive is C:\).

> **Tip** If you are upgrading from a previous version of Pervasive.SQL, you may want to select a different installation location. Using a different installation location reduces any potential conflicts that may occur as a result of using the same installation directory.

Figure 5-13 Select your Installation Directory
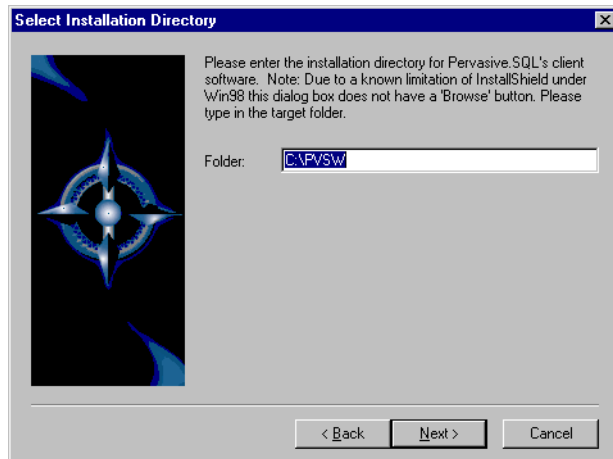
Click **Next** to accept the default installation directory and proceed with the installation.

> **Note** To select a different directory, click **Browse** and select or enter the new installation directory. If the folder you have chosen doesn't exist, you'll need to answer **Yes** to have setup create the folder for you.

6 Optimize your system with the Pervasive System Analyzer

Next the Pervasive System Analyzer (PSA) launches. PSA scans your system for previous versions of Pervasive Software products.

Figure 5-14 Pervasive System Analyzer Welcome Screen
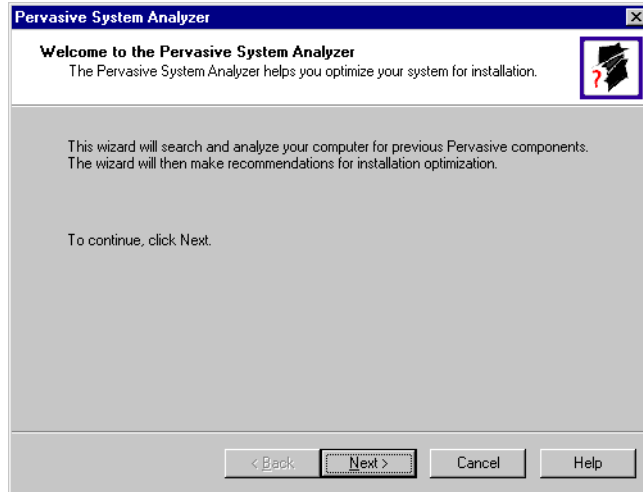
7 Analyze your system

Click **Next** to continue with the installation and begin system analysis. As PSA analyzes your system the following dialog box displays.

Figure 5-15 System Analysis Using PSA
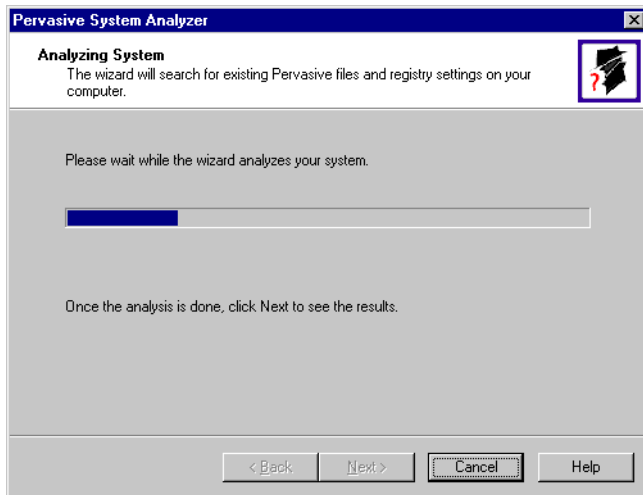
The progress bar completes as PSA analyzes your system. When PSA has completed the initial system analysis, click the **Next** button to continue and see the results.

If you have previous versions of Pervasive Software products, PSA displays those products by the software version. The next screen illustrates a system containing a previous installation of Pervasive.SQL.

Figure 5-16 Pervasive System Analyzer Results
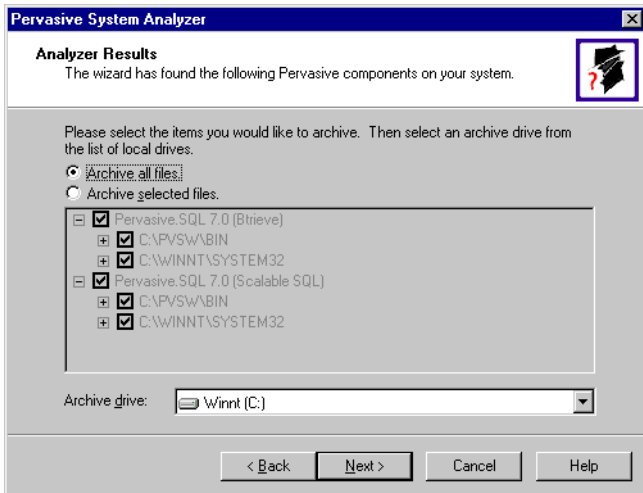
8 Select the Pervasive files to archive

Click **Next** to accept the default selection and archive all of your Pervasive components.

Note You can choose to archive only selected components, but Pervasive Software recommends that you archive all previous Pervasive components to reduce the possibility of product conflict and to ensure that your application(s) will run properly.

Figure 5-17 Select to Archive All Files
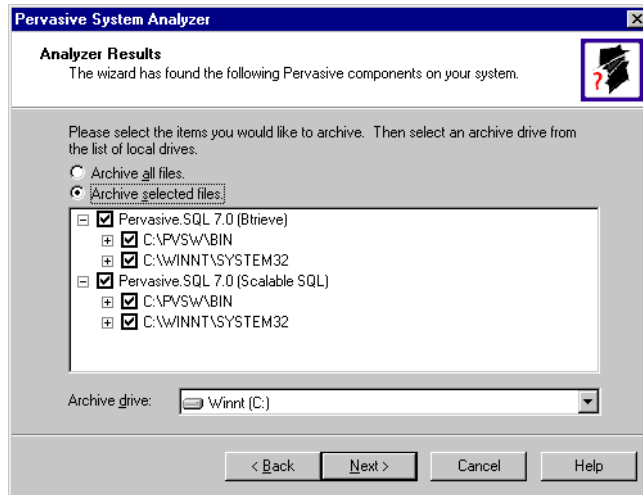
Note *If you are running an application that requires Scalable SQL:*
You may clear the selection **Pervasive.SQL 7.0 (Scalable SQL)**. You do
not have to leave Scalable SQL on your system in order to have ODBC
access to your data. Pervasive.SQL 2000i provides ODBC access.

Click **Next** to accept the default selection and archive all of your
Pervasive components. Otherwise, select **Archive Selected Files**
and deselect items that you do not wish to archive.

For example, if you wanted to keep the Scalable SQL interface for
your legacy applications, deselect the item 'Pervasive.SQL 7.0
(Scalable SQL)' as shown in the preceding dialog.

9 Archive your Pervasive components

Once you have selected to archive all the files, click **Next**.

10 Verify your installation settings

Figure 5-18 Your Selected Installation Settings
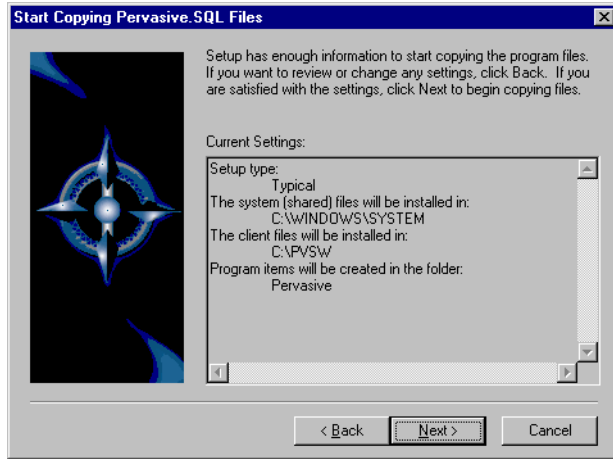
If the displayed selections are correct, click **Next** to continue.

Pervasive.SQL 2000i setup begins to copy the necessary files to complete your specified installation requirements.

Figure 5-19 Setup Begins Copying Files
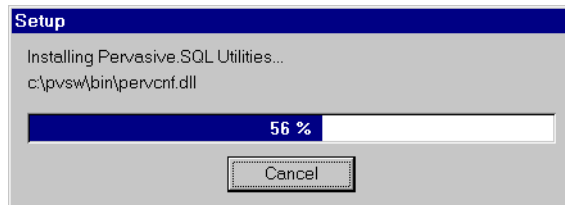
Note Only click **Cancel** if you want to stop installation and exit the Setup process.

11 Test the Transactional Engine

PSA now performs a series of tasks to ensure that your access to the transactional engine is working properly.

Figure 5-20 Transactional Engine Test
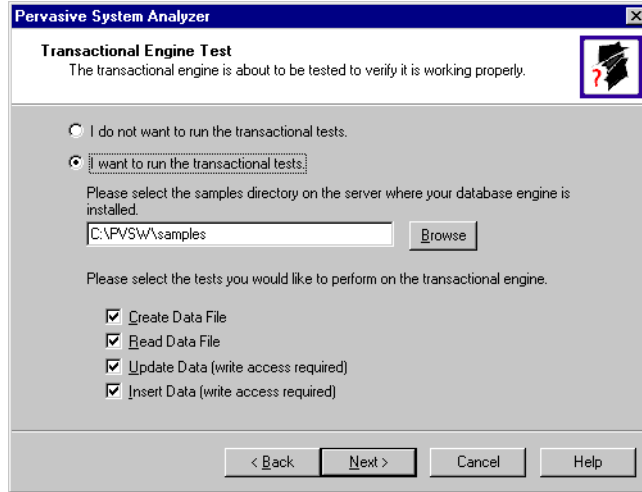
Select the samples subdirectory of the Workstation/Workgroup engine installation location. Then once you have selected to run the transactional tests and selected the tasks associated with that test, click **Next** to run the tests and see the results.

Note It is recommended that you perform this test to ensure that the transactional engine is communicating properly and the transactional databases are now functional.

12 View the Transactional Test Results

PSA displays a check mark for each test that passes and an X for each task that fails during the transaction engine test. A summary report is also provided in the information window, as shown in the following figure.

Figure 5-21 Transactional Engine Test Results from PSA
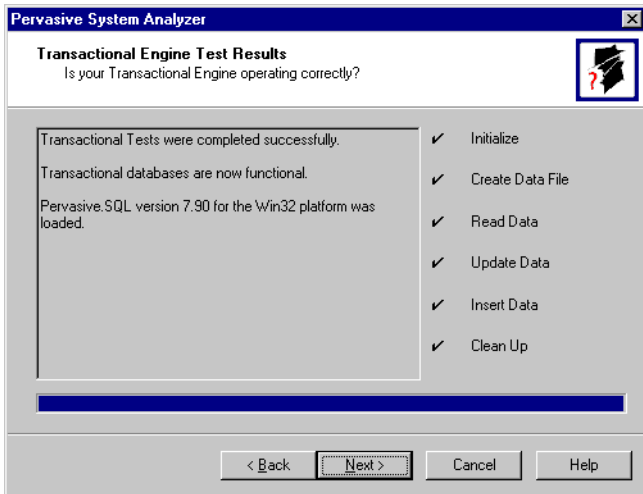
Once your transactional engine tests are complete, click **Next** to continue with installation and test the relational engine**.**

13 Test the Relational Engine

PSA now performs a series of tasks to ensure that the relational engine is working properly.

Figure 5-22 Relational Engine Test
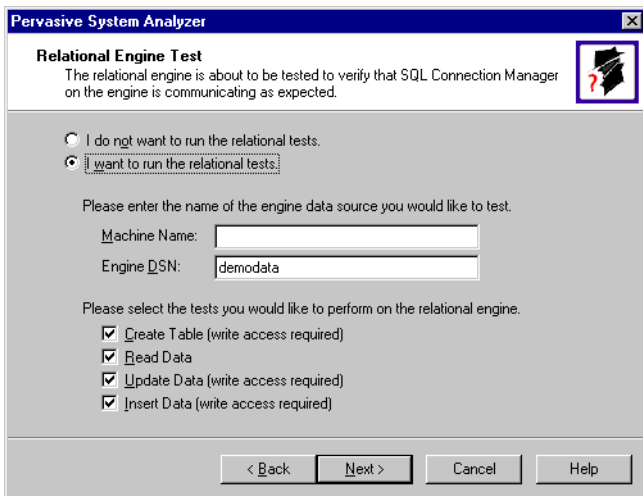
Enter the Machine Name and Engine DSN for the engine data you want to test. Then once you have selected to run the relational tests and selected the tasks associated with that test, click **Next** to run the tests and see the results.

Note It is recommended that you perform this test to ensure that the relational engine is communicating properly and the relational databases are now functional.

14 View the Relational Test Results

PSA displays a check mark for each test that passes and an X for each task that fails during the relational engine test. A summary report is also provided in the information window, as shown in the following figure.

Figure 5-23 Relational Engine Test Results from PSA
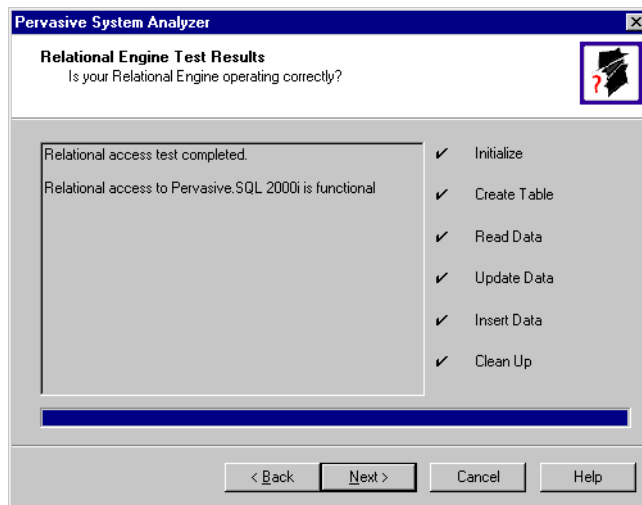
Once your relational engine tests are complete, click **Next** to finish the installation.

15 Read the README file

If you did not read the readme file prior to installation as described in "Before You Begin" on page 5-2, please do so now.

16 Your Pervasive.SQL 2000i workstation/workgroup installation is complete

A dialog box will appear asking if you want to read the README file. It is highly recommended that you leave the **View readme file** check box selected and click **OK**. The following dialog box informs you that the installation setup is complete.

Figure 5-24 Installation Completion Dialog Box
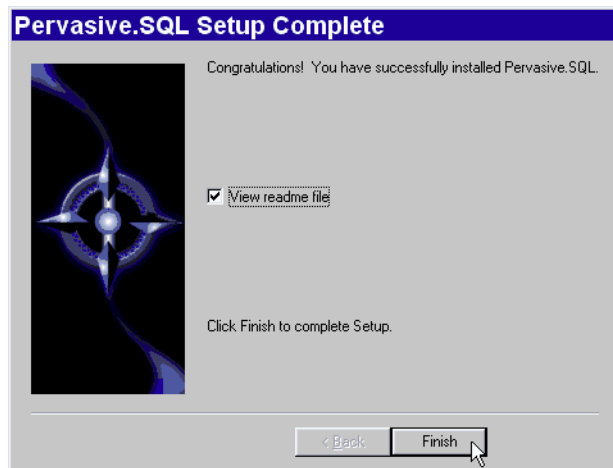
Click **Finish**.

Note If you had any trouble with installation, see Chapter 13, "Troubleshooting After Installation."

If installed files could not be copied because they were locked in memory, a reboot may be necessary at this point. Setup only prompts for a reboot if a locked file or some other event was detected that requires a reboot. Please reboot your system if prompted to do so in order to ensure proper operation of your Pervasive.SQL 2000i product.

The installation program modifies the PATH and CLASSPATH environment variables at the end of the installation process. These settings control how your Windows operating system finds Pervasive components. In Windows NT, the environment variables are located in the **Control Panel** under the **Environment** tab on the **System** icon. In Windows 2000, they are stored in the **Control Panel** under the **Advanced** tab on the **System** icon.

What to do Next   Skip over the next section for custom installations and see "Common Questions After Installing Pervasive.SQL 2000i" on page 5-33.

Custom Installation Path

This section describes how to customize your installation of Pervasive.SQL 2000i.

1   Choose the installation directory

The default installation location for Pervasive.SQL 2000i is C:\PVSW, assuming the Windows system drive is C:\. You can specify a different location for the installation as shown in the following dialog box.

Figure 5-25 Choosing Installation Directory of Pervasive.SQL 2000i
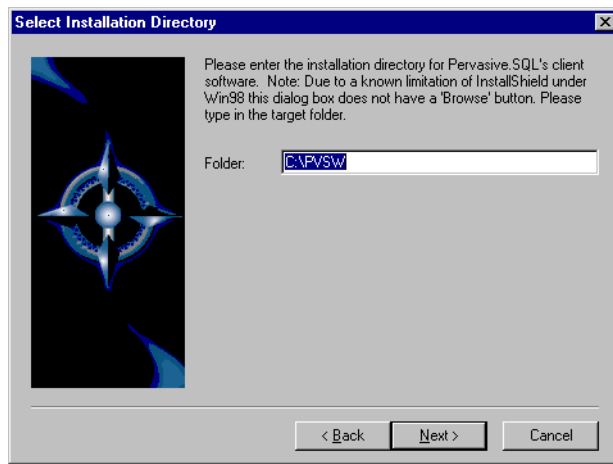
If you wish to change the default directory, click **Browse** and select the desired directory.

Click **Next** to proceed with the installation.

2   Select components to install

Pervasive.SQL 2000i Workstation/Workgroup edition currently contains only one optional component, which is online documentation.

Figure 5-26 Choose Components To Install
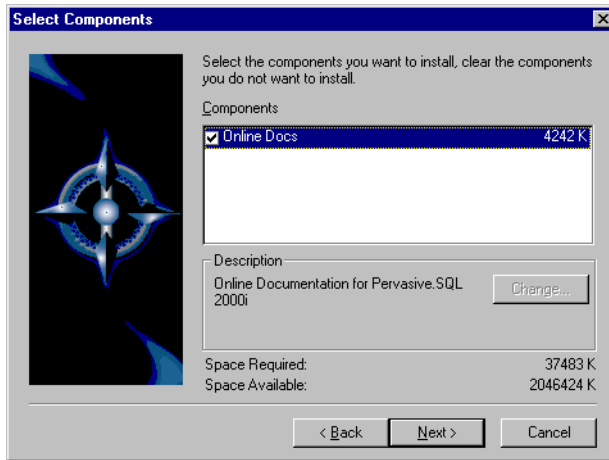
Select **Online Docs** to include online documentation in the product installation. When you are finished, click **Next**.

3 Optimize your system with the Pervasive System Analyzer

Next the Pervasive System Analyzer (PSA) launches. PSA scans your system for previous versions of Pervasive Software products.

Figure 5-27 Pervasive System Analyzer Welcome Screen
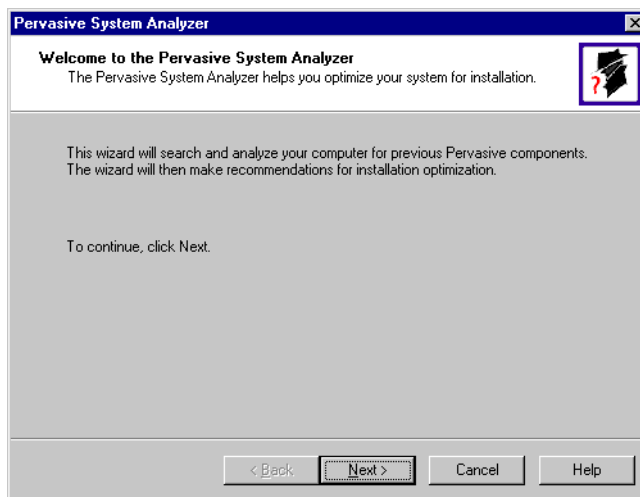
4 Analyze your system

Click **Next** to continue with the installation and begin system analysis. As PSA analyzes your system the following dialog box displays.

Figure 5-28 System Analysis Using PSA
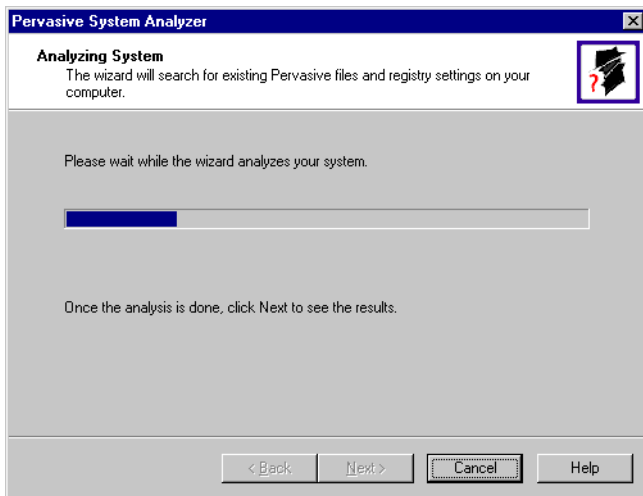
If you have previous versions of Pervasive Software products, PSA displays those products by the software version.

