IBM

Implementing PeopleSoft 8.4 on zSeries

- Overview of PeopleSoft on zSeries
- Installation of PeopleSoft 8.4
- Configuration of PeopleSoft 8.4

Dawn Fallon
Alain Roy
Marshall Murdock

Redbooks

ibm.com/redbooks

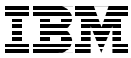

International Technical Support Organization

Implementing PeopleSoft 8.4 on zSeries

December 2002

SG24-6549-00

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (December 2002)

This edition applies to PeopleSoft Version 8.4 using DB2 UDB 7.1 on UNIX and Windows NT/2000.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	Home Director™	QMF™
AS/400®	IBM®	RACF®
Balance®	IBM eServer™	Redbooks™
CT™	Language Environment®	Redbooks (logo)™ 
Database 2™	MVS™	S/390®
DB2®	MVS/ESA™	SP™
DB2 Connect™	Net.Data®	WebSphere®
DB2 Extenders™	OS/390®	xSeries™
DB2 Universal Database™	PAL®	z/Architecture™
DRDA®	Parallel Sysplex®	z/OS™
e Strategy™	Perform™	zSeries™

The following terms are trademarks of International Business Machines Corporation and Lotus Development Corporation in the United States, other countries, or both:

Lotus®	Notes®	Word Pro®

The following terms are trademarks of other companies:

WebLogic, Tuxedo, and Jolt are trademarks of BEA Systems Inc. in the United States, other countries, or both.

Crystal Reports is a trademark of Crystal Decisions, Inc. in the United States, other countries, or both.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Preface

This IBM Redbook will help you install PeopleSoft 8.4 on zSeries with DB2. This information is based on installation experience gained while installing PeopleSoft 8.4 at the IBM Silicon Valley Laboratory in San Jose, California and on customer support experiences.

This book will be especially useful for those who are installing and implementing PeopleSoft 8.4 DB2 on zSeries for the first time. Basic knowledge of DB2, Windows NT/2000, and SQL is assumed

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Dawn Fallon is an IBM PeopleSoft zSeries Consulting I/T Specialist with the ERP Solutions Sales Organization, Americas. Her areas of expertise include PeopleSoft, ERP and zSeries. Dawn has seven years of experience with ERP applications with the last four years exclusively focused on PeopleSoft, as well as 22 years overall I/T experience.

Alain Roy is an IT Architect in the Products and Solution Support Center (PSSC) in Montpellier, France. He has 12 years of experience in the Project and Application Management field. He holds an MBA degree in Marketing from Saint-Etienne Business School. His areas of expertise include technical support and project assistance on benchmarks, proof of concepts for the IBM EMEA sales team in PeopleSoft-related projects and opportunities. He has written extensively about PeopleSoft.

Marshall Murdock is a Technical Consultant in the International Technical Support Organization, in San Jose, California. He has 25 years of experience in Project Management, e-business, e-commerce and CRM. He holds degrees in Law and Marketing. His areas of expertise include WebSphere, e-business, CRM, and CRM products. He has written extensively about CRM and WebSphere.

Thanks to the following people for their contributions to this project:

- Michael Curtis, IBM International Competency Center, San Francisco
- Lee Sigmund, PeopleSoft US
- Randall Hicks, PeopleSoft US
- Steven Sereg, PeopleSoft Canada
- Joe DeCarlo, ITSO San Jose
- Mike Ebbers, ITSO Poughkeepsie

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. 1WLB Building 80
  650 Harry Road
  San Jose, NY 12601-5400

1

Introduction

This IBM Redbook will help you install PeopleSoft V8.4 using DB2 UDB V7.1 on zSeries and Windows NT/2000 platforms. It is based on the installation experience gained while installing PeopleSoft V8.4 at the IBM Silicon Valley Laboratory in San Jose, California, and on customer support experiences.

In this chapter, we provide an overview of the complete installation and implementation of PeopleSoft 8.4 on zSeries and DB2 applications, and of the products involved in managing a PeopleSoft environment. We explore implementation in more detail in subsequent chapters.