The progress bar completes as PSA analyzes your system. When PSA has completed the initial system analysis, click the **Next** button to continue and see the results.

If you do not have any previous versions of Pervasive Software products, the following dialog box displays. (See Chapter 4, "Installing Pervasive.SQL 2000i with No Previous Installation.")

Figure 5-29 Pervasive System Analyzer Results
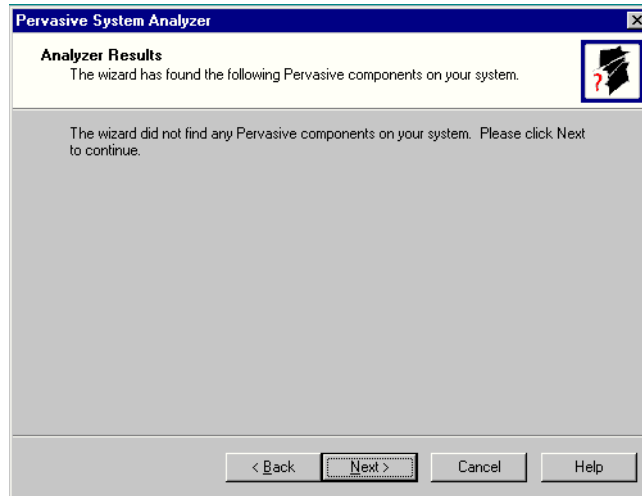
5   Choose the Program folder for menu items

Specify the folder in which the Start Menu items will be placed.
Pervasive installation suggests a name for the folder, which you
can accept by clicking **Next**.

If you have previously done an install on this machine, Install
recommends that you do not use the same folder name so that
your ability to cleanly uninstall the software in the future is
preserved.

Figure 5-30 Choose your Program Folder
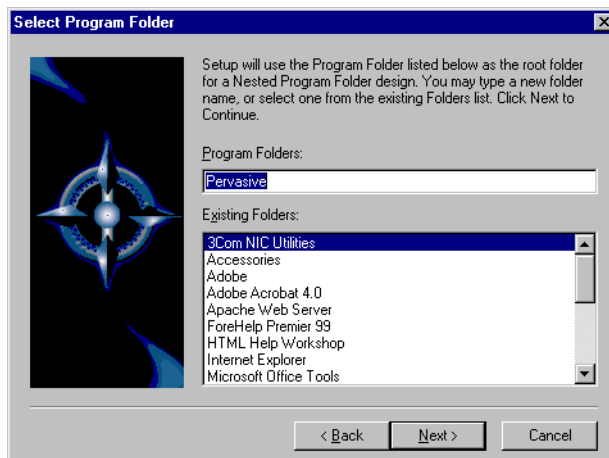
Click **Next** to accept the Program Folder name and display your current installation settings.

6 Verify your installation settings

Install now has enough information to start the installation process. A dialog box displays the list of options you have selected. If you want to change any of these selections, click **Back**.

If the displayed selections are correct, click **Next** to continue.

Figure 5-31 Your Custom Installation Selections
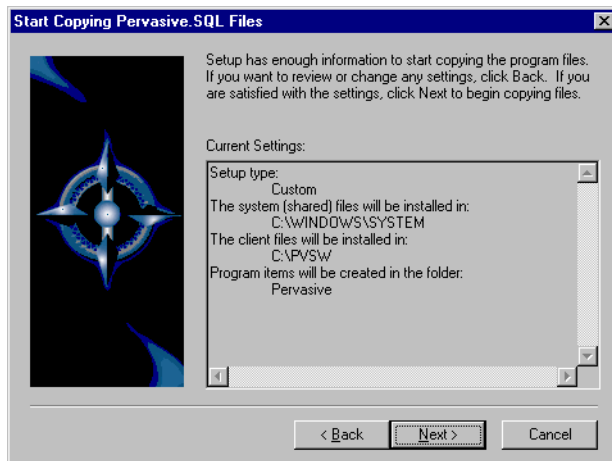
Pervasive.SQL 2000i setup begins to copy the necessary files to complete your specified installation requirements.

*Figure 5-32 Setup Begins Copying Files*
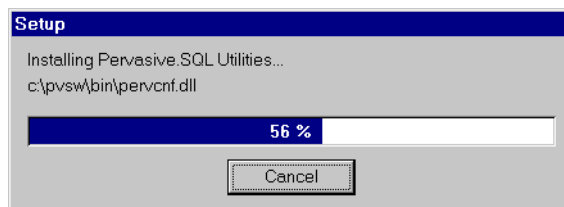
Note Only click **Cancel** if you want to stop installation and exit the Setup process.

The custom path rejoins the Typical installation path. See the step "Test the Transactional Engine," on page 5-20.

Upgrading from the Workstation Engine to the Workgroup Engine

> **Note** Be sure to close all files and stop the engine before you start this procedure.

➤ **To upgrade your Workstation Engine to the Workgroup Engine**

1 Launch the User Count Administrator

    a. On the **Start** menu, click **Programs,** then **Pervasive,** then **Pervasive.SQL**, then **Utilities**.

    b. Click **User Count Administrator** to open the utility**.**

The **User Count Administrator** utility appears:

Figure 5-33 User Count Administrator Dialog Box
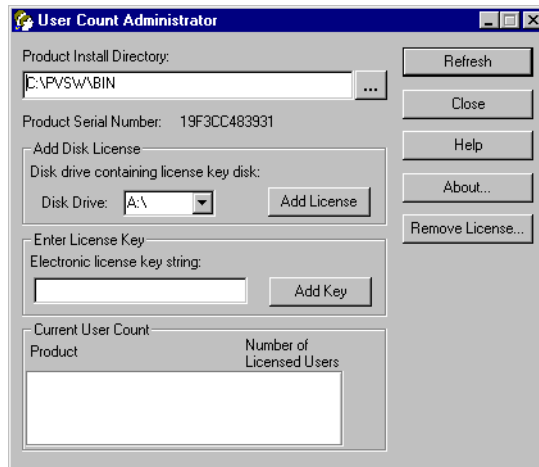
    c. Insert your workgroup license key diskette into drive A and click **Add License**.

> **Note** Workstation engines that are upgraded to Workgroup engines do not change their existing program icon labels.

When the update is complete, a dialog box informs you that you have increased your user count license.

d. Restart the engine.

The Workgroup engine is now available to the other workstations/clients for remote database requests.

Workstations on the network can access this new Workgroup engine by mapping a drive to it or by connecting to it through the Pervasive Control Center.

Refer to Chapter 11, "Configuring the Workgroup Engine," for information about how to set up your Workgroup engine.

Common Questions After Installing Pervasive.SQL 2000i

This section contains information that you should read after running the installation program. If you are having problems with your installation, go to Chapter 13, "Troubleshooting After Installation," or get help online from our Knowledge Base website at http://support.pervasive.com/eSupport.

If you are installing the Workgroup engine, you must install it on every computer that is expected to share data within your workgroup. Because every computer must have the Workgroup engine installed, there is no need to install the client software separately as there is for the Server engine.

How to Handle Data Source Names (DSNs)

After you have successfully installed the product, you have several choices depending on your situation, as shown in the table below.

Table 5-1   How to Proceed After Installing Workgroup/Workstation

If your situation is like this....	... then you should do this next:
You have existing Pervasive.SQL 7 DSNs already defined.	You must delete and re-create all existing Pervasive.SQL 7 DSNs before you can access existing databases. Follow the instructions provided in *Pervasive.SQL User's Guide* Chapter 2, sections "Deleting DSNs" and "Setting Up Database Access on a Windows Server or Workgroup/Workstation."
You have existing Pervasive.SQL 2000 DSNs already defined.	You should be able to access your databases by connecting to the existing DSNs. Follow the instructions provided in *Pervasive.SQL User's Guide*, Chapter 2, section "Accessing Data via ODBC From Other Applications."
You do not have any Pervasive.SQL DSNs defined	You should be able to connect to the sample DEMODATA database now. Refer to *Pervasive.SQL User's Guide* for general information on working with Pervasive.SQL. Refer to *Pervasive.SQL Advanced Operations Guide* for detailed information on working with databases and database engines.

How Do I Convert My Files From Previous Pervasive Products?

Converting your data files to 7.x format is not required, but you must convert them if you wish to take advantage of new features offered by the Pervasive.SQL 2000i engine.

Users of previous Btrieve versions: Use the Rebuild utility to convert your existing pre-6.0 or 6.x files to 6.x or 7.x format. For more information, refer to *Pervasive.SQL Advanced Operations Guide*.

How Do I Read the Online Documentation?

Pervasive.SQL 2000i documentation is in the form of WinHelp files for Win32 platforms. To read the documentation, select the desired book name from the **Documentation** subfolder of the **Pervasive** program folder on the **Start** menu.

Documentation for Win16 platforms is in the form of Adobe Acrobat (PDF) files. Shortcuts to these files are found in the **Pervasive.SQL** documentation file group. The installation program for an Acrobat reader is available on the Pervasive.SQL 2000i CD-ROM.

There is also a README file on the distribution media that has the latest changes and additions to Pervasive.SQL 2000i.

What Files Were Installed as Part of Pervasive.SQL 2000i?

Once installed, Pervasive.SQL 2000i files reside in the following default directory structure.

\WINNT\SYSTEM32	Microsoft system files Pervasive glue dlls (Win 9x only).
\PVSW	
\BIN	Pervasive.SQL engines, requesters and utilities.
\DEMODATA	Contains sample database.
\DOC	Online documentation.
\SAMPLES	Contains the sample database SAMPLE.BTR and UPPER.ALT, a sample Alternate Collating Sequence.

Working with User Counts

6

Understanding and Upgrading User Counts for Pervasive.SQL 2000i Workgroup

This chapter explains how to work with Pervasive.SQL 2000i licenses. This chapter contains the following topics:

Understanding Pervasive.SQL 2000i User Counts

Each user count allows one remote computer to connect to the Pervasive database engine.

What counts as a single user count?

A single workstation that is using TCP/IP or SPX protocol to access Pervasive.SQL on a given server counts as one user.

Users are counted by being given a unique Media Access Control (MAC) address. On a local area network or Ethernet, the MAC address is your computer's unique hardware address. Multiple applications on a single client computer are not counted as separate users.

Where do I obtain a user count license?

User counts are obtained as a software key issued by Pervasive Software or your application vendor if the Pervasive.SQL database is embedded.

Your Pervasive.SQL 2000i Workstation product includes no user count licenses and is intended to work in standalone mode. You can upgrade this product to a Workgroup version by purchasing a user license key. The Workgroup product contains a three-user license that is automatically applied during installation.

Your Pervasive.SQL 2000i Server product includes a user count key disk. In addition, you can obtain a key on disk by purchasing a User Count Increase kit from your Pervasive Software product distributor or your software vendor. You can increase the user count incrementally by a given number of users, usually 10 or more at a time for a client/server and one or more at a time above 3 users for the Workgroup engine, up to an unlimited number of users.

Once you approach 10 user licenses with the Workgroup product, you may want to consider upgrading to a Server product. The Server product is more cost effective for 10 or more user licenses and offers increased performance for large numbers of users. When you reach 200 users, you effectively have an unlimited user count environment.

How long does the upgrade take?

Using the User Count Administrator, you can upgrade user counts almost instantaneously: apply the user count key and restart the engine to activate the new user count.

The Pervasive.SQL 2000i Workgroup Engine includes a 3-user count license. This license is installed automatically with the product installation.

Upgrading from Workstation to Workgroup License

> **Note** Some applications are not configured to support a Workgroup configuration. Before you upgrade from Workstation to Workgroup, check with your application vendor to ensure that your application configuration will operate correctly in a Workgroup configuration.

The upgrade from a Workstation to a Workgroup license can be accomplished in just a few minutes. You will notice that Workstation engines that are upgraded to Workgroup engines do not change their existing program icon labels.

Any Workgroup engine that is acting as a Gateway must have the appropriate number of user counts applied. In a floating Gateway configuration, the user count upgrade must be applied to all Workgroup engines in the configuration. If you have a fixed Gateway configuration, you need only apply a user count upgrade to that particular engine.

➤ **To upgrade your Workstation Engine license to a Workgroup Engine license**

1 Close all files and stop the Workstation engine.

2 Launch the User Count Administrator from the **Start** menu by choosing:
   Programs | Pervasive | Pervasive.SQL 2000 | Utilities.

3 From the list of utilities, select **User Count Administrator**. The following dialog box opens:

Figure 6-1 User Count Administrator Dialog Box
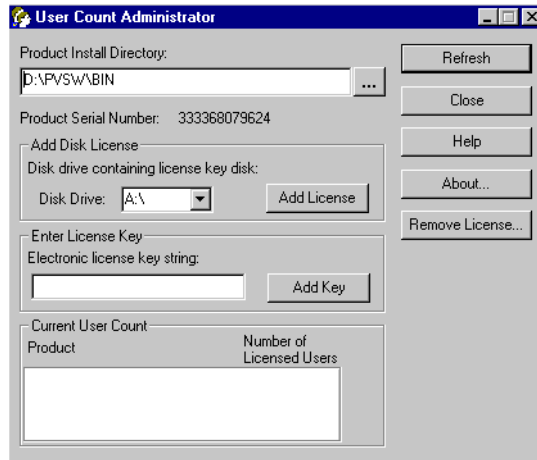
4   Insert your workgroup license key diskette into the diskette drive
and click **Add License**.

5   When the update is complete, a dialog box informs you that you
have increased your user count license.

6   Restart the engine.

How do other Workstations access the new Workgroup engine?

After you complete the upgrade instructions, the Workgroup engine
is available to the other workstations/clients for remote database
requests. Workstations on the network can access this new
Workgroup engine by mapping a drive to it or by connecting
through the Pervasive Control Center.

Applying Additional User Count Licenses

You can purchase additional user count licenses and apply them to your Workgroup engine at any time.

You are permitted to apply a new Workgroup license to all Workgroup engines in your workgroup. For example, if you have five computers with the Workgroup engine installed, and you buy a Workgroup 5-user license upgrade, you may apply that upgrade to all five computers.

To perform the upgrade, simply follow the instructions provided in "Upgrading from Workstation to Workgroup License" on page 6-4. The steps for applying an additional license are the same as specified for upgrading Workstation to Workgroup.

Note You must restart the engine after applying the license for it to take effect.

Using the User Count Administrator

As shipped, Pervasive.SQL 2000i Workgroup has an initial user count of three.

After product installation, you can use the User Count Administrator, a PCC utility, to increase the Pervasive.SQL 2000i user count. (The User Count Administrator can be run outside of the PCC, also.) You can obtain a disk key by purchasing a User Count Increase kit from one of the Pervasive Software product distributors.

You must run the User Count Administrator locally on each Workgroup engine. The figure below shows the **User Count Administrator** dialog box.

Figure 6-2 User Count Administrator Dialog Box
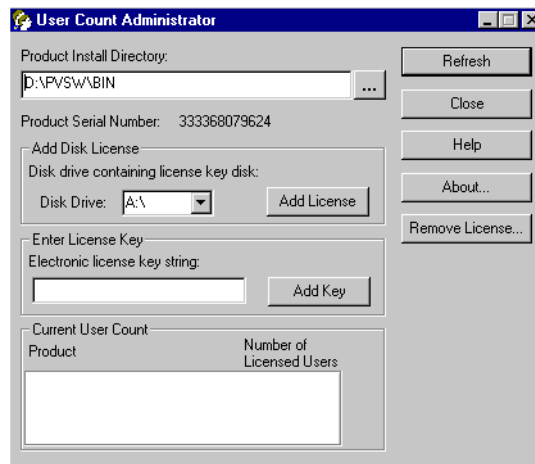
Click **Refresh** to update the information in the dialog box. (You may need to refresh the information if, for example, you change the Product Install Directory.)

Click **Remove License** to display a list of all incremental user counts you have applied. You can use this dialog box to remove individual user count licenses. Because the User Count Administrator displays only the highest version of each product in the current user count, you may want to remove a license if you inadvertently install a user count key from a higher version while still using a previous version of the product.

The table below describes the elements in the User Count Administrator dialogs.

Table 6-1   Elements in the User Count Administrator Dialog Box

Product Install Directory	Specify the path where the Pervasive.SQL engine is installed.
Product Serial Number	Displays the unique serial number assigned to the product.
Disk Drive	Specifies the disk drive letter that contains the user count key disk.
Add License	Applies the specified disk key.
Electronic License Key String	Specifies the 16 character license key string used to increase your user count. This string works only with special Web-generated OEM keys.
Add Key	Applies the specified electronic license key.
Product	Lists the products installed.
Number of Licensed Users	Lists the current user counts for the specified products.

Updating Your User Count and Licenses in Windows

➤ **To start the Pervasive.SQL User Count Administrator**

1   On the **Start** menu, click **Programs | Pervasive | Pervasive.SQL 2000 | Utilities.**

2   Select **User Count Administrator** from the list.

➤ **To increase your user count using a license key disk**

1   Start the **User Count Administrator** as described above.

2   If the user count license is in a disk drive other than A, then use the drop down box to specify the drive on which the license is located.

3   Click **Add License.**

➤ **To increase your user count using an electronic license key string**

1   Start the **User Count Administrator** as described above.

2   In the **Electronic License key string** field, enter your 16 character license key.

3   Click **Add Key.**

➤ **To remove a user count license**

1   Start the **User Count Administrator** as described above.

2   Click **Remove License.**

3   Select the license you want to remove from the list and click **Remove Single License.**

PERVASIVE®

S O F T W A R E
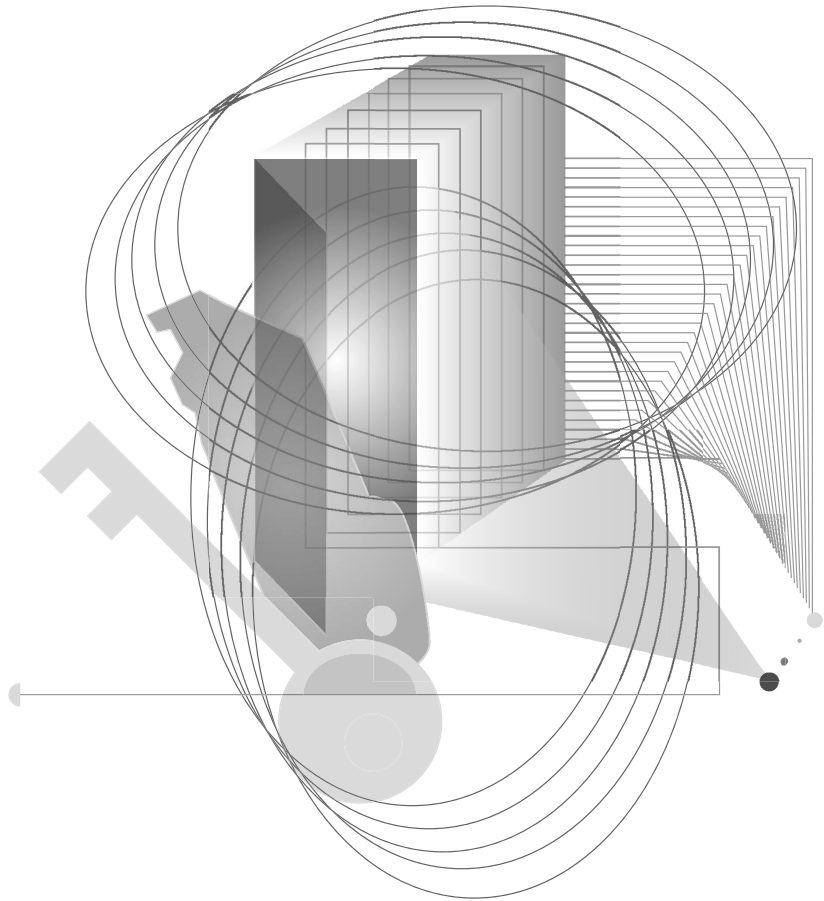
CUSTOMIZING THE WORKGROUP ENGINE

Network Settings for Workgroup Engine

How to Set up Your Workgroup Engine on Your Network and Understand the Default Settings

This chapter explains the default network settings for your database engine and how to customize them if your network requires different settings.

- "Determining What Kind of Network You Have" on page 7-2
- "Default Engine Settings" on page 7-3
- "Setting Up TCP/IP Support" on page 7-5
- "Setting Up SPX Support" on page 7-7
- "Setting Up NetBIOS Support" on page 7-9
- "Avoiding Unavailable Protocols" on page 7-10

Determining What Kind of Network You Have

This section explains how to determine the network protocol that you should use with the database engine. If you already know what protocol or protocols are supported on your network, you can skip this section.

Database Engine on Windows

If your network is 100% Microsoft, and you have a Server engine, then your network probably uses TCP/IP. The Server engine does not support NetBIOS.

It is possible to run Btrieve applications over SPX on Microsoft networks, but ODBC applications over SPX require a NetWare server for name resolution. If you do not have a NetWare server on your network, you may not use SPX for ODBC applications.

If your network is 100% Microsoft, and you are using Workgroup engines, then you can use either NetBIOS or TCP/IP.

Default Engine Settings

This section explains what server Configuration settings are related to networking support, and what the default values mean.

You can change these settings within PCC by double-clicking **Configuration** for the given database engine, then double-clicking **Server**, then **Communication Protocols**.

AutoReconnect Timeout

Default: 180 seconds
If Pervasive AutoReconnect is enabled, this setting specifies how long the engine and client attempt to contact each other after a network interruption has occurred. By default, the client, server and multiple workgroup engines attempt to reconnect for three minutes before giving up. If the value of **Enable AutoReconnect** is **Off**, then the **AutoReconnect Timeout** value is ignored.

Enable AutoReconnect

Default: Off
This setting determines whether the client, server or multiple workgroup engine attempts to reconnect to each other in the event of a network outage. If set to **On**, it allows the database connections to recover from intermittent or temporary network interruptions. If this value of this setting is **Off**, then the client returns a status code to the application immediately upon any failure to connect to the server, and the connection context is not preserved.

If you have this setting turned **On**, you can specify how long the client and database engine should attempt to reconnect by using the setting **AutoReconnect Timeout**.

Listen IP Address

Default: 0.0.0.0
This setting specifies the IP address of the network interface card that the MicroKernel should listen on when the server computer has two network cards installed. This value is ignored if the server computer has only one network card, or if the value of **TCP/IP Multihomed** is **On**. This value is also ignored if TCP/IP is not in use by the database engine.

NetBIOS Port

Default: 66
This value specifies the port number the Workgroup engine listens on when using NetBIOS communications. The Server engine does not support NetBIOS, so this setting is ignored by the Server engine.

Note This port setting must be the same for all Workgroup engines on the network. Workgroup engines with different port settings cannot communicate with one another.

Supported Protocols

Default: SPX, TCP/IP, NetBIOS
This setting specifies the vendor protocols that the database engine should attempt to use. When more than one protocol is specified, upon start up, the engine attempts to connect on all specified protocols. The protocol that connects first is then used for the remainder of the session.

You can often improve initial connection performance by removing the protocols that are not used on your network. For example, if you have a Server engine on an all-TCP/IP network, removing NetBIOS and SPX support may reduce the wait time during initial connections to the engine.

Note The Server engine does not support the NetBIOS protocol.

TCP/IP Multihomed

Default: Off (Server); On (Workgroup)
If your server computer has two network interface cards installed, you can use this setting to specify whether the server engine should listen on both network connections. If your computer has two network cards and this setting is turned off, you must use the setting **Listen IP Address** to specify which network card the server engine should use.

Setting Up TCP/IP Support

By default, TCP/IP is supported between Pervasive.SQL clients and remote engines or between multiple workgroup engines. If you have modified the default settings or need to verify that TCP/IP support is available, refer to this section.

> **Note** In order to perform any of the tasks in this section, you must be seated at the console of the machine running the Workgroup engine. You cannot remotely configure the Workgroup engine.

➤ Enabling TCP/IP Support

Follow this procedure to ensure that the database engine can communicate with clients over TCP/IP networks.

1   From the **Start** menu, choose **Programs | Pervasive | Pervasive Control Center** (PCC).

2   In the PCC window, double-click on **Pervasive.SQL 2000 Engines** to display a list of the engines registered with PCC.

3   Double-click on the icon representing your target engine. Double-click **Configuration**. Login if prompted. Double-click **Server.** Click on **Communication Protocols**. In the window to the right, double-click on **Supported Protocols**.

4   A smaller dialog box appears, showing a list of **Available** protocols and a list of **Selected** protocols. If the list of Selected protocols includes the value **Microsoft TCP/IP,** then TCP/IP is already supported.

   If the list of Selected protocols does not include **Microsoft TCP/IP**, then you should click on this value in the left-hand window, and then click the > button to move it into the list of Selected protocols. Once the value appears in the right-hand window, click **OK**.

5   Click ❗ to apply the changes. You need to restart the engine for the changes to take effect.

6   Remember that you also need to confirm that the client software on your other Workgroup computers are configured to use TCP/IP, as well. Please refer to Chapter 8, "Client Network Configuration."

➤ Enabling Multihomed TCP/IP Support

Follow this procedure when your Windows server computer has two network cards installed. By default, Pervasive.SQL assumes that you only have one network card installed.

1   From the **Start** menu, choose **Programs | Pervasive | Pervasive Control Center** (PCC).

2   In the PCC window, double-click on **Pervasive.SQL 2000 Engines** to display a list of the engines registered with PCC.

   If the engine you want to configure is not listed, right-click on **Pervasive.SQL 2000 Engines** and choose **Register New Engine**.

3   Double-click on the icon representing your target engine. Double-click **Configuration**. Login if prompted. Double-click **Server**. Click on **Communication Protocols**. In the window to the right, double-click on **TCP/IP Multihomed**.

4   If you want the server engine to listen for client connections on both network cards, change the value to **On**. If you want the server engine to listen on only one card, change the value to **Off**. Click **OK**. If you only have one network card, this setting is ignored.

5   If your server computer has two network cards, and you set the value of **TCP/IP Multihomed** to **Off**, you must double-click the setting **Listen IP Address** and specify the TCP/IP address of the card you want the database engine to listen to. If you do not specify an IP address, the database engine will not receive communications from either network card.

6   Click ❗ to apply the changes. You need to restart the server engine for the changes to take effect. You do not need to make any changes to client settings.

Setting Up SPX Support

By default, SPX is supported between Pervasive.SQL clients and servers. If you have modified the default settings or need to verify that SPX support is available, refer to this section.

Your network's SPX Frame Type setting does not have any effect on Pervasive.SQL.

> **Note** In order to perform any of the tasks in this section, you must be seated at the console of the machine running the Workgroup engine. You cannot remotely configure the Workgroup engine.

➤ Enabling SPX Support

Follow this procedure to ensure that the server engine can communicate with clients over SPX networks.

> **Note** In an all-Microsoft environment, SPX can only be used with pure Btrieve applications. With ODBC applications, a NetWare server must be present to provide name resolution for SPX. Pure Btrieve applications do not require name resolution with SPX.

1   From the **Start** menu, choose **Programs | Pervasive | Pervasive Control Center** (PCC).

2   In the PCC window, double-click on **Pervasive.SQL 2000 Engines** to display a list of the engines registered with PCC.

3   Double-click on the icon representing your target server. Double-click **Configuration**. Login if prompted. Double-click **Server**. Click on **Communication Protocols**. In the window to the right, double-click on **Supported Protocols**.

4   A smaller dialog box appears, showing a list of **Available** protocols and a list of **Selected** protocols. If the list of Selected protocols includes the value **Microsoft SPX II** or **Novell SPX II**, then SPX is already supported.

If the list of Selected protocols does not include one of the above values, then you should click on this value in the left-hand window, and then click the > button to move it into the list of Selected protocols. Once the value appears in the right-hand window, click **OK**.

5   Click ! to apply the changes. You need to restart the server engine for the changes to take effect.

6   Remember that you also need to confirm that the client software on your other Workgroup computers are configured to use SPX, as well. Please refer to Chapter 8, "Client Network Configuration."

Setting Up NetBIOS Support

By default, NetBIOS is supported between Pervasive.SQL clients and Workgroup engines. If you have modified the default settings or need to verify that NetBIOS support is available, refer to this section.

> **Note** In order to perform any of the tasks in this section, you must be seated at the console of the machine running the Workgroup engine. You cannot remotely configure the Workgroup engine.

➤ Enabling NetBIOS Support

Follow this procedure to ensure that the database engine can communicate with clients over NetBIOS networks.

1 From the **Start** menu, choose **Programs | Pervasive | Pervasive Control Center** (PCC).

2 In the PCC window, double-click on **Pervasive.SQL 2000 Engines** to display a list of the engines registered with PCC.

3 Double-click on the icon representing your target engine. Double-click **Configuration**. Login if prompted. Double-click **Server.** Click on **Communication Protocols**. In the window to the right, double-click on **Supported Protocols**.

4 A smaller dialog box appears, showing a list of **Available** protocols and a list of **Selected** protocols. If the list of Selected protocols includes the value **Microsoft NETBIOS**, then NetBIOS is already supported.

   If the list of Selected protocols does not include **Microsoft NETBIOS**, then you should click on this value in the left-hand window, and then click the > button to move it into the list of Selected protocols. Once the value appears in the right-hand window, click **OK**.

5 Click ❗ to apply the changes. You need to restart the engine for the changes to take effect.

6 Remember that you also need to confirm that the client software on your other Workgroup computers are configured to use NetBIOS, as well. Please refer to Chapter 8, "Client Network Configuration."

Avoiding Unavailable Protocols

It may be possible to improve performance on the initial connection to the database by disabling database communications support for any protocols that are not available on your network or that you do not wish to use.

Note In order to perform any of the tasks in this section, you must be seated at the console of the machine running the Workgroup engine. You cannot remotely configure the Workgroup engine.

➤ **To Remove Network Protocol Support**

Note This procedure does not affect your operating system configuration in any way. This procedure only prevents the database communications system from attempting communications on unavailable or undesired protocols.

1 From the **Start** menu, choose **Programs | Pervasive | Pervasive Control Center** (PCC).

2 In the PCC window, double-click on **Pervasive.SQL 2000 Engines** to display a list of the engines registered with PCC.

3 Double-click on the icon representing your target server. Double-click **Configuration**. Login if prompted. Double-click **Server**. Click on **Communication Protocols**. In the window to the right, double-click on **Supported Protocols**.

4 A smaller dialog box appears, showing a list of **Available** protocols and a list of **Selected** protocols. If the list of Selected protocols includes a value that is not supported on your network or that you do not wish to use, click on the value and then click the < button to move it into the list of **Available Protocols**. Repeat for any other protocols you do not wish to use. You must leave at least one protocol in the **Selected protocols** list.

Once the given value(s) appear in the left-hand list, click **OK**.

5 Click ![!] to apply the changes. You need to restart the server engine for the changes to take effect.

6   Remember that you also need to confirm that your client computers are configured to use the protocol remaining in the **Selected protocols** list. Please refer to Chapter 8, "Client Network Configuration."

Client Network Configuration

8

How to Configure the Network Settings for your Pervasive.SQL Clients

To access network files from a workstation using a Pervasive.SQL application, you must use the appropriate requester at that workstation. Your application's Pervasive.SQL calls go through the Requester, which sends them to the server for processing and then returns the reply to your application.

This chapter contains the following sections.

Default Settings

This section lists the settings that are relevant to client networking.

TCP/IP is the default protocol for Pervasive.SQL. That is, TCP/IP is tried first if more than one protocol is set to be available.

Enable AutoReconnect

Default: Off
This setting determines whether the client and server attempt to reconnect to each other in the event of a network outage. If set to **On**, it allows the database connections to recover from intermittent or temporary network interruptions. If this value of this setting is **Off**, then the client returns a status code to the application immediately upon any failure to connect to the server, and the connection context is not preserved.

If you have this setting turned **On**, you can specify how long the client and database engine should attempt to reconnect by using the setting **AutoReconnect Timeout** in the Server configuration.

Note This feature is not available on 16-bit clients.

Supported Protocols

Default: SPX, TCP/IP, NetBIOS
This setting specifies the vendor protocols that the database engine should attempt to use. When more than one protocol is specified, upon start up, the engine attempts to connect on all specified protocols. The protocol that connects first is then used for the remainder of the session.

You can often improve initial connection performance by removing the protocols that are not used on your network. For example, if you have a Server engine on an all-TCP/IP network, removing NetBIOS and SPX support may reduce the wait time during initial connections to the engine.

On UNIX platforms, TCP/IP is the only supported protocol on the server.

Note The Server engines do not support the NetBIOS protocol.

TCP/IP Timeout for Communications Requester

Default: 15 seconds
This setting specifies the number of seconds the requester should
wait for a TCP/IP connect request to succeed before timing out.

How to Configure the Pervasive Clients

This section contains information on configuring Pervasive client software.

- Configuring the Win32 Requester
- Configuring the Win16 Requester

Note For information on configuring DOS requesters, see "Using the DOS Requesters" on page 8-26.

Configuring the Win32 Requester

The Win32 Requester serves as the Pervasive.SQL client for Windows 9X/ME and Windows NT/2000 workstations. You configure the Win32 Requester using the Configuration Utility within the Pervasive Control Center.

➤ **To Configure the Win32 Requester**

1 Click **Start**, point to **Programs** and select **Pervasive**.

2 Select **Pervasive Control Center**.

3 Click **Pervasive.SQL Engines** and select the local machine.

4 Select **Configuration** and then **Client**.

5 Refer to the *Pervasive.SQL Advanced Operations Guide* in section "Configuration Reference" for detailed information on the Configuration utility and settings.

Win32 Configuration Notes

WOW Application Users: The Win16 Requester provides thunking over to the Win32 Requester, which then handles your request as it does one from a Win32 application. To use these thunking capabilities, be sure to use the requesters included with the Pervasive.SQL package, and not those from an earlier version of Btrieve.

Configuring the Win16 Requester

The Win16 Requester serves as the Btrieve client for Windows 3.x environments. It can also be used in conjunction with Thunking in Windows 95 and Windows NT workstations.

> **Note** Clients using the 16-bit requesters will have only transactional access to the data files. No relational access will be available.

➤ To Configure the Win16 Requester from Windows 3.x:

1 Double-click on the **Pervasive SQL 2000** Program Group.

2 Double-click on Setup (Win16).

3 Refer to the *Pervasive.SQL Advanced Operations Guide* in section "Configuration Reference" for detailed information on the Configuration utility and settings.

➤ To Configure the Win16 Requester from Windows 9X/ 98/ME/NT/2000:

1 Click Start, point to Programs, then Pervasive, and select Pervasive Control Center.

2 Expand **Pervasive.SQL 2000 servers** and expand the desired server.

3 Expand **Configuration** and expand **Client-16 bit**.

4 Refer to the *Pervasive.SQL Advanced Operations Guide* in section "Configuration Reference" for detailed information on the Configuration utility and settings.

Win16 Configuration Notes

To access a Windows NT server, your workstation must have the Microsoft Network Client for DOS and Windows loaded. You also need the NetWare SPX protocol stack. Pervasive Software recommends using the ODI drivers from Novell. You do not need to run NETX.EXE or be logged in to a NetWare server.

WOW Application Users: You must install both the Win16 Requester and the Win32 Requester. The Win16 Requester provides thunking over to the Win32 Requester, which then handles your request as it does one from a Win32 application. To use these thunking capabilities, be sure to use the requesters included with this package, and not those from an earlier version of Btrieve. Also, be sure to enable the Requester's **Use Thunk** option.

Network Path Formats Supported by Pervasive Requesters

When using your Requester, you connect to the Pervasive server engine to access data files. This section shows the variations on network file syntax you can use to access files on your network using Btrieve or SQL applications.

Pervasive.SQL supports the Universal Naming Convention (UNC) and Drive path formats (explicit and current) across the majority of operating environments, including:

Table 8-1   Supported UNC and Drive Path Formats

Application Types	Environments	Network Client	Novell Login Types
Transactional	Windows (32-bit)	Microsoft	Bindery
Relational	Windows (16-bit)	Novell	NDS
	DOS		

For more information on the path formats, see the sections that follow:

- "Universal Naming Convention (UNC) Path Formats" on page 8-6
- "Drive-based Formats" on page 8-7
- "NetWare Specific Formats" on page 8-7
- "UNIX Path Formats" on page 8-12

Universal Naming Convention (UNC) Path Formats

The following UNC path formats are supported on all clients to all servers:

```
\\ServerName\volume\path\file
\\ServerName\volume:[\]path\file
```

UNC syntax is resolved correctly regardless of the actual type of network operating system (NOS) that is operating on the target server.

Note In all instances above, backslashes (\) can be interchanged with forward slashes (/) except for the double backslash (\\). The syntax [\] indicates that the backslash is optional.

Drive-based Formats

The following drive representations are supported on all clients to all servers:

Note Do not map drive letters to the \\Tree\VolumeObject format found in Network Neighborhood. See "NetWare Directory Services (NDS) Formats" on page 8-7.

```
drive:file
drive:[\]path\file
file
[\]path\file
..\file
```

NetWare Specific Formats

NetWare path formats on all clients only to NetWare servers:

```
volume:[\]path\file (ServerName is derived from current
    drive)
```

NetWare Directory Services (NDS) Formats

Pervasive.SQL provides enhanced Novell Directory Service support over previous versions of Pervasive.SQL through complete integration with Novell's IntranetWare Client for Windows NT, Windows 95, and Windows 3.x, and Microsoft's Service for Novell Directory Services for Windows 95 (MSNDS). These enhancements provide the following benefits to users and network administrators:

- Better integration with NDS authentication capabilities when accessing the Pervasive.SQL Server Engine for NetWare.
- Ability to resolve server names into network address by querying NDS instead of relying on the NetWare Bindery.
- Support for NDS Volume and Directory Map objects in Btrieve file operations.

> **Note** Pervasive does not support the \\Tree\VolumeObject format
> found in Network Neighborhood. See "NetWare Directory Services
> (NDS) Formats" on page 8-7.

If you want to run 32-bit applications in an NDS environment, the
Pervasive.SQL requesters work best with the Novell IntranetWare
clients installed on the workstation. The following sections detail the
level of support that Pervasive offers.

Support for Novell Clients for NetWare

Pervasive.SQL was tested with the following Novell clients:

- IntranetWare Client for Windows NT v4.11
- IntranetWare Client for Windows 95 v2.2
- IntranetWare Client for DOS/Windows 3.x v2.2

Each of these clients provide the set of NDS specific APIs needed by
the Pervasive.SQL requesters to integrate fully with the NDS
environment. All Pervasive.SQL requesters will detect these clients'
presence automatically; no additional Pervasive.SQL configuration
changes are required. When these clients are available to the
Pervasive.SQL requesters, the following operations are supported:

- Network login via NDS (bindery context does not need to be set
  on the target server)
- NetWare server address resolution for SPX addresses
- Drive letters mapped to NDS Volume Objects through either the
  Map utility or Network Neighborhood may be used to specify
  Btrieve file names or SRDE dictionary and data directories
- NDS Volume Objects and Directory Map Objects may be used to
  specify Btrieve file names with Win32, and Win16 applications

The file name and path formats supported by the Pervasive.SQL requesters through the Novell clients are listed in the table below:

Table 8-2   File name and Path formats for Novell Clients for NetWare

Format	Btrieve Apps	Notes	SQL Apps
`<drive letter>:[\]path\file`	✔	(1)	✔
`\\server\volume\path\file`	✔		✔
`\\server\volume:[\]path\file`	✔		✔
`server\volume:[\]path\file`	✔		✔
`volume:[\]path\file`	✔	(4)	✔
`VolumeObject:[\]path\file`	✔	(2) (3)	✘
`\\VolumeObject\path\file`	✔	(2) (3)	✘
`DirectoryMap:[\]path\file`	✔	(2) (3)	✘
`\\DirectoryMap\path\file`	✔	(2) (3)	✘
`\\Tree\VolumeObject\path\file`	✘		✘
`\\Tree\DirectoryMap\path\file`	✘		✘

1 <drive letter> can be redirected to a server\volume:, or NDS Volume or Directory Map object.
2 Not supported by the Btrieve DOS requester.
3 See Novell documentation for rules about specifying NDS context when using NDS objects.
4 Requires that the current drive is mapped to the target NetWare server.

Support for Microsoft Clients for NetWare

Pervasive.SQL was tested with the following Microsoft clients for NetWare:

- Microsoft Service for Novell Directory Services (Windows 95)

- Microsoft Client Service for NetWare (Windows NT/2000). This client limits your NDS support; Pervasive recommends using the Novell IntranetWare client.

- Microsoft Client Service for NetWare (Windows 95)

Note The Win16 Requesters require the Novell components NWCALLS.DLL and NWNET.DLL. Win32 Requesters using the Microsoft Service for NDS also require these files.

The file name and path formats supported by the Pervasive.SQL requesters using Microsoft's clients for NetWare on Windows 9X/ME and Windows NT/2000 are listed in the following table.

Table 8-3   File name and Path formats for Microsoft Clients for NetWare

Format	Btrieve Apps Win16	Btrieve Apps Win32	Notes	SQL Apps
`<drive letter>:[\]path\file`	✔	✔	(1)	✔
`\\server\volume\path\file`	✔	✔		✔
`\\server\volume:[\]path\file`	✔	✔		✔
`server\volume:[\]path\file`	✔	✔		✔
`volume:[\]path\file`	✔	✔	(4)	✔
`VolumeObject:[\]path\file`	✔	✘	(2) (3) (5)	✘
`\\VolumeObject\path\file`	✔	✘	(2) (3) (5)	✘
`DirectoryMap:[\]path\file`	✔	✘	(2) (3) (5)	✘
`\\DirectoryMap\path\file`	✔	✘	(2) (3) (5)	✘
`\\Tree\VolumeObject\path\file`	✘	✘		✘
`\\Tree\DirectoryMap\path\file`	✘	✘		✘

1 <drive letter> can be redirected to a server\volume:, or NDS Volume or Directory Map object on Windows 95 with MSNDS. On Windows NT 4.0 with CSNW, <drive letter> can only be redirected to a server\volume: name; using Network Neighborhood to map a <drive letter> to an NDS object is not supported for Win32 applications.
2 Not supported by the Btrieve DOS requester.