1.1 Overview of PeopleSoft

PeopleSoft is a major provider of e-business applications that build a collaborative network of customers, suppliers and employees. PeopleSoft helps organizations optimize every interaction with their customers, employees, and suppliers to create more loyal, collaborative, and profitable relationships.

PeopleSoft Inc. was founded in 1987 by David Duffield and Ken Morris. The company started out as a human resource client/server package application vendor and developed and delivered its first client/server human resource management system (HRMS) at the end of 1988. In April 1992, the company began branching out beyond its original human resource stronghold by shipping its second major product line, PeopleSoft Financials. In November 1992, the company completed its initial public offering on NASDAQ (PSFT). In December 1994, it announced PeopleSoft Distribution, its third major horizontal offering. In 1996, PeopleSoft acquired RedPepper Software and launched PeopleSoft Manufacturing.

In 1999, PeopleSoft's product range expanded to include PeopleSoft Enterprise Performance Management (EPM), PeopleSoft eStore and PeopleSoft eProcurement. It acquired Vantive, the number two Customer Relationship Management (CRM) vendor. PeopleSoft has embedded analytic capabilities into its new generation of eBusiness products. This means it delivers “smart” eBusiness applications such as Customer Profitability, Workforce Analytics, Supply Chain Analytics, and Balanced Scorecard. PeopleSoft’s pure Internet Customer Relationship Management, Supply Chain Management and Enterprise Management solutions provide the industry’s most open and flexible eBusiness platform.

In 2000, the latest version of PeopleSoft was delivered. PeopleSoft 8 represents a generational shift to the Internet in enterprise software architecture. Each PeopleSoft 8 application is delivered completely in HTML, with no code on the client. This pure Internet architecture means that the right data is available to the right person in the enterprise anytime, anywhere from any Web browser. Customers are able to reduce the costs of their IT infrastructure, network administration, and information access.

PeopleSoft’s Enterprise Solution is based on PeopleTools, a core set of tools that provides the technology foundation for all PeopleSoft applications. These tools enable customers to implement, tailor, and maintain PeopleSoft applications as well as to extract, analyze, and manipulate data. PeopleTools includes several tools for reporting, customizing, and workflow. Based on a multi-tier client/server

architecture and using advanced workflow technology, PeopleSoft products support clients running both Microsoft Windows and popular Web browsers, as well as a range of mainframe, midrange, and LAN relational database server platforms.

More than 4,700 organizations in 107 countries run on PeopleSoft eBusiness applications. PeopleSoft now employs more than 8,000 people worldwide.

Details of the PeopleSoft applications which are currently available can be found in Appendix A, “PeopleSoft Products” on page 361.

More information on PeopleSoft products and services can be found on the World Wide Web at:

http://www.peoplesoft.com

1.2 Overview of zSeries

The zSeries is well positioned as the deployment platform for e-business data and transaction workloads. It offers traditional enterprise computing strengths - scalability, high availability, low total cost of ownership and robust security - necessary to create the kind of flexible computing infrastructure required for enterprise-wide e-business solutions.

The zSeries is based on the entirely new z/Architecture, which can eliminate bottlenecks associated with the lack of addressable memory. This is achieved through its virtually unlimited addressing capability for unpredictable workloads and growing enterprise applications.

The zSeries servers are supported by a multitude of operating systems and z/OS is the foundation for future of the zSeries, an integral part of the z/Architecture designed and developed to quickly respond to the demanding quality of service requirements for enterprise e-business. z/OS is a new, robust operating system that is based on the new 64-bit z/Architecture. It delivers the highest qualities of service for enterprise transactions and data, and extends these qualities to new applications using the latest software technologies. It provides a highly secure, scalable, high-performance base on which to build and deploy Internet and Java enabled applications, providing a comprehensive and diverse application execution environment.

1.2.1 Scalability

zSeries is the first mainframe class server designed to meet the high performance transaction requirements for the next generation of e-business. zSeries offers highly scalable servers and proven technology that lowers costs, reduces complexity, and speeds deployment of e-business infrastructure. It has the ability to consolidate and manage hundreds of virtual servers on a single mainframe, and presents uncompromising flexibility in selecting, building and deploying e-Business applications.