3 Supported for applications using the Win16 Pervasive requesters only (with no thunking to 32-bit).

4 Requires that the current drive is mapped to the target NetWare server.

5 See Novell documentation for rules about specifying NDS context when using NDS objects.

Implementation Notes for the Microsoft NetWare Clients

The Microsoft Service for Novell Directory Services (MSNDS) allows Windows 95 workstations to login to NDS trees and redirect drive letters to NDS servers. MSNDS also supports Windows 16-bit 'NDS-aware' applications by supporting the Novell Windows 16-bit client DLLs, NWCALLS.DLL and NWNET.DLL. MSNDS does not provide 32-bit API support for NDS-aware applications. Windows 95, however, does provide a thunk mechanism allowing 32-bit applications to call 16-bit applications. The Pervasive.SQL requesters take advantage of that mechanism to provide NDS support, although currently the 32-bit requesters do not allow NDS object names to be used in UNC-style paths as they do with the Novell clients. On Windows 95 workstations with MSNDS installed, the Pervasive.SQL requesters support the following NDS operations:

■ Network login via NDS (bindery context does not need to be set on the target server)

■ NetWare server address resolution for SPX addresses

■ Drive letters mapped to NDS Volume Objects or Directory Map Objects through either the Map utility or Network Neighborhood may be used to specify Btrieve file names, or SRDE dictionary and data directories when logging into a Pervasive.SQL database, only with WIN16 applications

■ NDS Volume Objects and Directory Map Objects may be used to specify Btrieve file names, only with Win16 applications

Like MSNDS, the Microsoft Client Service for NetWare (CSNW) on Windows NT 4.0 allows Windows NT 4.0 workstations to log in to NDS trees and redirect drive letters to NDS servers, as well as supporting Windows 16-bit 'NDS-aware' applications by supporting the Novell Windows 16-bit client DLLs. CSNW does not provide 32-bit API support for NDS-aware applications. Unlike Windows 95, Windows NT does not provide a thunk mechanism allowing 32-bit applications to initiate calls to 16-bit applications (it does allow 16-bit applications to call 32-bit applications, which the Pervasive requesters support). The lack of support for 32-bit 'NDS-aware' applications coupled with the lack of a thunk mechanism for 32-bit

to 16-bit calling conventions significantly reduces the ability of the Pervasive.SQL requesters to support NDS environments. On Windows NT 4.0 workstations with CSNW installed, Pervasive.SQL supports the following NDS operations:

■ Network login via NDS for DOS applications and Win16 applications that do not thunk to the Win32 requester (bindery context does not need to be set on the target server)

■ Network login via NDS for Win32 applications, when accessing the Pervasive.SQL for NetWare server engine using SPX as the transport

■ NetWare server address resolution for SPX addresses

■ NDS Volume Objects and Directory Map Objects may be used to specify Btrieve file names for Win16 applications that do not thunk to the Win32 requester

UNIX Path Formats

Incoming paths on a UNIX server using Samba will be processed as follows in order of relative priority:

Share names

`\\<server>\<sharename>\<path>`

The smb.conf file must be configured to accept `<sharename>`, otherwise it will default to the following:

Absolute paths

`\\<server>\<absolute_path>`

If the smb.conf file is not configured properly or not found on the target server, the absolute path is used.

For more information on the UNIX versions of Pervasive.SQL 2000i, see the "Unix Supplementary Documentation" section of *Pervasive.SQL User's Guide*.

Using TCP/IP to Connect to a NetWare Server

This section documents the use of TCP/IP when connecting to a Pervasive.SQL server running on a NetWare machine.

Configuring Pervasive.SQL 2000 to use TCP/IP

TCP/IP is the default protocol when connecting to NetWare servers, meaning that if both TCP/IP and SPX protocols are available, TCP/IP is the first one the Communications Requester attempts to use.

If you have a working TCP/IP setup from your client to your NetWare server, you should not have to do anything in addition to make Pervasive.SQL work with TCP/IP and NetWare.

Configuring a Client For the Server IP Address

You must have your server's IP address configured properly on your clients for Pervasive.SQL to function. There are two ways to do this for Windows clients.

➤ **To configure your clients to use the server IP address, do one of the following:**

For Windows 9X/NT clients:

■ Click **Start**, point to **Settings** and select **Control Panel**. Double-click **Network**, select Protocols, select **TCP/IP** and click **Properties**. Click the **DNS** tab. Enable DNS and enter the appropriate server information.

For Windows 2000 clients:

■ Click **Start**, point to **Settings** and select **Control Panel**. Double-click **Network and Dial-up Connections**, select **Local Area Connection** and click **Properties**. From the component list, select **TCP/IP** and click **Properties**. Enable DNS and enter the appropriate server information.

or

■ Enter your Server IP address/name in the hosts file:
`\WINNT\SYSTEM32\DRIVERS\ETC\HOSTS` file (Windows NT/ 2000).
`\WINNT\SYSTEM\DRIVERS\ETC\HOSTS` file (Windows 95).

8-13

Using SPX to Connect to a NetWare Server

This section documents the use of SPX when connecting to a Pervasive.SQL server running on a NetWare machine.

Configuring Pervasive.SQL 2000 to use SPX

SPX is a native protocol to NetWare. There should be no special requirements for operating Pervasive.SQL with SPX on NetWare. However, note the following:

SPX is not the default protocol for Pervasive.SQL. Consequently, if you have both TCP/IP and SPX installed, Pervasive.SQL attempts to use TCP/IP first. If you want SPX to be used, follow these steps:

➤ **To Configure your Clients to use SPX with NetWare Servers if Both TCP/IP and SPX are Installed:**

1 Click **Start**, point to **Programs** and select **Pervasive**.

2 Select **Pervasive Control Center**.

3 Expand **Pervasive.SQL 2000i Engines** and expand the desired engine.

4 Expand **Configuration** and expand **Client**.

5 Select the **Communication protocols** and, in the right-hand pane, right-click on **Supported protocols**.

6 Highlight **Microsoft TCP/IP** from the list of selected protocols and click the < button to move it off the selected list.

7 Click **OK**.

To use SPX with NetWare, also ensure that your network packet size is appropriate for your network. By default the **Receive Packet Size** option of the Configuration utility (detailed in the *Pervasive.SQL Advanced Operations Guide*) is set to 1500, which is the correct value for Ethernet networks. If, however, you have a Token-ring network, use the Configuration utility to modify your **Receive Packet Size** to 4096.

Using Either Microsoft or Novell's SPX when Connecting to NetWare Servers

If you select Novell SPX using the **Supported Protocols** setting in the Configuration utility, this means that all Pervasive SPX calls will be made through the Novell component NWIPXSPX.DLL. Both Novell and Microsoft support this DLL, so you could be using Microsoft SPX even though the Pervasive setting shows that you are using Novell SPX.

➤ **To determine which version of SPX (Microsoft or Novell) you are using from your Windows workstation:**

In Windows NT:

■ Click **Start**, point to **Settings** and select **Control Panel**. Double-click **Network**, and look for **IPX/SPX** in the list. If the entry says "IPX/SPX compatible-protocol", you have Microsoft's SPX installed. Otherwise, you are using Novell's IPX/SPX protocol.

In Windows 2000:

■ Click **Start**, point to **Settings** and select **Control Panel**. Double-click **Network and Dial-up Connections**, select **Local Area Connection** and click **Properties**. Look for **IPX/SPX** in the list. If the entry says "IPX/SPX compatible-protocol", you have Microsoft's SPX installed. Otherwise, you are using Novell's IPX/SPX protocol.

Pervasive.SQL supports both the Microsoft and Novell versions of SPX. If you do not use the Novell Client 32 to communicate with the NetWare server, you could encounter a Status Code 3119 when you first attempt a connection to the NetWare server. Subsequent connections should clear that condition.

Changing the Default Server Listening Port (Advanced)

Most users should never have to use this procedure. This procedure should only be performed by a system administrator.

In normal TCP/IP operation, clients use a particular port number to communicate with servers. Pervasive.SQL by default listens on port 1583. For TCP/IP to function properly, all clients and all servers must use the same port.

Pervasive recommends that you do not modify the server listening port unless you have a conflict with port number of 1583. If you do

have a conflict, you can override the Pervasive Database port settings with an entry in the SERVICES file for each machine. The SERVICES file is located at:

```
SYS:\ETC\SERVICES
```

For example, if you wanted to change the listening port to 80 for Pervasive.SQL, you would add the following entry in your SERVICES file:

```
psql 80/tcp
```

After changing your server listening port, you must stop and restart your Pervasive.SQL services for the changes to take effect. See "Using Pervasive.SQL" in *Pervasive.SQL User's Guide*.

Note Entries in the NetWare SERVICES file are NOT case-sensitive.

Using TCP/IP to Connect to a Windows NT/2000 Server

This section documents the use of TCP/IP when connecting to a Pervasive.SQL server running on a Windows NT/2000 machine.

Tasks   **Configuring a Client for the Server IP Address**

You must have your server's IP address configured properly on your Pervasive.SQL requesters for Pervasive.SQL to function. You can do this in one or two ways depending on which client requester platform you are using:

➤ **To configure your clients to use the server's IP address, do one of the following:**

For Windows 95/NT clients:

■   Click **Start**, point to **Settings** and select **Control Panel**. Double-click on the **Network icon**, double-click on **TCP/IP** in the list. Click the **DNS Configuration** tab. Enable DNS by selecting the corresponding button and entering the appropriate server information. This is the preferred method for Windows clients.

For Windows 2000 clients:

■   Click **Start**, point to **Settings** and select **Control Panel**. Double-click on the **Network and Dial-up Connections** icon, double-click on **Local Area Connection**, and click on the **Properties** button. From the component list, select **TCP/IP** and click the **Properties** button. Enable DNS by selecting the corresponding button and entering the appropriate server information.

or

■   Enter your Server IP address/name in the hosts file:
`\WINNT\SYSTEM32\DRIVERS\ETC\HOSTS` file (Windows NT/2000).
`\WINDOWS\SYSTEM\DRIVERS\ETC\HOSTS` file (Windows 95).

Preventing the Windows Dial-Up Network Dialog Box from Displaying When Using a Pervasive Application with TCP/IP

The **Windows Dial-Up Networking** dialog box can display when a TCP/IP request is made to Windows. Usually, this is to make an

Internet connection, but this feature can be an annoyance when using Pervasive applications and TCP/IP.

➤ To Prevent the Dial-Up Networking Dialog Box from Displaying Automatically:

1 Click **Start**, point to **Settings** and select **Control Panel**.

2 Double-click the **Internet** icon in Control Panel.

3 Click the **Connection** tab.

4 Clear the check box titled **Connect to the Internet As Needed**. In Windows 2000, this button is titled **Dial whenever a network connection is not present**.

Note While this stops the dialog box from displaying with Pervasive applications, this also has the side effect that other applications such as Internet browsers will no longer automatically spawn the **Dial-Up Networking** dialog box when a connection to the Internet is needed. In that case, you need to connect to the Internet manually using Dial-Up Networking.

According to Microsoft, the **Connect to the Internet As Needed** check box is designed to launch Dial-Up Networking whenever TCP/IP is used by an application, so this behavior is correct.

Changing the Default Server Listening Port (Advanced)

Most users should never have to use this procedure. This procedure should only be performed by a system administrator.

In normal TCP/IP operation, clients use a particular port number to communicate with servers. By default, Pervasive.SQL listens on port 1583. For TCP/IP to function properly, all clients and all servers must use the same port numbers.

Pervasive recommends that you do not modify the server listening port unless you have a conflict. If you do have a conflict, you can override the Pervasive.SQL port settings with an entry in the

SERVICES file for each machine, client and server. The SERVICES file is located at:

Windows NT/2000   \WINNT\SYSTEM32\DRIVERS\ETC\SERVICES

Windows 9X   \WINDOWS\SERVICES

For example, if you wanted to change the listening port to 80 for Pervasive.SQL, you would add the following entry in your SERVICES file:

```
psql 80/tcp
```

Note Entries in the Windows NT SERVICES file are case-sensitive.

After changing your server listening port, you must stop and restart your Pervasive.SQL services for the changes to take effect. See the section "Using Pervasive.SQL" in *Pervasive.SQL User's Guide*.

Using SPX to Connect to a Windows NT/2000 Server

This section documents the use of SPX when connecting to a Pervasive.SQL server running on a Windows NT/2000 machine.

Name Resolution

If you are using Btrieve applications only, then you can use SPX in an all-Windows environment. If you have an ODBC application, however, Pervasive.SQL depends on NetWare for name resolution. You will need at least one NetWare server to perform name resolution for your ODBC applications.

Tasks

Configuring Pervasive.SQL 2000 to use IPX/SPX

IPX/SPX is not a native protocol on the Windows NT/2000 platforms. If you want to use IPX/SPX, perform *all* of the following procedures to ensure proper operation with Pervasive.SQL 2000.

Changing Pervasive's configuration to use IPX/SPX with Windows NT/2000:

➤ **If you have both TCP/IP and IPX/SPX Installed in the Network Icon of the Control Panel, you must remove TCP/IP from the client configuration to make IPX/SPX function with Pervasive applications:**

1 Click **Start**, point to **Programs** and select **Pervasive**.

2 Select **Pervasive Control Center**.

3 Expand **Pervasive.SQL 2000i Engines** and expand the desired engine.

4 Expand **Configuration** and expand **Client**.

5 Select the **Communication protocols** and, in the right-hand pane, right-click on **Supported protocols**.

6 Highlight **Microsoft TCP/IP** from the list of selected protocols and click the < button to move it off the selected list.

7 Click **OK**.

Changing Windows Configuration to Make IPX/SPX Run with Pervasive.SQL:

➤ **Ensure that your IPX/SPX settings are correct**

In the Windows NT Control Panel:

1   Click **Start**, point to **Settings** and select **Control Panel.**

2   Double-click on the **Network** icon.

3   Select **IPX/SPX Compatible Transport** and click **Properties**.

4   Click the **Advanced** tab.

5   Under **Frame Type** setting, ensure that the correct frame type for your network is selected. Do <u>not</u> use **Auto Frame Type Detection**.

6   Still under the **Advanced** tab, click **Network Address**.

7   Enter a non-zero value for your network address. For information about what your network address should be, contact your system administrator.

In the Windows 2000 Control Panel:

1   Click **Start**, point to **Settings** and select **Control Panel.**

2   Double-click on the **Network and Dial-up Connections** icon.

3   Select **Local Area Connection** and click **Properties**.

4   Scroll down to **IPX/SPX/NetBIOS Compatible Transport**, highlight and click the **Properties** button.

5   In the **Frame Type** field, ensure that the correct frame type for your network is selected. Do <u>not</u> use **Auto Detect**.

6   In the **Network number** field, enter a non-zero value for your network address. For information about what your network address should be, contact your system administrator.

➤ **Ensure that your IPX/SPX Maximum Packet size (MaxPktSize) is set correctly in the Windows NT/2000 registry:**

1   Click **Start** and select **Run**.

2   Type regedit and press **Enter.**

3 Find the registry entry using the following paths:
Windows NT:
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Servic
es\NwlnkIPX\NetConfig*driver_name*\MaxPktSize.
Windows 2000:
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Servic
es\NwlnkIPX\Parameters\Adapters*name*\MaxPktSize.

4 Ensure that the **MaxPktSize** setting in the Windows registry is set
to 576 decimal or 240h.

NetBIOS and the Pervasive.SQL Server Engine

Pervasive.SQL does not support NetBIOS using the server engine. Only the Workstation and Workgroup engines support NetBIOS.

Using TCP/IP to Connect to a Linux or Solaris Server

Your Samba must be properly configured on your Linux or Solaris server in order to properly network with Windows-based clients. See Chapter 13, "Network Settings for Linux and Solaris" for more information.

Tasks

Configuring a Client for the Server's IP Address

You must have your server's IP address configured properly on your Pervasive.SQL requesters for Pervasive.SQL to function. You can do this in one or two ways depending on which client requester platform you are using:

➤ **To configure your clients to use the server's IP address, do one of the following:**

For Windows 95/NT clients:

■   Click **Start**, point to **Settings** and select **Control Panel**. Double-click on the **Network icon**, double-click on **TCP/IP** in the list. Click the **DNS Configuration** tab. Enable DNS by selecting the corresponding button and entering the appropriate server information. This is the preferred method for Windows clients.

For Windows 2000 clients:

■   Click **Start**, point to **Settings** and select **Control Panel**. Double-click on the **Network and Dial-up Connections** icon, double-click on **Local Area Connection** and then click the **Properties** button. From the component list, select **TCP/IP** and click **Properties**. Enable DNS by selecting the corresponding button and entering the appropriate server information. This is the preferred method for Windows clients.

or

■   Enter your Server IP address/name in the hosts file `\WINNT\SYSTEM32\DRIVERS\ETC\HOSTS` file (Windows NT/2000).
`\WINDOWS\SYSTEM\DRIVERS\ETC\HOSTS` file (Windows 95).

Preventing the Windows Dial-Up Network Dialog Box from Displaying When Using a Pervasive Application with TCP/IP

The **Windows Dial-Up Networking** dialog box can display when a TCP/IP request is made to Windows. Usually, this is to make an Internet connection, but this feature can be an annoyance when using Pervasive applications and TCP/IP.

➤ To Prevent the Dial-Up Networking Dialog Box from Displaying Automatically:

1   Click **Start**, point to **Settings** and select **Control Panel**.

2   Double-click the **Internet** icon in Control Panel.

3   Click the **Connection** tab.

4   Clear the check box titled **Connect to the Internet As Needed**. In Windows 2000, this button is titled **Dial whenever a network connection is not present**.

Note While this stops the dialog box from displaying with Pervasive applications, this also has the side effect that other applications such as Internet browsers will no longer automatically spawn the **Dial-Up Networking** dialog box when a connection to the Internet is needed. In that case, you need to connect to the Internet manually using Dial-Up Networking.

According to Microsoft, the **Connect to the Internet As Needed** check box is designed to launch Dial-Up Networking whenever TCP/IP is used by an application, so this behavior is correct.

Using the DOS Requesters

See Chapter 10, "Using the DOS Requesters."

Application Configuration Scenarios

9

Common Scenarios for Setting up Your Database Engine

This chapter explains how to configure your database engine for several common application scenarios.

■ "Using Multiple Applications" on page 9-2

■ "Using a Server Engine and a Workstation Engine Concurrently" on page 9-4

■ "Accessing Data on Other Computers" on page 9-6

Using Multiple Applications

Sometimes, two or more applications may use the same database engine. You will need to configure the database engine differently depending on whether the applications are used at the same time.

If your vendors supply configuration guidelines for Server configuration parameters, you will need to adjust your configuration based on these guidelines.

If the applications run concurrently (that is, if two or more applications are using the database engine at the same time) ...

You should configure the engine by adding together all the recommended values for each parameter. For example, if one application vendor suggests **Access | Number of Sessions** should be set to 400, and another application vendor suggests this parameter should be set to 600, then you should set it to 1000.

If the default value is higher than the sum of the recommended settings, then do not change the default value.

Do not add up the recommended values for any buffer size settings, or log file size settings. Use the largest recommended setting. Again, do not change the default if it is larger than any vendor recommendation.

If the applications do not run concurrently (that is, if only one application is running at any given point in time) ...

You should configure the server by using the largest recommended value for each parameter. For example, if one application vendor suggests **Access | Number of Sessions** should be set to 400, and another application vendor suggests this parameter should be set to 600, then you should set it to 600.

If the default value is higher than the largest recommended setting, then do not change the default value.

Settings Affected by Multiple Applications

Most Server settings are not affected when you are running multiple applications. This section explains the settings that may need to be adjusted for multiple applications.

Access | Number of Sessions

All the other settings in the **Access** category are managed dynamically by the server engine. Because each client may open multiple sessions with the database, this setting may need to be

adjusted if you have a hundred or more concurrent clients accessing the server.

You can determine whether you need to increase this setting by using Monitor to inspect the maximum and peak values of Total Remote Sessions. To start Monitor, click the **Start** menu and select **Programs | Pervasive | Pervasive.SQL 2000 | Utilities | Monitor**. From the Monitor menu, select **MicroKernel | Communications**. If the Peak value of **Total Remote Sessions** is at or near the maximum, you should increase the maximum value.

Compatibility | Create File Version

Some applications may require that new files be created with version 6.x file format, while other applications may require version 7.x file format (the default).

These applications can run concurrently only if new files are not created during run time. There is no way to toggle the setting back and forth for each application, unless you wish to do it by hand or write a program to do so using the Distributed Tuning Objects.

If the applications do not create new files during runtime, then this setting is not relevant for multiple applications.

Data Integrity | Transaction Durability

Some applications may require durable transactions, while others may not. If you have two application vendors recommending different values for this parameter, then you should set it to **On**. Generally, having transaction durability turned on does not affect applications that do not use transactions.

Performance Tuning | Cache Allocation Size

If the installed memory allows, the cache should be set to a size equal to the sum of all data files that will be open at the same time.

Using a Server Engine and a Workstation Engine Concurrently

You can use your Workstation/Workgroup engine to access files on a remote file server through a mapped drive on either a Windows NT or NetWare file server.

You can also use the client software provided with your workstation engine to connect to other server engines on a remote machine as described in the following sections.

Using a Workstation Engine to Access Local Files and a Remote Server Engine to Access Remote files

If you want to use your local engine for local file access and a remote server for access to files being serviced by the remote Pervasive server, you must change the settings in your MicroKernel Router. There are two ways to access these settings: through the Pervasive Control Center and through the Registry Editor.

➤ **Accessing the MicroKernel Router Settings Through the Pervasive Control Center**

1   Pull up the Pervasive Control Center (from the **Start** menu, select **Programs,** then **Pervasive,** and **Pervasive Control Center**).

2   Under **Control Center Root**, double-click **Pervasive.SQL 2000i engines**.

3   Double-click the icon representing your computer.

4   Double-click **Configuration**.

5   Double-click **Client** and click on **Access.**

The MicroKernel Router settings appear to the right. Be sure the following settings are set correctly:

```
Use Remote MicroKernel Engine: On
Use Local MicroKernel Engine: On
```

➤ **Accessing the MicroKernel Router Settings Through the Registry Editor**

1   Go to the Registry Editor (from the **Start** menu, select **Run** and type `Regedit`).

2 Go to the `HKEY_LOCAL_MACHINE\SOFTWARE\Pervasive Software\Microkernel Router\Version 7\Settings` directory within the Registry Editor.

3 The MicroKernel Router settings will appear in the right hand frame of the Registry Editor. Be sure the following settings are set to the following:

```
Local: Yes
Requester: Yes
```

If you previously used Scalable SQL and choose to keep it in conjunction with Pervasive.SQL 2000i, then the following settings for your SQL requester will also need to be changed. These settings can be found in the Registry Editor in the `HKEY_LOCAL_MACHINE\SOFTWARE\Pervasive Software\SQL Requester\Version 7.0\Settings` directory.

```
Local Usage: Yes
Remote Usage: Yes
```

Note See *Advanced Operations Guide* for more information on changing settings using the Control Center utility.

Accessing Data on Other Computers

The Workgroup engine provides great flexibility for a variety of small networked environments. The table below explains the most common configurations and where to look for more information. In any of the configurations below, a Workgroup engine must be installed on every computer that is expected to access data.

Table 9-1   Summary of Network Configurations

Configuration	Where to look for more information
Small client/server: Data resides on a single computer where a Workgroup engine is installed.	"Setting Up a Small Client/Server Configuration" on page 11-6
Peer-to-Peer: Data resides on two or more computers where Workgroup engines are installed.	"Setting Up a Peer-to-Peer Configuration" on page 11-8
Gateway: Data resides on a file server where no database engine is installed, or it is not running.	"Setting Up a Gateway Configuration" on page 11-10

Using the DOS Requesters

How to Use the DOS Requesters to Run DOS Applications

This chapter explains how to use the DOS requesters to run Pervasive.SQL-based DOS applications in Windows. The DOS requesters support Btrieve applications only, not ODBC applications.

- "Overview" on page 10-2
- "DOS Box Support on Windows" on page 10-4
- "Verifying the DOS Configuration" on page 10-10

Overview

Pervasive.SQL 2000i supports DOS Btrieve applications in several different ways. The exact method, however, depends on the current configuration and environment of the workstation.

Requesters Available to Access Workgroup Engines

The following choices are available:

■ Win32 DOS Box support: Allows a DOS application to run in a DOS box on a Windows NT/2000 or Windows 9x/ME workstation. This enables direct communication to the Windows 32-bit workstation components rather than to the database engine. This configuration can be used with either a local Pervasive.SQL 2000 workstation or workgroup engine, or a remote Pervasive.SQL 2000 server engine. The TCP/IP or SPX protocol supported for client/server access depends on the configuration of the Windows 32-bit components.

Note DOS Box support must be used for Workstation and Workgroup engines. These engines do not support the DOS requesters described below.

Requesters Available to Access Remote Server Engines

You can use Win32 DOS Box support described previously to connect to remote server engines. In addition, you can use these other requesters:

■ DOS TCP/IP requester (BREQTCP): Allows a DOS application to run in any Windows DOS box or on a DOS workstation, and communicate to a remote Pervasive.SQL 2000 server engine via the TCP/IP protocol.

■ DOS SPX requester (BREQUEST or BREQNT): Allows a DOS application to run in any Windows DOS box or on a DOS workstation, and communicate to a remote Pervasive.SQL 2000 server engine via the SPX protocol.

Note BREQTCP, BREQUEST, and BREQNT are explained in *Getting Started with Pervasive.SQL* (Server edition).

There are different reasons for using these different options, including:

- Pervasive.SQL 2000i Workstation and Workgroup products support only the Win32 DOS Box configuration.
- DOS and Windows 3.1x operating systems require the DOS TCP/IP or SPX requesters when accessing a remote server engine.
- Any given network environment must have only one of the supported protocols available.

Supported Configurations

The table below shows the DOS requester configurations that are supported.

Table 10-1 Supported DOS Requester Configurations

Requester	Workstation/ Workgroup on Windows 95/98/ME	Workstation/ Workgroup on Windows NT/2000	Client on Windows 95/98/ME to Remote Server Engine	Client on Windows NT/2000 to Remote Server Engine
DOS Box	✔	✔	✔	✔
BREQTCP	✘	✘	TCP/IP only	TCP/IP only
BREQNT	✘	✘	SPX only	SPX only
BREQUEST	✘	✘	SPX to NetWare only	SPX to NetWare only

DOS Box Support on Windows

The preferred Requester for Windows is BTRBOX. You can use this Requester even for legacy DOS applications, and for Workstation/ Workgroup applications.

> **Note** Use the BREQUEST/BREQNT/BREQTCP Requesters ONLY if you experience a problem with BTRBOX. These Requesters are supported for access to Server engines only, not to Workstation or Workgroup engines.

Win16 and DOS Workstations

The Requesters for non-Win32 boxes are:

- BREQUEST
- BREQNT
- BREQTCP

The following sections document the steps for running DOS applications on Windows NT/2000 and Windows 9x.

Running DOS applications on Windows NT or 2000

All of the components needed to run DOS applications using BTRBOX on Workgroup engine are installed with your client.

After the Windows client component installation, you have everything you need to run a DOS, Windows 16-bit, or Windows 32-bit application. The default DOS application support installed is the Win32 DOS Box configuration.

Using/Disabling Win32 DOS Box Support

Under Windows NT/2000, the DOS Box Install configures the drivers to be completely transparent. Thus, you are able to immediately open a command prompt and run a DOS Btrieve application. The CONFIG.NT file, located in the WINNT\SYSTEM32 directory, contains the command that enables DOS application support. This file is similar to CONFIG.SYS in DOS. Windows NT/2000 loads it for each DOS session opened. In this file, the Install places the following path to load the Win32 DOS Box driver:

DEVICE = C:\PVSW\BIN\BTRDRVR.SYS

➤ **To disable the Win32 DOS box support to use the DOS TCP/IP or SPX requester instead**

1   Remove the files BTRDRVR.SYS and BTRVDD.SYS from your system. These files are installed by default in the \pvsw\bin directory.

2   Edit the CONFIG.NT file (using Notepad) in the folder WINNT\SYSTEM32. Delete or invalidate the line:

```
DEVICE=C:\PVSW\BIN\BTRDRVR.SYS
```

Note Each configuration (Win32 DOS Box, DOS TCP/IP and DOS SPX Support) is mutually exclusive. You must disable the existing configuration before configuring your machine for a new configuration.

Using DOS TCP/IP Support (BREQTCP)

To use the DOS TCP/IP requester on a Windows NT/2000 workstation, you must first disable the Win32 DOS Box support as described above, and then enable the DOS TCP/IP requester configuration.

Note You cannot use BREQTCP with a Workstation or Workgroup engine. It can only be used to connect to a remote server engine.

➤ **To enable BREQTCP support on Windows NT/2000**

1   Copy the following six files into your %SystemRoot%\SYSTEM directory. %SystemRoot% refers to your Windows directory, typically C:\WINNT. You can determine the value for %SystemRoot% by opening a DOS box and entering the SET command.

```
BREQTCP.EXE
BREQTCP.MSG
JSBDOSWS.EXE
JSBDOSWS.DLL
MSOCKLIB.RC
VSLDOS.INI
```

10-5

Click **Start | Settings | Control Panel**.

2 Double-click **System** and click the **Environment** tab.

3 Add the following environment variable:

```
VSL=%SystemRoot%\SYSTEM
```

Note: There must not be a trailing semicolon in this statement.

Once the DOS TCP/IP requester is setup, use the following steps to run a DOS application:

a. Start a DOS session.

b. Run JSBDOSWS.EXE

c. Load BREQTCP.EXE from the \pvsw\bin directory; this needs to be loaded in each DOS session running a Btrieve application.

d. Run your application.

Note: The commands to load JSBDOSWS and BREQTCP can be put in your autoexec.NT file. You can load JSBDOSWS with the Load High command.

To stop a DOS application, unload the requester components by entering the command BREQTCP /u followed by JSBDOSWS /u in the DOS box before closing it.

Using DOS SPX Support (BREQUEST/BREQNT)

Note You cannot use BREQUEST or BREQNT with a Workstation or Workgroup engine. They can only be used to connect to a remote server engine.

There are two DOS SPX Btrieve requesters:

■ BREQNT.EXE is the "dual mode" requester that can be used by a DOS application to communicate to either a Windows NT or a NetWare server.

■ BREQUEST.EXE is a NetWare-only version of the DOS SPX requester. This exists for historical purposes before Windows NT support was available, and is still provided because it requires smaller memory than BREQNT. If you are accessing

Pervasive.SQL on a NetWare server and if you are running low on DOS conventional memory, you may need to use BREQUEST. If you are accessing a Windows NT server, you must use BREQNT.

In order to use the DOS SPX requester on a Windows NT workstation, you must first disable the Win32 DOS Box support as described above; there are no special steps required to enable the DOS SPX requester configuration.

➤ **To run a DOS application with the SPX protocol to a remote server engine**

1   Start a DOS session.

2   Load BREQNT.EXE or BREQUEST.EXE from the \pvsw\bin directory; this needs to be loaded in each DOS session running a Btrieve application.

3   Run your application.

To stop a DOS application, unload the requester by entering the command BREQUEST /u or BREQTCP /u in the DOS box before closing it.

Running DOS Applications on Windows 9x/ME

All of the components needed to run DOS applications using BTRBOX on Workgroup engine are installed with your client.

After the Windows client component installation, you will have everything you need to run a DOS, Windows 16-bit, or Windows 32-bit application. The default DOS application support installed is the Win32 DOS Box configuration.

This section explains how to use and disable the DOS Box support on Windows 9x/ME.

Note Legacy DOS applications must use the BTRBOX configuration on Windows 9X/ME. In fact, it is the only supported configuration with the Workstation and Workgroup engine products. However, with client/server products, you can run DOS applications with either the BTRBOX configuration, or the legacy 16-bit DOS requesters.

If the DOS application is being run on a DOS or Windows 16-bit operating system, then the legacy 16-bit DOS requesters must be used. BTRBOX works on Windows 32-bit operating systems only.

➤ **To run a DOS application with the Win32 DOS Box configuration**

1   Start BTRBOX95:

Select the BtrBox95 icon to launch from the **Start | Programs | Pervasive | Pervasive.SQL 2000 | Utilities** menu.

Load from the \pvsw\bin directory. A minimized dialog appears indicating active support. You must leave this dialog running.

Closure of this dialog will unload BTRBOX95. You only need to run BTRBOX once—you may have multiple DOS sessions open using a single instance of the driver.

2   Start a DOS session.

3   Load BDOSTUB.EXE from the \pvsw\bin directory; this needs to be loaded in each DOS session running a Btrieve application.

4   Run your application.

Note You must reboot after the initial installation of the Windows client components in order for the Win32 DOS Box support to function.

A device driver (VxD) is registered and will not load until reboot.

The proper order to stop a DOS application is as follows:

1   Stop the DOS application.

2   Close the DOS box.

3   Stop BTRBOX95.

If you do not follow this order, BTRBOX95 or the database engine may not unload properly.

In situations where stopping and restarting all clients is recommended, you should perform all three steps above. Simply stopping and starting your client applications is not sufficient.

➤ **To disable the Win32 DOS box support in order to use the DOS TCP/IP or SPX requester**

Note You cannot use the DOS TCP/IP or SPX requesters with a Workstation or Workgroup engine. They can only be used to connect to a remote server engine.

1 Remove the files BTRBOX95.EXE, BDOSSTUB.EXE, and BTRBOX95.VXD from your system; these are installed by default in the \pvsw\bin directory.

2 Using REGEDIT, remove the following key from your registry:

```
HKEY_LOCAL_MACHINE\
 System\
     Current Control Set\
         Services\
             VXD\
                 BtrBox95
```

Note Each configuration (Win32 DOS Box, DOS TCP/IP and DOS SPX Support) is mutually exclusive. You must disable the existing configuration before configuring your machine for a new configuration.

Caution Editing your registry is an advanced procedure and, if done improperly, could cause you to reinstall your entire computer. If you do not feel comfortable doing this, please obtain the services of a qualified technician. Pervasive cannot accept responsibility for a damaged Registry.

Verifying the DOS Configuration

To verify that the install completed successfully, you need a pure DOS Btrieve application. One of the simplest application ships with the server engines. You will find BUTIL.EXE in the PVSW\Clients\DOS directory. Running this command by itself will show a list of available commands. Do this one time to ensure the copyright information says Butil for DOS, not Windows NT. You will then want to access a Btrieve file. You can use the BUTIL -STAT command to do this. If the command completes successfully, your DOS support is functioning as designed.

Try a command like this:

```
BUTIL -STAT f:\pvsw\samples\sample.btr
```

Here, "f:" is a drive letter mapped to your server and "f:\pvsw\samples\sample.btr" is the path and filename of an existing Btrieve data file. A successful completion will return information about the data file. Otherwise, you will see a status code indicating the problem. Pervasive Technical Support can help you resolve any problems.

Configuring the Workgroup Engine

How to Take Advantage of Major Workgroup Features

This chapter explains how to get the most out of your Workgroup engine.

Overview

This section explains the basic concepts and requirements of Workgroup engines. For additional technical details about the Workgroup engine and the Gateway configuration, please see *Advanced Operations Guide*.

Installation Requirements

Every computer that may be used to access the same data at the same time must have a Workgroup engine installed on it.

Operating System Security

Only NetWare and Windows database server engines can enforce OS level file security based on the privileges assigned to the login user name. The workgroup product does not attempt to do this. In a small office, where workgroup engines are most common, this can be considered a plus because they are usually short on networking experts, and the fewer barriers to successful data access the better.

When to Use Workgroup

There are three main configurations in which you would want to use the Workgroup engine.

See *Advanced Operations Guide* for additional technical details about Workgroup engine.

Small Client/Server Configuration

The first configuration takes place when all the data is located on a single computer with a Workgroup engine installed, and there is limited sharing of data. The Workgroup is licensed for up to three concurrent connections, and additional user count licenses may be purchased and applied. This configuration is roughly equivalent to a small client/server configuration.

Peer-to-Peer Configuration

Another situation when you would want to use the Workgroup engine is when the data is distributed among the workstations. This is called a peer-to-peer topology. This configuration is used when each application typically stores much of its own data on the local hard drive, but periodically needs to access data from other workstations or share its own data with others.

In this configuration, each computer shares its data directory or directories. Any computer that needs access to that data maps one or more drives to the shared data directories. Then the Workgroup

engine on each computer acts as a mini-server engine to read/write all changes to the data files on that machine.

Gateway Configuration

See *Advanced Operations Guide* for additional information about setting up a Gateway configuration and about re-directing locator files.

The third topology requiring the use of the Workgroup engine is when the data is stored on a file server where there is no MicroKernel engine. This can be a UNIX server or other type of network file server that gets backed up regularly, but cannot support a MicroKernel engine. In this situation, the first Workgroup engine that opens files in a directory on the server becomes the Gateway to each file in that directory. The other workstations access the data in a client-server fashion through that Gateway engine.

The Gateway engine for a given directory identifies itself by creating a file named ~PVSW~.LOC in that directory. This file is called a *Gateway locator file* and contains the network name of the computer where the Gateway engine is located. Other Workgroup engines attempting to access this data read the locator file to find the name of the engine they must communicate with in order to access the data.

You can ensure that the same engine always services the files in a given directory by making the locator file read-only. This is called a *static gateway,* also referred to as a *fixed gateway*. See "To Set up a Fixed Gateway" on page 11-11 for more information.

The Gateway engine acts as a server engine as it reads and writes pages to the data files, allowing it to make the most use out of its cache. The Gateway feature is designed so that the ownership of any particular directory can change whenever the current gateway engine has no more client applications with any files open in that directory. When the last file is closed in a directory by a given database engine, the engine releases and deletes the locator file. Then the first engine which opens a file after that in that directory becomes the new gateway.

What is a Gateway Engine?

A *Gateway engine* is a Workgroup engine that acts as the sole point of access to all data files in a particular directory on a remote file server. If several Workgroup engines are accessing the same database at the same time, they do not all open the files simultaneously, nor do they share the files. Rather, the first Workgroup engine to access that database becomes the temporary "owner" of those files, and all other Workgroup engines must access the data by contacting the Gateway

engine. Only the Gateway engine has the files open and reads/writes the files. The other Workgroup engines act as clients, making requests to the Gateway engine acting as a mini-server engine.

> **Note** A Gateway engine only comes into play when no database engine is installed on the machine where the data files are, or when the database engine on that machine is not operating.

Multiple Engine File Sharing

Multiple Engine File Sharing (MEFS) is a technology that was used for over a decade in Btrieve workstation engines to share data files. It had a logistical benefit in that each engine worked independently of the other engines. However, each engine had to coordinate it's own actions with any other engine that might be sharing a data file. To perform this task, each engine had to lock bytes in a lock file using operating system locking.

Pervasive chose not to continue with MEFS technology because of some inherent problems. MEFS required a lot of extra network traffic to coordinate access to data by multiple database engines. It also required changes to be written to disk, over the network, as soon as possible so that other engines could get access to those changes. Also, during concurrent user transactions and system transactions, much more of the file was locked than just the portion changed.

Why replace MEFS—Increased Performance

We have found the performance of the Workgroup engine to be dramatically better than version 6.15 workstation engines using MEFS.

The new workgroup solution uses the well-tested Pervasive client/server technology to allow only one engine to open files at a time. Change operations to files shared by multiple workstations can be over a hundred times faster.

Since only one engine is accessing the data, checkpoints do not have to happen as often to the file, making the best use of the MicroKernel's cache. Disk I/O is greatly reduced, and so is the chance for data corruption.

Our benchmark tests using the TPC-B standard show that, for two or more clients, the Workgroup client-server architecture performs anywhere from 20x to 100x faster than the old standard of MEFS.

The TPC-B is a change-intensive test which tracks how fast can we modify the database.

Why replace MEFS—Increased Database Reliability

With MEFS, the chance for lost database changes was greater since database pages needed to be written to the file soon after they were changed so that other engines could gain access to the changes. This situation meant that a lot more data pages were travelling back and forth from the engine to the data file across the network, which greatly increased the chance for corruption.

Setting Up a Small Client/Server Configuration

As explained in "Small Client/Server Configuration" on page 11-2, you should use this set up when you have only a few workstations sharing data located on a central computer where you have a Workgroup engine installed.

If you have data located on many computers, or if you do not or cannot install a database engine on the computer where the data is located, you should use one of the other configurations.

➤ **To Set Up a Small Client/Server Configuration**

1   You must have the Workgroup engine installed on the computer where the data is, and you must have the Workgroup engine installed on all computers expected to access the data.

Be sure that the Workgroup engine is in the folder **Start | Programs | Startup** on the computer where the data is located. You want to be sure that this engine is operational each time the computer is started, before any other database engines attempt to access the data. By placing the Workgroup engine icon in the Startup folder, you ensure that the database engine is loaded when the computer starts up.

Note You may inadvertently fall into a Gateway configuration if the database engine on the machine where the data is located is not started up when the computer is started up. If another Workgroup engine attempts to access the data and the local database engine is not running, the other database engine may establish itself as the Gateway for those data files.

You can resolve this situation by shutting down the computer where the data is located, and starting it up again, while ensuring that no other computers request data before the local Workgroup engine is restarted. You may need to remove the file ~PVSW~.LOC from the data directory to ensure the Gateway is not re-established.

The best way to ensure that only the Gateway services the file is to set a static gateway locator file using the **Gateway Locator Utility**.

2   Share the directory where the data is located so that other computers can map a drive to the data directory.

3   On each Workstation expected to access the data, create an Engine DSN using a drive mapped to the data directory.

4   Set up is complete. The Workgroup engine on the machine where the data is located now acts as a mini-server, to fulfill all requests for data on that machine.

Setting Up a Peer-to-Peer Configuration

As explained in "Peer-to-Peer Configuration" on page 11-2, you should use this set up when you have workstations sharing local data as well as data located on many different machines, and each machine has the Workgroup engine installed.

This configuration is similar to the small client/server configuration discussed above, except that now every Workgroup engine is sharing data as a server.

If you have data located on only one computers, or if you do not or cannot install a database engine on the computer where the data is located, you should use one of the other configurations.

➤ To Set Up a Peer-to-Peer Configuration

1 You must have the Workgroup engine installed on each computer where data is located, and you must have the Workgroup engine installed on all computers expected to access the data.

Be sure that the Workgroup engine is in the folder **Start | Programs | Startup** on each computer where data is located. You must be sure that this engine is operational each time the computer is started, before any other database engines attempt to access the data. By placing the Workgroup engine icon in the Startup folder, you ensure that the database engine is loaded when the computer starts up.

Note You may inadvertently fall into a Gateway configuration if the database engine on a machine where data is located is not started up when the computer is started up. If another Workgroup engine attempts to access the data and the local database engine is not running, the other database engine may establish itself as the Gateway for those data files.

You can resolve this situation by shutting down the computer where the data is located, and starting it up again, while ensuring that no other computers request data before the local Workgroup engine is restarted. You may need to remove the file ~PVSW~.LOC from the

data directory to ensure the Gateway is not re-established.

The best way to ensure that only the Gateway services the file is to set a static gateway locator file using the **Gateway Locator Utility**.

2   On each computer where data is located, share the directory where the data is located so that other computers can map a drive to the data directory.

3   On each Workstation expected to access the data, create an Engine DSN using a drive mapped to the data directory.

Also, create an Engine DSN for any local data that each Workgroup engine needs to access from its own hard drive.

4   Set up is complete. The Workgroup engine on each machine where data is located now acts as a mini-server, to fulfill all requests for data on that machine.

Each Workgroup engine also handles any local data access, that is, database requests from applications on that machine for data that resides on the same machine.

Setting Up a Gateway Configuration

See *Advanced Operations Guide* for additional information about setting up a Gateway configuration and about re-directing locator files.

As explained in "Gateway Configuration" on page 11-3, you should use this set up only when you have data files on a computer where no database engine is installed.

If you have database engines installed on all machines, you should use one of the other configurations.

Note You may inadvertently fall into a Gateway configuration if the database engine on a machine where data is located is not started up when the computer is started up. If another Workgroup engine attempts to access the data and the local database engine is not running, the other database engine may establish itself as the Gateway for those data files.

You can resolve this situation by specifying a permanent Gateway as described in this section, or by shutting down the computer where the data is located, and starting it up again, while ensuring that no other computers request data before the local Workgroup engine is restarted. You may need to remove the file ~PVSW~.LOC from the data directory to ensure the Gateway is not re-established.

The best way to ensure that only the Gateway services the file is to set a static gateway locator file using the **Gateway Locator Utility**.

Floating or Fixed Gateway

You can set up two different Gateway configurations. The default behavior is a *floating* Gateway configuration. In this configuration, the first engine to open the remote data files becomes the Gateway engine for that directory until all files in the directory are closed. Then the next engine to open the data files becomes the new Gateway. This configuration is the most flexible, but also can entail delays upon initial connection to the database, as the engine tries the different network protocols and checks for an existing Gateway engine.

Note Using a floating Gateway in a peer-to-peer configuration with multiple shared data sources is not recommended. This configuration is supported and it operates as designed, however, with multiple engines shuffling ownership among multiple data locations, connection delays may be significant. It is also possible to create a

situation where a Workgroup engine on a remote machine serves as a the Gateway for data located on your local hard drive. Obviously there is no reason to endure this delay when your local Workgroup engine can serve this data with higher performance.

You can avoid this situation by ensuring that the Workgroup engine on every computer is started up when the computer is started up. You must also ensure that someone logs on to each computer, because normally the Workgroup engine doesn't start up until a user logs on.

You can also avoid this situation by permanently assigning each machine as the Gateway for the data files located on it. See "To Set up a Fixed Gateway" on page 11-11 for information on how to perform this task.

The second configuration is called a *fixed* or *permanent* Gateway configuration. In this configuration, a specific engine is permanently assigned as the Gateway engine for a specific directory. If that engine is not running when another engine attempts to access the data, an error code results and the data is not available.

➤ To Set up a Floating Gateway

1 This is the default behavior. All you need to do is set up an Engine DSN on each Workgroup computer. When creating the Engine DSN for the database, use a mapped drive or a UNC pathname to identify the remote data directory.

The Gateway assignment now floats dynamically as different Workgroup engines access the remote data.

➤ To Set up a Fixed Gateway

To specify a permanent Gateway engine for a given directory, you need only change the attributes of the ~PVSW~.LOC file to read-only, once it contains the name of the desired engine. There are several ways to perform this task.

Use the Gateway Locator Utility

1 From the **Start** menu, choose **Programs | Pervasive | Pervasive.SQL 2000 | Utilities | Gateway Locator**.

2 In the **Target Directory** area, type in or browse for the directory containing the data files for which you wish to set up a permanent Gateway.

3 In the **Directory Status** area, click **Change**. In the dialog box that appears, click **Assign a Gateway**, then type in or browse for the network name of the computer that you want to be the Gateway. Click **OK**.

4 Back in the main Gateway Locator window, check **Permanent assignment**. Click Exit.

Use the DOS Command Line

1 Use the ATTRIB +R command at a DOS command prompt to change the attributes of the ~PVSW~.LOC file.

For example, if your current directory is the directory where the file is located, you can type the following command:

```
ATTRIB +R ~PVSW~.LOC
```

Use the Windows Explorer

1 Right-click on the ~PVSW~.LOC file in the directory for which you want to make a permanent Gateway assignment. Choose **Properties** from the pop-up menu.

2 In the Properties window, on the General tab, click **Read-only** in the section labeled Attributes. Click **OK**.

Working with the Gateway Locator Utility

The Gateway Locator Utility provides control of and insight into any Gateway configuration you have on your network. This section explains how to use the utility for a variety of purposes.

This utility enables users to determine or change the Workgroup Engine which is being used as the gateway for the data files in a particular directory. The Gateway Locator utility is used only with Pervasive.SQL 2000i Workgroup Engine.

The Gateway Locator operates by reading and manipulating the locator file, ~PVSW~.LOC, which resides in any directory which is assigned a Gateway engine. If this file is locked (in use), the Gateway Locator can only locate, not change, the Workgroup engine being used as a Gateway for that particular directory.

To start the Gateway Locator, From the **Start** menu, choose **Programs | Pervasive | Pervasive.SQL 2000 | Utilities | Gateway Locator**.. The following dialog box will appear:

Figure 11-1 Gateway Locator Main Dialog Box
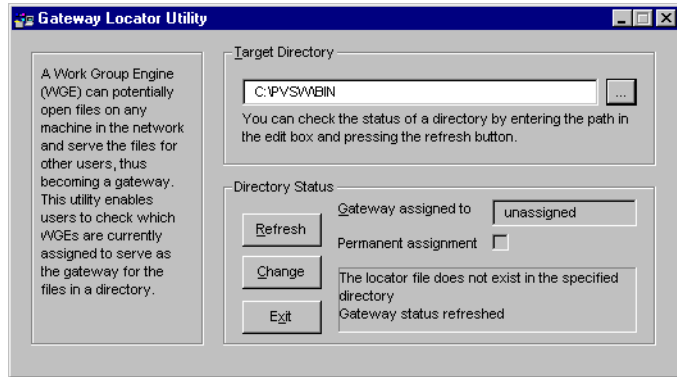
In the **Target Directory** box, enter the path of the directory which contains the data files for which you wish to locate or change the Gateway engine. The default target directory is the current working directory. Clicking the "..." button allows you to browse for the target directory, by bringing up the following dialog box:

Figure 11-2 Gateway Locator Browse Dialog Box
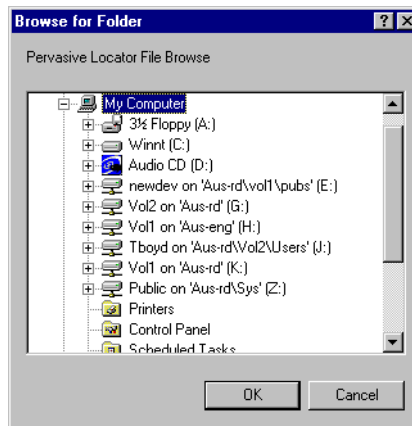
➤ Locating the Gateway Workgroup Engine

Once the target directory is selected, clicking the **Refresh** button causes the name of the Gateway engine for that directory (if such exists) to appear in the "Gateway Assigned To" box. If no Gateway exists for a particular directory, the box reads "unassigned".

➤ **Changing the Gateway Workgroup Engine**

Once the target directory is selected, clicking the **Change** button allows you to choose the Workgroup engine which you wish to serve as Gateway for a particular directory (this button is disabled if the locator file for that directory is locked.) The following dialog appears:

Figure 11-3 Gateway Assignment Dialog Box
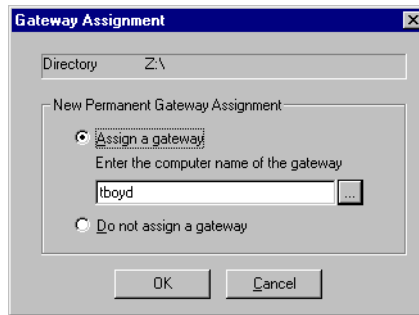
You may enter the name of the machine you wish to serve as gateway, or click on the "..." to browse for it. The following dialog box appears:

Figure 11-4 Browse for Computer Dialog Box
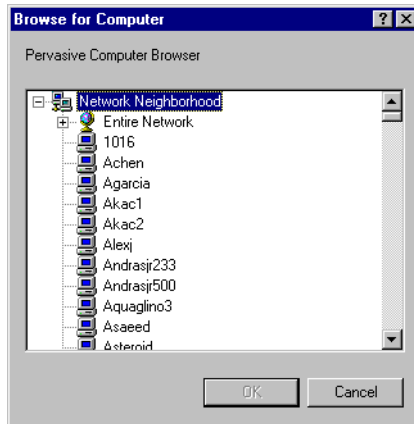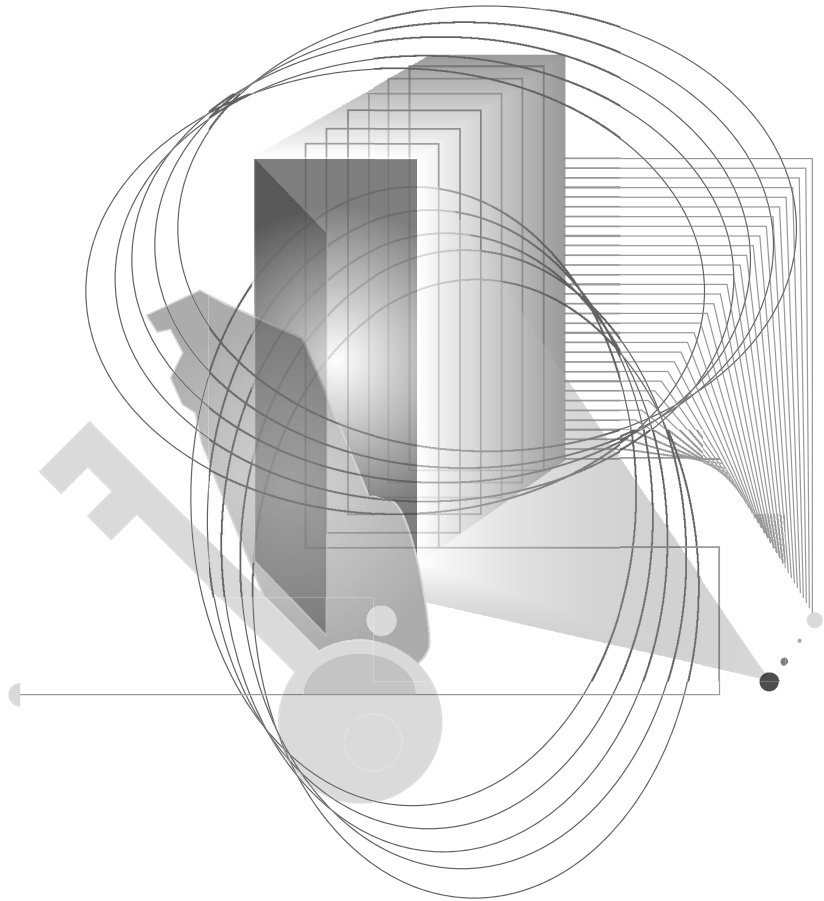
AFTER INSTALLATION

Connecting Clients to a Pervasive.SQL Database

How to Connect Clients to the Sample Database on a Pervasive.SQL Workstation or Workgroup Machine

A sample database called DEMODATA is included on every Pervasive.SQL 2000i installation. This chapter shows how to use a Pervasive.SQL 2000i client to:

❏ Connect to a machine that is running a Pervasive.SQL engine.

❏ Use the Pervasive Control Center to find the DEMODATA database.

❏ Access information in DEMODATA tables using Pervasive Control Center.

This chapter contains the following sections:

Basics About Pervasive.SQL Servers and Clients

You should note the following about Pervasive.SQL servers and clients.

Every Engine is Also a Client

- The client components of Pervasive.SQL 2000i are installed with every server install. So if you have a Pervasive.SQL 2000i Workstation engine installation, you can use your machine connect to other servers as a client.

Clients

- You can use Pervasive Control Center to connect to remote machines on which a server engine of Pervasive.SQL 2000i is installed.
- You will need to login as an administrative user on the remote server to perform most functions.

Connect a Client to a Pervasive.SQL Server

1 Start PCC

Click **Start** | **Programs** | **Pervasive** | **Pervasive Control Center**

2 Register a Server Engine in your Namespace.

Right-click on **Pervasive.SQL 2000 Engines** (displayed in your namespace) and select **Register New Engine**.

The Namespace is the column on the left side of the Control Center window that contains a list of machines to which you are connected.

Note The machines listed in your namespace will remain between Control Center sessions. To remove a machine, right-click on the machine name and select **Delete**.

Figure 12-1 Registering a New Engine
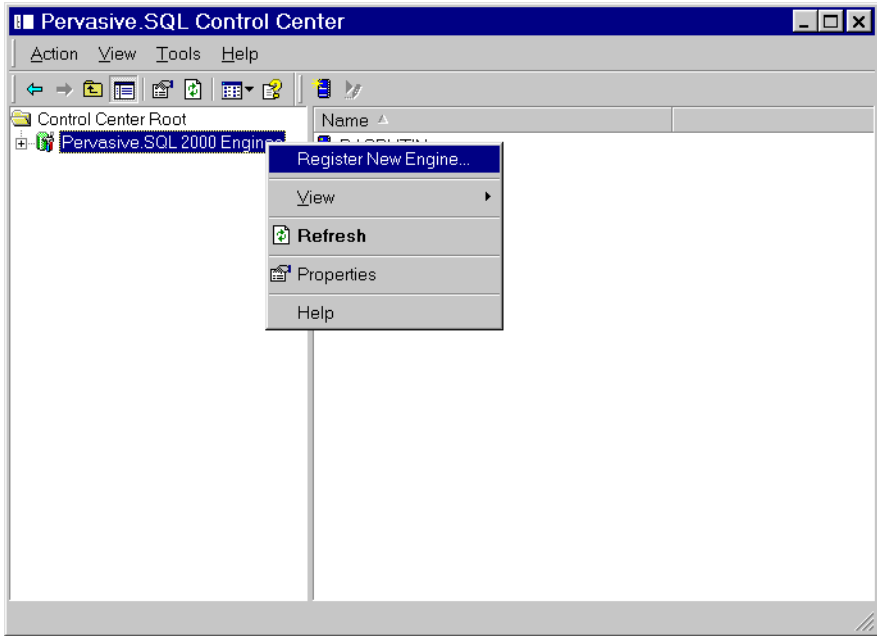
A dialog is displayed that allows you to choose the machine name where the Pervasive.SQL 2000i server resides.

Figure 12-2 Choosing a Computer Name
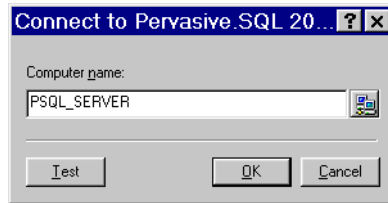
3 If you are unsure of the computer name, click the button to the right of the name field and select a computer name from the Network list that appears.

Figure 12-3 Choosing a Computer Name
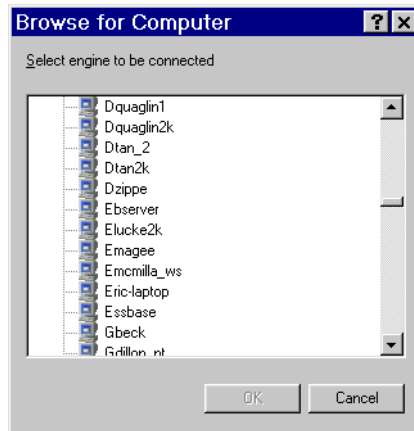
4 Click **OK** to add the machine to your namespace, or click **Test** to test the connection first.

5 You need to be authenticated on the remote server, and a dialog displays prompting you for a user name and password. Enter the values in the appropriate field and click OK.

6 You are now connected to the remote Pervasive.SQL server.

Use the Client to Obtain Data from the Server

Once you are connected to the database:

1 Expand the Databases list for your server as shown in Figure 12-4 on page 12-5.

Figure 12-4 Expanding the Databases List for a Machine
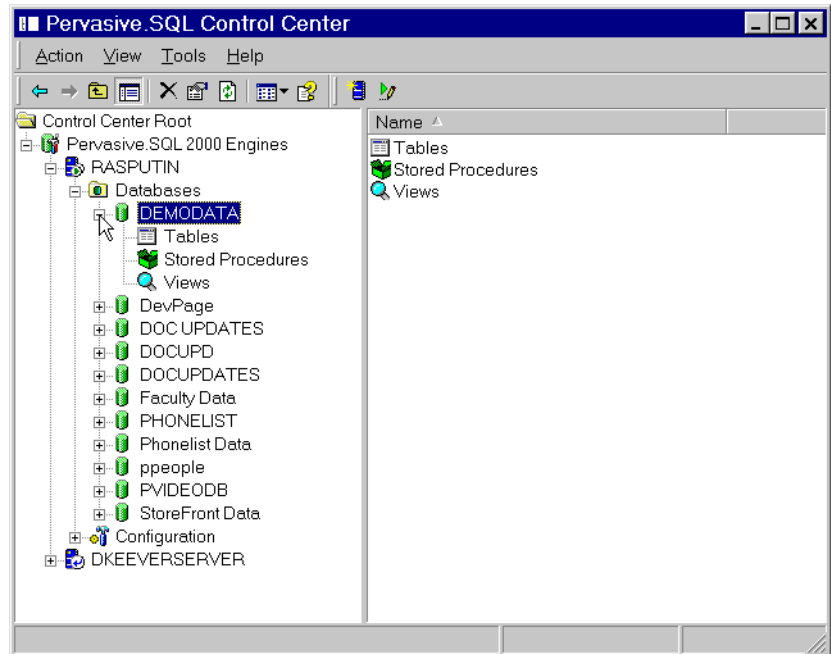
2 Find DEMODATA in the list and click on the plus (+) sign.

3 Select Tables from the list.

4 Click on the table DEPARTMENT as shown in Figure 12-5 on page 12-6.

Figure 12-5 Selecting the DEPARTMENT Table in DEMODATA
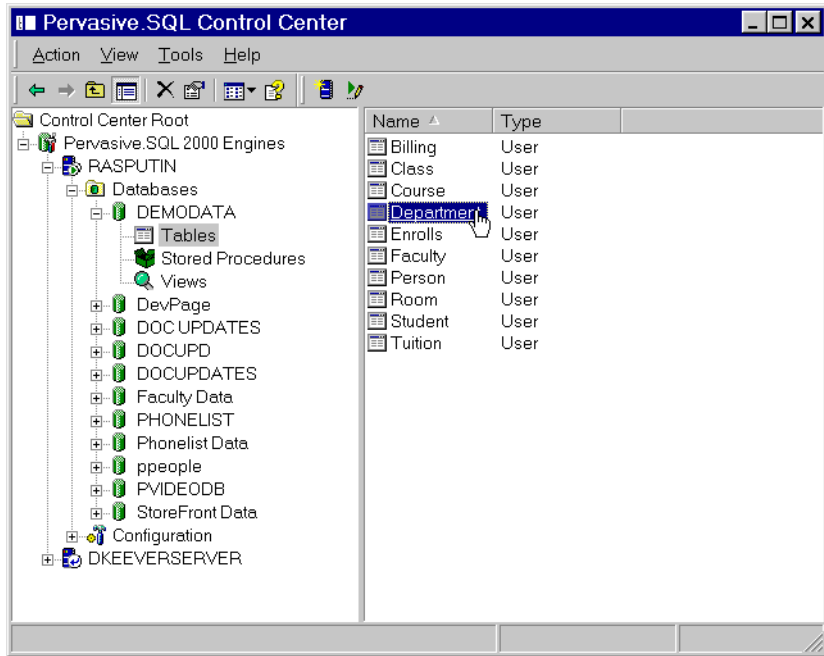
5   By default, a "SELECT * FROM" query is run and the table results are displayed in an active grid as shown in Figure 12-6 on page 12-7

The active grid that loads when your data is displayed is updateable. That is, changes you make to the data in that grid are stored to the database.

If you wish to browse the data without the risk of accidentally updating data, click the Text Results tab at the lower-left of the SQL Data Manager window, as shown in Figure 12-6.

Figure 12-6 Displaying the DEPARTMENT Table in DEMODATA
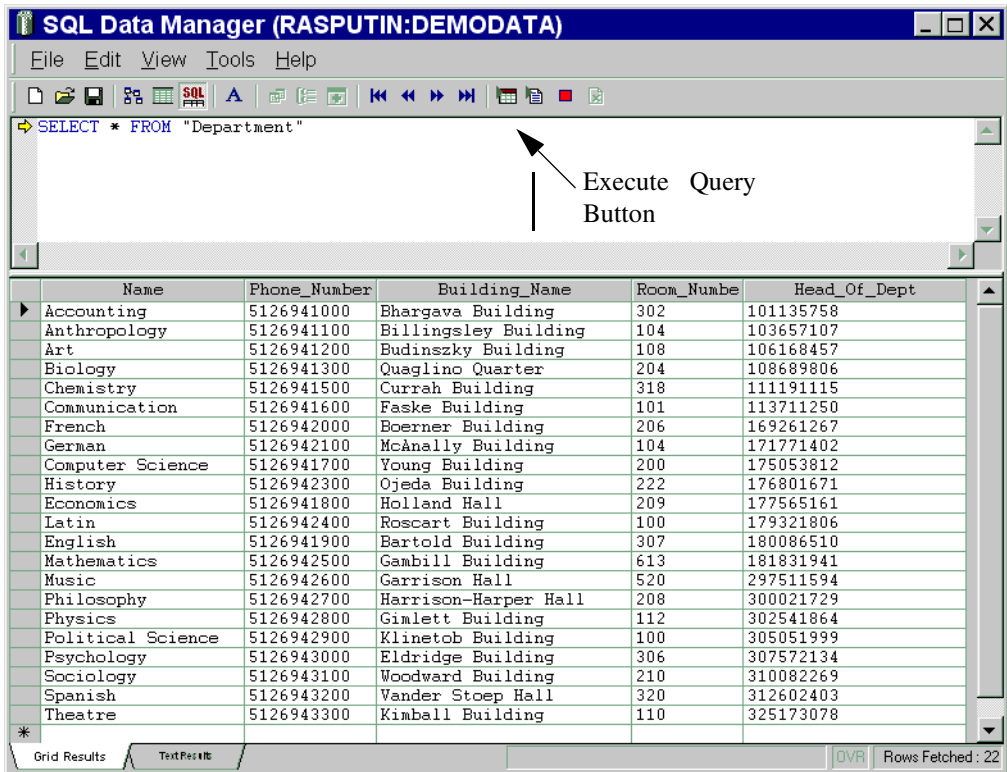
6   Refine your query

To show how to modify a basic query, restrict the results to only departments that start with the letter 'M' by altering the query at the top half of the screen and clicking the Execute Query button, which is shown in Figure 12-6.

```
SELECT * FROM Department WHERE Name LIKE 'M%'
```

You can see the results of the query in Figure 12-7 on page 12-8.

Figure 12-7 Refining Your Query - DEPARTMENT Table in DEMODATA
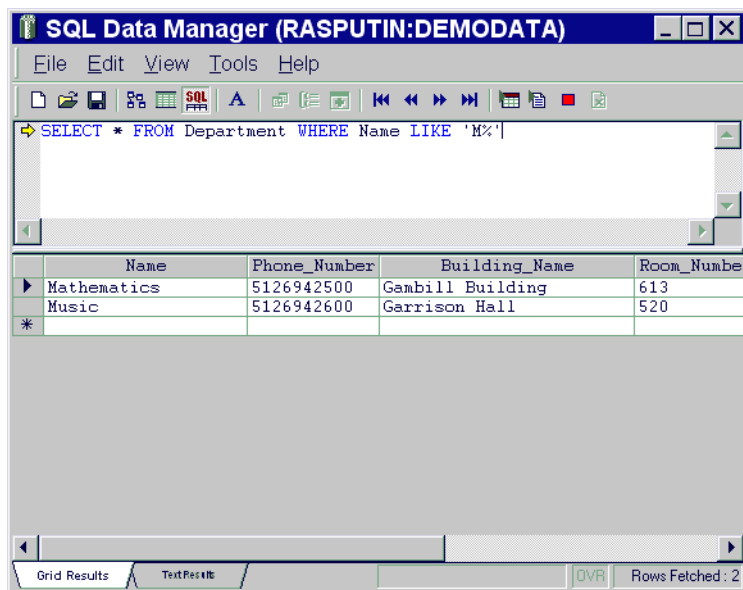
Become an Expert User

Now that you can do some basic database access using Pervasive Control Center, you may want to continue learning about other features.

Users's Guide

This manual contains an overview of the Pervasive Control Center and instructions for performing basic tasks.

Advanced Operations Guide

This manual contains advanced operating and maintenance tasks, including database operations.

Troubleshooting After Installation

How to Proceed When You Encounter Errors During Installation

Pervasive Software provides several features and tools in Pervasive.SQL 2000i that help to prevent configuration and installation problems.

Some of these utilities are installed and run as part of the installation process and all can be run later to evaluate configuration and registry settings and to troubleshoot problems. They are shown in Table 13-1 on page 13-2.

This chapter contains the following sections:

- "Troubleshooting Tools" on page 13-2
- "Troubleshooting Strategies" on page 13-3
- "Diagnosing Problems with Pervasive System Analyzer (PSA)" on page 13-4
- "Checking Network Connectivity Using PSA" on page 13-5
- "Verifying Database Engine is Running" on page 13-8
- "Checking Client Functionality" on page 13-9
- "Obtaining File, Client, and Engine Version Number" on page 13-12
- "Troubleshooting Common Pervasive.SQL 2000i Issues" on page 13-17
- "How to Get Additional Help" on page 13-22

Troubleshooting Tools

The following table describes some tools that can help you avoid or solve problems.

Table 13-1 Pervasive Tools that Assist in Installation and Problem Determination

Feature/ Component	Function	For More Information
Pervasive System Analyzer	Analyzes system components, runs communication tests, and archives or restores previous database engine files on your system.	See "Diagnosing Problems with Pervasive System Analyzer (PSA)" on page 13-4.
Gateway Locator	A utility which can be used to determine or change the Gateway being used for a particular data dictionary (only in Pervasive.SQL 2000i Workgroup Edition.)	See "Configuring the Workgroup Engine" on page 11-1.
Smart Components	Smart Components is an internal design in Pervasive.SQL that ensures that Pervasive software components always load with compatible components.	Refer to the *Pervasive.SQL Advanced Operations Guide*.
Knowledge Base	Provides information about many Pervasive software configurations and common environments.	Search the Pervasive Knowledge base at: http://support.pervasive.com/kb/

Troubleshooting Strategies

Pervasive Software hopes that your installation process completes without experiencing any problems. However, this depends on a number of factors, including proper network support, and operating system configuration.

If something does go wrong during an installation, Pervasive offers some tools that can help in diagnosing the problem. This chapter explores some of the troubleshooting techniques that you can use.

Checklist for Problems

❏ Did you see any error messages displayed during installation?

❏ Does the Network function correctly?

❏ Is the Engine Running?

❏ Is the Client Software correctly functioning?

❏ Are there errors in your PVSW.log file?

Troubleshoot the Problem

The rest of this section contains procedures that you can use in verifying your installation.

- "Checking Network Connectivity Using PSA" on page 13-5
- "Verifying Database Engine is Running" on page 13-8
- "Checking Client Functionality" on page 13-9
- "Obtaining File, Client, and Engine Version Number" on page 13-12
- "How to Get Additional Help" on page 13-22

Diagnosing Problems with Pervasive System Analyzer (PSA)

Pervasive.SQL 2000i introduces you to a new diagnostic tool: PSA.

Pervasive System Analyzer (PSA) is conveniently integrated into the product installation and available as a stand-alone diagnostic tool to help you with a the following tasks:

- Troubleshoot network problems
- Detect previous installations of Btrieve or Pervasive.SQL on your system
- Note other factors that influence your networking environment
- Perform archives and/or restorations of previously-installed versions of Pervasive.SQL

PSA replaces the features that were previously offered by SmartScout and InstallScout.

How to Start PSA

➤ **To run PSA, follow these steps:**

1 Start PSA by doing one of the following:

- Go to:
  Start | Programs | Pervasive | Pervasive.SQL 2000i | Utilities | Pervasive System Analyzer.
- Double-click the PSA icon in Windows Explorer (95/NT) or File Manager (Windows 3.*x*).
- Enter `psawizrd.exe` (Win32) from the command line in a Win32 DOS box (Windows 9.*x*, Windows NT/2000 only). This assumes that the Pervasive installation directory is in your path.
- Select **Run** from the **Start** menu (all Windows) or from the **File** menu (Windows 3.*x*), and enter `psawizrd.exe` (Win32) and press **Enter**. This assumes that the Pervasive installation directory is in your path.

SmartScout Replaced by PSA

Note SmartScout was replaced by Pervasive System Analyzer starting in Pervasive.SQL 2000i Service Pack 3. The documentation on SmartScout has been removed.

Checking Network Connectivity Using PSA

If you suspect that network connectivity may be the source of your problem, the Pervasive System Analyzer Utility can help you check this.

➤ **To check Network Connectivity using Pervasive System Analyzer:**

1   Load Pervasive System Analyzer.

Click **Start** | **Programs** | **Pervasive** | **Pervasive.SQL 2000i** | **Utilities** | **Pervasive System Analyzer**

2   Click **Next** after the Welcome screen displays

3   Choose the **Network Communications** test as shown in Figure 13-1. Before starting the test by clicking **Next**, choose the path where you want the log file written. The log file will contain information that may help you in diagnosing the problem.

Figure 13-1 Select Network Test
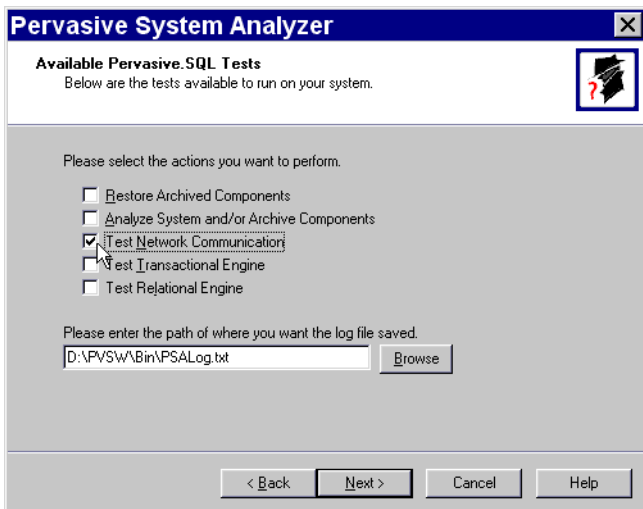
4   Enter the path to your network server in the dialog that follows as shown in Figure 13-2.

Figure 13-2 Select Network Path
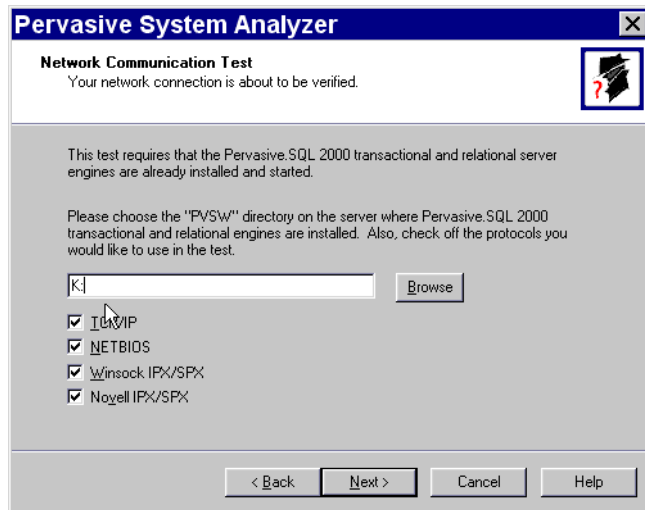
You can enter a path based on a mapped drive, or a UNC path. On Unix platforms, specify the Samba share for Pervasive.

- Windows-based platforms: \\my_server\pvsw
- Unix-based platforms: \\my_server\psqldata (psqldata is default Samba share).

Click **Next** to start the test.

5 While the test is running, you can see its progress as shown in Figure 13-3 on page 13-7.

Figure 13-3 Network Test In Progress
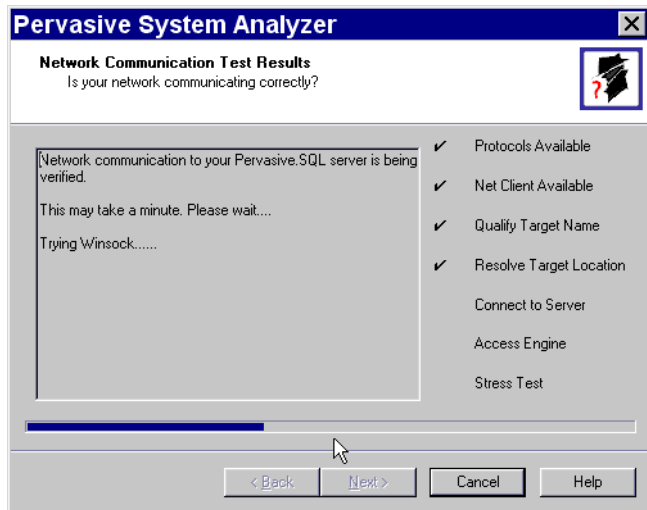
6   When the test is complete, this screen will show the results of the test. Click **Next** to display the finish page.

Figure 13-4 Choosing to Display the Log File
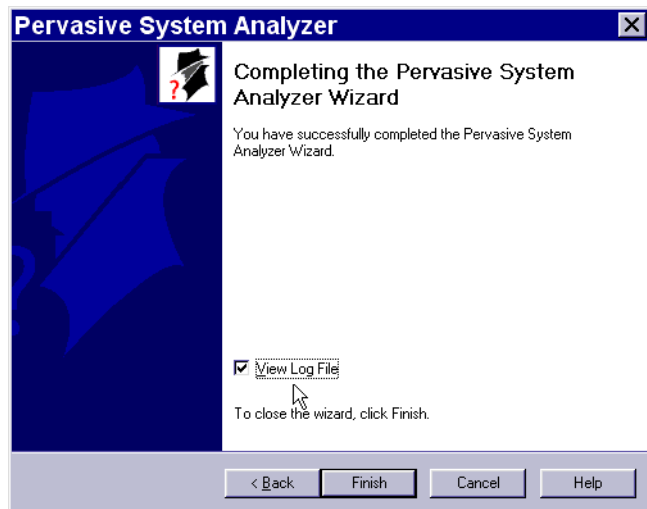
7   Select the check box to view the log, as shown in Figure 13-4. The log file contains information that may help you diagnose the problem. The log file remains in the directory you chose earlier (see Figure 13-1 on page 13-5).

Verifying Database Engine is Running

To verify that the Pervasive.SQL 2000i server engine is running, see the procedure for your platform:

➤ To Start the Pervasive Workstation engine:

1 Click **Start**, arrow to **Programs**, to **Pervasive**, then **Pervasive.SQL 2000i**, and **Engines**.

2 Select the **Pervasive.SQL Workstation Engine** icon.

By default, the MicroKernel allocates resources and is ready to service local application database requests.

➤ To Stop the Pervasive Workstation engine:

1 On the Windows taskbar, right-click the **Pervasive Database** icon.

2 Select **Stop Engines and Exit**.

Note You will receive a warning message when trying to stop the engine if any of the following is true:

- There are active clients.
- No activity took place since the engine loaded.
- 10 seconds has not elapsed since the last operation took place.

Note Starting the Pervasive.SQL Workgroup Engine is exactly the same as starting the Workstation Engine, except that the engine icon is labeled "**Pervasive.SQL Workgroup Engine**".

Checking Client Functionality

If you want to test the client components as the possible source of problems, use the Pervasive System Analyzer utility.

➤ **To check Client Software Functionality using Pervasive System Analyzer:**

1 Load Pervasive System Analyzer.

Click **Start** | **Programs** | **Pervasive** | **Pervasive.SQL 2000i** | **Utilities** | **Pervasive System Analyzer**

2 Click **Next** after the Welcome screen displays.

3 Choose the **Test Transactional Engine** and **Test Relational Engine** tests as shown in Figure 13-5. Before starting the tests by clicking **Next,** choose the path where you want the log file written. The log file will contain information that may help you in diagnosing the problem.

Figure 13-5 Select Transactional and Relational Engine Tests
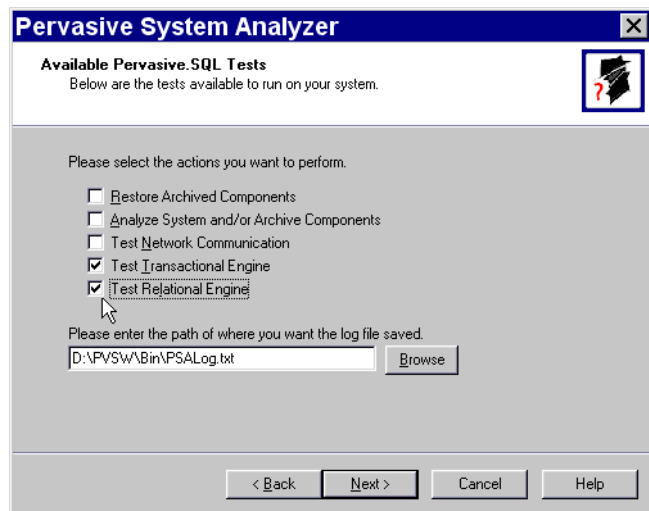
4 Click Next to set up the Transactional test. The resulting dialog is shown in Figure 13-6.

Figure 13-6 Transactional Test Setup Dialog
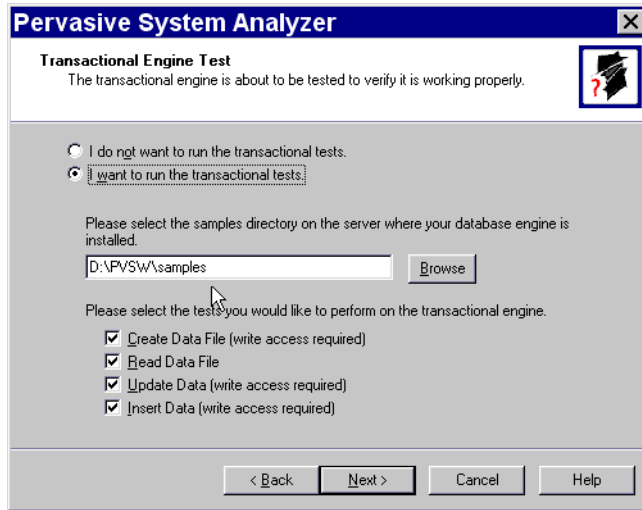
5 The dialog should already be configured to the location of the SAMPLES directory of your Pervasive installation. If not, enter the path of your SAMPLES directory in the appropriate field.

To target a Pervasive.SQL server on a Unix platform, include the Samba path share in the field. For example, `\\unix_server\psqldata\samples`.

Click **Next** to start the test.

Figure 13-7 Transactional Test Results Dialog
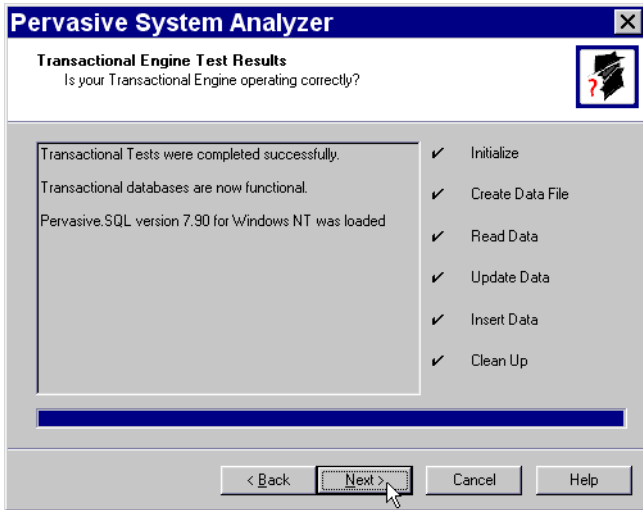
6   After the Transactional test completes, the results are displayed
on the screen, noting any errors that occurred. Click **Next** to set
up the Relational test.

Figure 13-8 Relational Test Set Up Dialog
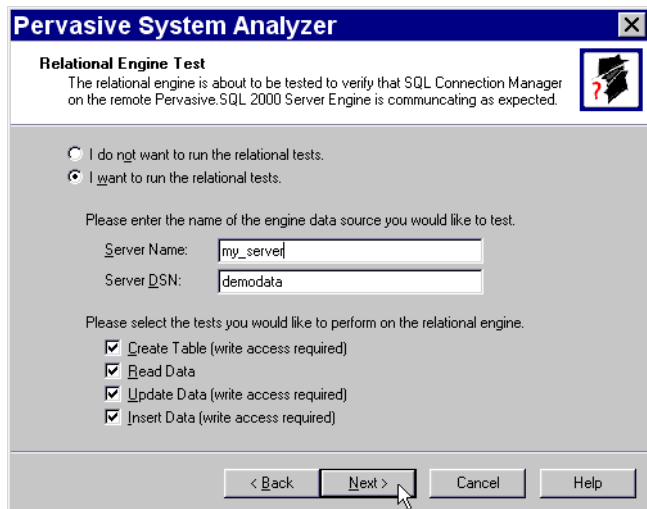
7   In the Server Name field, enter the name of the machine to
which you want to run the Relational test if it is not the local
machine. Click Next to run the test. A dialog similar Figure 13-7
displays when the Relational test is completed.

Obtaining File, Client, and Engine Version Number

You can use Pervasive.SQL utilities to verify that the client and engines have the version number you expect, or to check the version of a particular file.

Determining Client and Engine Version

You can check the engine and client versions using Function Executor on Windows platforms or using the BUTIL command-line utility on all platforms:

Using Function Executor

1   Load the Function Executor Utility.

Click **Start | Pervasive | Pervasive.SQL 2000 | Utilities | Function Executor**

Function Executor is a utility that simulates Btrieve client operations using the Pervasive.SQL requesters.

2   Select the **Btrieve Version Info** button, which appears as a clock icon. The correct button is shown in Figure 13-9.

Figure 13-9   Selecting the Btrieve Version button
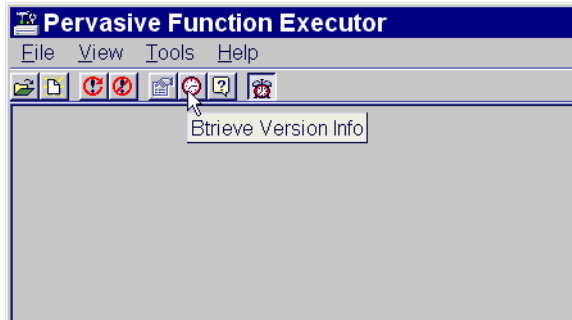
3   After choosing the **Btrieve Version Info** button, a dialog displays that indicates the version of the client requesters and the engine.

Figure 13-10 Btrieve Version Info Display
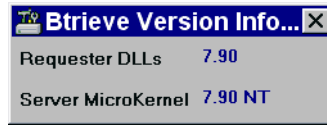
Using the BUTIL Utility

From a command prompt, enter the following:

```
BUTIL -VER
```

The requester and engine versions are then displayed.

Determining a File Version

You can determine the file version of a MicroKernel data file using the Pervasive Control Center or Function Executor on Windows platforms, or using the BUTIL command-line utility on any platform:

- Using the Pervasive Control Center to select the properties of a table.
- Using Function Executor to open and then query the version of a file.
- Using BUTIL command-line utility to query the statistics of the file.

Using the Pervasive Control Center

You can use the Pervasive Control Center to perform this task.

➤ To Determine the File Version of a Table Using Pervasive Control Center

1   Click **Start** | **Programs** | **Pervasive** | **Pervasive Control Center**

2   Find the database by expanding its name in the Control Center namespace on the left.

3   Right-click on a table name and select Properties as shown in Figure 13-11.

Figure 13-11 Obtaining a File Version with the Pervasive Control Center
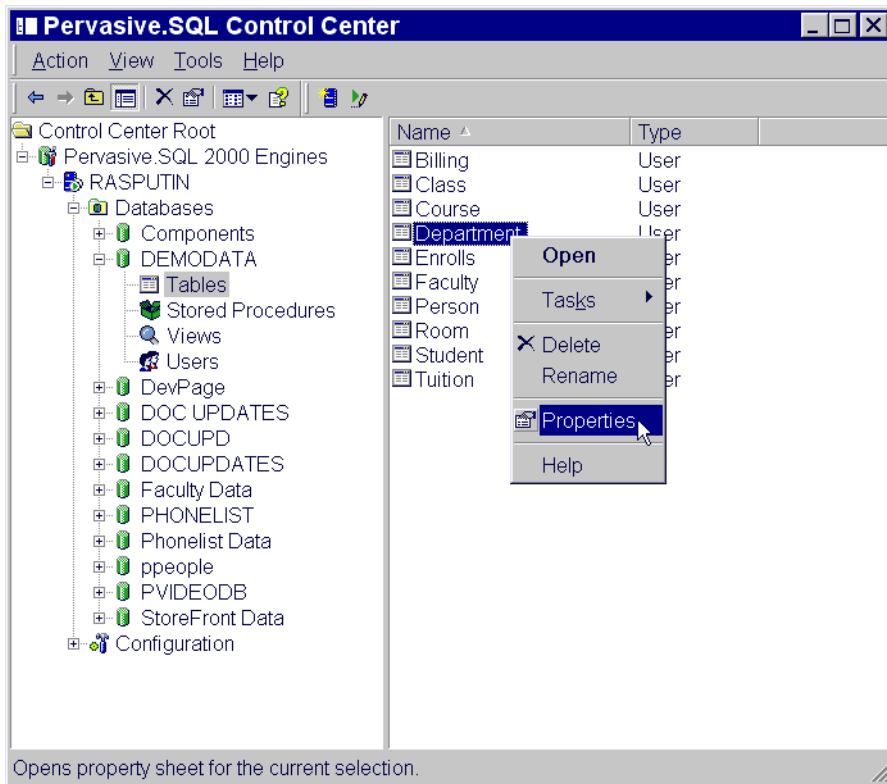
4   The table properties are displayed, which includes the file
version of the underlying MicroKernel data file.

Figure 13-12 Table Properties Page
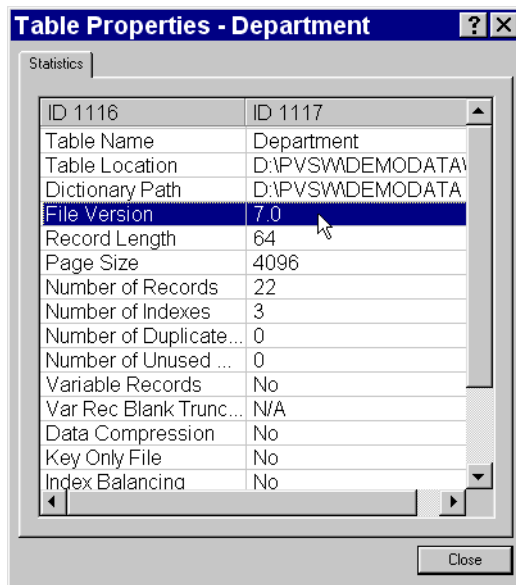
Using Function Executor

The Function Executor utility can simulate Btrieve operations and can be used to determine the file version by performing the following operations:

1 Open (0)

2 Version (26)

The Function Executor utility is documented in more detail in *Pervasive.SQL Advanced Operations Guide*.

Using BUTIL command-line utility

Use the -stat parameter of BUTIL to query the file statistics, which includes information about:

- File version
- Pages
- Records
- Keys

Type the following at a command prompt:

```
butil -stat <filename>
```

For example, to query the statistics of the file DEPT.MKD of the DEMODATA database included with Pervasive.SQL:

```
butil -stat dept.mkd
```

The BUTIL utility (available on Windows, NetWare, Linux, and Solaris) is documented in more detail in *Pervasive.SQL Advanced Operations Guide*.

Troubleshooting Common Pervasive.SQL 2000i Issues

This section outlines problems you may encounter during the installation or when first using the Workstation product.

I fail to see the effects of my configuration changes.

Try stopping and then restarting the MicroKernel. Whenever you make a change to engine configuration components, you must stop and restart the MicroKernel engine for the changes to take effect. For information on how to start and stop the MicroKernel engine, see "Verifying Database Engine is Running" on page 13-8.

Why do I (now) receive Status 95, after running my application successfully?

Your application has lost its session with the MicroKernel. This can happen if you make changes to your configuration settings and must restart the MicroKernel, as in the troubleshooting example given above. At the moment the MicroKernel is stopped, any application that is running loses its session with the MicroKernel. You must stop all those utilities and restart them in order to reestablish communication.

See the *Status Codes and Messages* manual for more cases in which this status code can be returned.

I'm having problems connecting to a NetWare server with a 16-bit client.

You may be missing two Novell system DLLs on your system. The Win16 client requires these files to connect using SPX.

1   Check for the existence of NWCALLS.DLL and NWIPXSPX.DLL.

2   If you do not have them on your system, you can download the patch files from Novell's web site (http://www.novell.com) and copy them to your system directory.

Installing a Pervasive.SQL application has rendered another application unusable.

If the latest DLLs have been overwritten, it is possible to restore the overwritten DLLs using a backup directory that is automatically created when you install Pervasive.SQL 2000i.

Pervasive.SQL uses a system called Smart Components to avoid component mismatches and the occurrence of an older program overwriting newer components. To provide an interface for legacy applications, however, some component names are still capable of being overwritten.

➤ **To restore the compatibility DLLs:**

1 Open the directory
C:\WINDOWS\SYSTEM\BTCOMPAT.BCK.

2 Copy the files located inside this directory to your Btrieve requester's system directory. For example, C:\WINDOWS\SYSTEM for a Windows 95 machine or C:\WINNT\SYSTEM32 for a Windows NT/2000 machine.

Note You only need to perform this procedure if your Pervasive.SQL engine was made inoperable by installing an incompatible or outdated application.

How do I verify that my DOS components are functioning properly?

Pervasive provides a DOS version of BUTIL.EXE in C:\PVSW\BIN\DOS for purposes of verifying that your DOS components are functioning properly.

➤ **To verify using the DOS version of BUTIL:**

1 Enter the following at a DOS prompt:

```
cd \pvsw\bin\dos
butil -stat ..\..\samples\sample.btr
```

In addition, use of the BDOSSTUB.EXE is needed for some applications using Windows 95 DOS Box for Btrieve if you encounter a Status 20, indicating the requester is not loaded. Some applications built with Turbo Pascal have this problem, which can be

avoided by loading the BDOSSTUB.EXE TSR before starting the DOS application.

Why can't I restart my application after an improper program exit?

Database engine components may remain in memory if the engine is interrupted improperly.

➤ **If you cannot restart your program after improperly aborting the application by using Ctrl-C or stopping the process:**

1 Shut down and restart your system.

2 Avoid terminating applications in an abnormal manner.

Why isn't my application using the Workstation engine?

If you previously installed Pervasive.SQL requesters and later installed the Pervasive.SQL 2000i Workstation engine but your application is only using the requesters, you may have an outdated configuration that sets Local Access to Off. The Pervasive.SQL 2000i Workstation engine's installation does not overwrite existing settings. To reset Local Access to On, see "Using a Workstation Engine to Access Local Files and a Remote Server Engine to Access Remote files" on page 9-4.

I have previously installed requesters and then Pervasive.SQL 2000i. How do I resolve a Status 3012 returned while running a 16-bit application?

A 16-bit Pervasive.SQL application that functions properly going to a remote server will require an adjustment when running against the Workstation engine. The Pervasive.SQL Workstation engine requires that 16-bit applications use the ODBC thunk mechanism.

Thunking with the Microsoft-provided ODBC thunk DLLs is required for remote Pervasive.SQL engine access from a Win16 ODBC application running on a Win32 workstation.

You may also receive a Status 802 from the Workstation engine in this scenario.

Status Code 3012 (Local engine is not accessible to the MicroKernel router) can be corrected by examining your Thunk settings.

➤ **How to set up Win16 Applications to use 32-bit Components (Thunking):**

1   On the Start menu, go to:
Programs | Pervasive | Pervasive Control Center.

2   Double-click **Pervasive.SQL 2000 Engines** and double-click the icon representing your client computer.

3   Double-click **Configuration** and double-click **Client-16 bit**.

4   Click on **Access**, and in the right-hand pane, right-click **Use Thunk** and make sure it is set to **On**.

5   In the right-hand pane, right-click on **Use Remote MicroKernel Engine** and ensure that it is set to **Off**.

6   Click on **Use Local MicroKernel Engine** and ensure that it is set to Off.

Because the Remote and Local parameters are both set to Off, the application passes control to the "Thunking Layer." After this procedure is completed, your Win16 applications use Win32 components to access the database engine.

I do not see the Configuration Utility when I open the PCC. How do I display this utility?

➤ **To display the Configuration Utility from inside the PCC:**

1   Ensure that you are connected to the proper server.

2   Double-click **Pervasive.SQL 2000 servers**.

3   Double-click the desired database engine icon.

4   **Configuration** is one of the options under the desired server name.

How Do I Access the Pervasive.SQL 2000i Online Manuals?

➤ **If you wish to read the online documentation:**

1   On the Start menu, choose:
Programs | Pervasive | Pervasive.SQL 2000 | Documentation | Pervasive.SQL Documentation.

2 Click on the desired manual under the **Contents** tab or search for a topic.

Note Pervasive.SQL 2000i Documentation is in the form of HTML Help or WinHelp files for Win32 platforms. Documentation for Win16 platforms is in the form of Adobe Acrobat files. Shortcuts to these files are found in the Pervasive.SQL Documentation file group. The Acrobat Reader program is available on the Pervasive.SQL 2000 CD-ROM. Alternately, the application may be downloaded at http://www.adobe.com.

How to Get Additional Help

Pervasive Software strives to ensure that your product installation is easy and successful. If you encounter problems during the installation that are not covered in this manual, please contact Pervasive Software in one of the following ways and we will address your problem promptly.

For general questions, common problem resolution, client/server issues, your first line of support should be the Pervasive Software Knowledge Base, a web-based searchable index of all Pervasive technical information. This is located at http://support.pervasive.com/kb.

For developer-related issues, visit the following web sites:

■   Pervasive Developer Center
    http://www.pervasive.com/developerzone
■   PervasiveDevWire: Free developer-oriented e-mail news service at
    http://www.pervasive.com/developerzone/emailnews

If your installation is not successful, or you encounter problems not documented in the Pervasive.SQL manuals or on the Knowledge Base, contact Pervasive Software Customer Support in one of the following ways:

■   Open an electronic incident at http://www.pervasive.com/support/Email_Support.taf
■   Fax or telephone the Support Center in your region. See *Pervasive.SQL User's Guide* for a complete list of technical support contacts.

Thirty-Day Free Technical Support

If you still have questions or problems relating to your Pervasive.SQL 2000i installation, you can obtain help from the Pervasive Customer Support department.

Your purchase of Pervasive products entitles you to 30 days of free technical support for installation and configuration problems.

See the following section, "Pervasive.SQL Resources and Contacts" on page 14-1 for information on how to contact Pervasive Software Customer Support.

Pervasive.SQL Resources and Contacts

A Guide to Pervasive.SQL Customer Information Resources

Pervasive Software strives to ensure that your experience with Pervasive.SQL 2000i is successful. This chapter describes the resources and information available to you as a valued customer of Pervasive Software.

The following variety of resources can help you get answers to your questions, troubleshoot problems, and interact with the Pervasive team as well as with other customers:

Printed Documentation

Pervasive.SQL 2000i SP3 comes with a printed copy of *Getting Started* and *Status Codes Quick Reference*. A complete suite of online documentation is installed on Windows when you choose the **Typical** installation procedure. It is available as an option in the **Custom** installation procedure. The content is accessible through the **Start** menu:
Programs | Pervasive | Pervasive.SQL 2000i | Documentation | Pervasive.SQL 2000 Documentation.

Printed versions of the following titles are available for purchase:

- *Getting Started with Pervasive.SQL* (Server or Workstation/ Workgroup edition)
- *Pervasive.SQL User's Guide*
- *SQL Engine Reference*
- *Status Codes and Messages*
- *Pervasive Products and Services*
- *Advanced Operations Guide*
- *What's New in Pervasive.SQL 2000i*

To order manuals, please contact Pervasive Software using one of the following methods:

Online: http://www.pervasive.com/products/manuals/

E-mail: salessupport@pervasive.com

Phone: 1-800-287-4383

Developer Center

The Pervasive Software Web site is a great source for Pervasive.SQL information: http://www.pervasive.com. It is your most immediate source for assistance with the product.

The link shown below is commonly referred to as Developer Center. It is a great starting point from which to navigate to available downloads, documentation, product updates, news articles, sample code and tutorials. Developer Center also provides access to an expansive technical library and training information.

http://www.pervasive.com/developerzone/

Pervasive.SQL Knowledge Base

The Pervasive.SQL Knowledge Base Online is a searchable database for technical information regarding installation, configuration, component management, product defect status, and answers to the frequently asked questions (FAQs). The Knowledge Base, shown below, uses an associative problem-solving technology to perform contextual searches and can be used to quickly find specific answers to your questions about Pervasive products.

http://support.pervasive.com/eSupport/
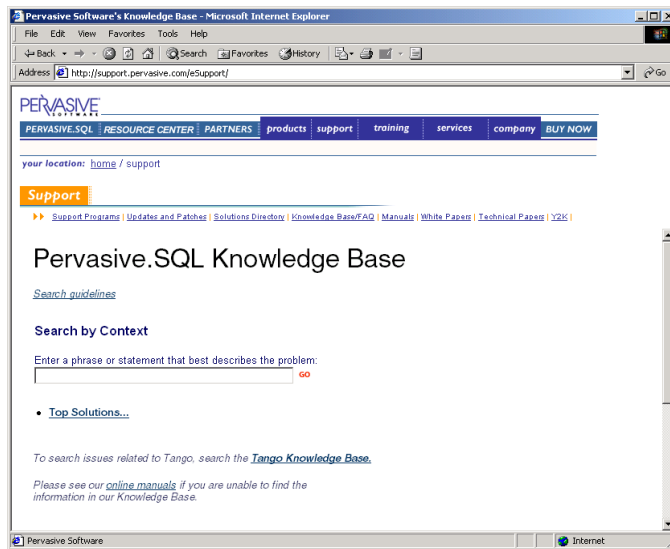
FTP Site

Pervasive Software strives to maintain close ties to developers using Pervasive.SQL for their database applications. On the Pervasive FTP site, you can find practical resources such as downloadable updates and patches to our product offerings as well as additional debugging tools, documentation, third-party tools, and beta releases.

ftp://ftp.pervasive.com/support/

Online Documentation

The latest versions of Pervasive.SQL product manuals in CHM, HLP, and PDF format are available for download from the Pervasive Software web site. These titles include:

- *Getting Started with Pervasive.SQL* (Server or Workstation/ Workgroup edition)
- *What's New in Pervasive.SQL 2000i*
- *Pervasive.SQL User's Guide*
- *SQL Engine Reference*
- *Status Codes and Messages*
- *Pervasive Products and Services*
- *Advanced Operations Guide*

http://www.pervasive.com/support/technical/product/

DevWire

DevWire is a monthly newsletter on all things Pervasive: learn about Beta cycles and releases, Service Pack releases, current topics, FYIs, Q & As, Pervasive Software Events in your area, trade shows where you can find us, and much more!

To subscribe to DevWire, send an e-mail message with "**add**" as the subject line to devwire@pervasive.com.

DevTalk

Pervasive Software's DevTalk discussion forums are a great way to share ideas with other customers, get technical questions answered, and give feedback directly to Pervasive Software. Check out the figure below to see some of the topics currently being discussed and watch the web site for new topics as they appear!

http://www.pervasive.com/devtalk/

Newsgroup

Many Pervasive.SQL customers enjoy participation in a newsgroup—a learning environment in which users help users, with some participation by Pervasive Software. The newsgroup is managed by the end-user community, posting and answering questions as they wish.

Pervasive Software is represented in the worldwide network of news discussion groups at:

news://comp.databases.btrieve

E-Mail

Pervasive Software welcomes your comments, suggestions and requests for assistance via e-mail. Please submit to the following contacts:

- docs@pervasive.com

  For comments or concerns regarding the content of Pervasive.SQL documentation. Requests for manuals should be directed to your sales rep or "salessupport@pervasive.com".

- techsupport@pervasive.com

  For technical support for Pervasive.SQL. For a faster response, we recommend submitting your request via the e-mail support form located at: http://www.pervasive.com/support/Email_Support.taf

- salessupport@pervasive.com

  For information about Pervasive.SQL sales matters such as contacts, pricing, and product specifications. You may also submit the sales information form located at: http://www.pervasive.com/contact/email_sales.tml

- developer@pervasive.com

  For developer relations. A great way for developers to communicate their ideas about all Pervasive products, interfaces and programs.

- beta@pervasive.com

  For questions regarding beta products and general beta information, including Pervasive's Beta Program.

- info@pervasive.com

  For general information about the company, marketing efforts, public relations, and other general questions.

- investor.relations@pervasive.com

  For questions by investors.

Technical Support

Pervasive Software Technical Support Call Centers

- Support Center (Headquarters-Austin)

   Hours: 7 A.M. to 7 P.M. Central Standard Time.

   Phone: 512.231.6000

   Toll Free: 800.287.4383

- European Service and Support Center (Belgium)

   Hours: 9 A.M. to 6 P.M. Central Europe Time.

   Phone: +32 2 710-1660

For detailed contact information for Pervasive Software offices worldwide, refer to *Pervasive Products and Services.*

To receive a faster response to your technical support questions, we recommend using the electronic support method located at: http://www.pervasive.com/support/Email_Support.taf

Index

Symbols