1.2.2 Parallelism

The Parallel Sysplex cluster is an integral part of the zSeries platform and is the foundation on which a growing number of subsystems and operating system enhancements are based. Parallel Sysplex technology is a highly advanced processing clustered system. It supports high-performance, multisystem, read/write data sharing, enabling the aggregate capacity of multiple z/OS and OS/390 systems to be applied against a common workloads. This, in turn, facilitates dynamic workload balancing, helping to maximize overall systems throughput and providing consistent application response times.

Further, through data sharing and dynamic workload balancing, continuous availability and continuous operations characteristics are significantly improved for the clustering system, as servers can be dynamically removed or added to the cluster in a nondisruptive manner. If the processing demands grow and exceed the capacity of the existing server systems, it is possible to add an additional system to the Parallel Sysplex cluster and grow the application workload transparently. This can be accomplished across multiple servers.

Parallel Sysplex clustering provides capacity you eventually cannot outgrow; continuous computing and single-system management is a very open, flexible package. It provides real business value by allowing you to integrate your enterprise and focus on your strategic goals.

1.2.3 Availability

In addition to providing more utilities, DB2 for OS/390 and z/OS V7.1 improve availability with the following enhancements.

With version 7.1 you can change the value of many subsystem parameters without stopping DB2.

Data sharing customers can benefit from a new restart option. Restart lets you restart DB2 with a smaller storage footprint so that you can quickly recover retained locks following an abnormal termination.

1.3 Overview of DB2 7.1

The history of DB2 stretches back almost to the first days of relational databases. The relational database concept was originally developed by Dr. E.F Codd, an IBM researcher. In June 1970, Codd published an article entitled *A relational model of data for large shared data banks*, which triggered a flurry of database research, including a major research project called System /R within IBM. Work on System /R began in the mid-70s at IBM's Santa Teresa laboratories in San Jose, California, and led to two major developments: the development of the Structured Query Language (SQL), and the production of various commercial relational database management system (RDBMS) products.

In 1983, IBM introduced an RDBMS for its operating systems called Database 2 (DB2). DB2 operated under IBM's MVS operating system. The first release of DB2 began shipping in 1985, and soon became IBM's flagship RDBMS. DB2's SQL language became the de facto standard database language. DB2 technology migrated across all IBM product lines, from personal computers to network servers to mainframes. In 1997 IBM took the DB2 cross-platform strategy further by announcing DB2 versions for computer systems made by Sun Microsystems, Hewlett Packard and other IBM hardware competitors.

Now the DB2 Universal Database, at the heart of the DB2 product family, runs on a multitude of platforms.

1.3.1 Parallelism

Several enhancements deliver more efficient processing of SQL queries and better performance for complex queries. The way DB2 for OS/390 and z/OS Version 7.1 process queries can improve performance of your SQL queries. Among the improvements are the following items:

- Fewer sort operations for queries that have an ORDER BY clause and WHERE clauses with predicates of the form COL= constant.
- IN-list index access parallelism, which can improve performance for queries involving IN-list index access.
- The ability to use indexes with a correlated subquery so that ERM applications, such as PeopleSoft applications, will perform and scale much better.

1.3.2 High availability and data sharing

The DB2 for zSeries Version 7.1 delivers improved performance, availability, and scalability. When DB2 for zSeries runs in a Parallel Sysplex environment, the availability of data and applications is very high. If one DB2 subsystem is unavailable because of maintenance, for example, other DB2 subsystems in the Sysplex take over the workload. Users are unaware that part of the system is unavailable because they have access to the data and applications that they need.

DB2's high availability features ensure maximum system availability and minimum downtime. DB2 uses standard operating system failover software and hardware configuration for high availability. The operating system failover software includes High Availability Clustered Multi-Processing (HACMP), Microsoft Cluster Server (MSCS) and Sun Cluster 2.x (SC 2.x) respectively for AIX, Windows NT and Sun Solaris.

1.4 Overview of DB2 Connect 7.1

DB2 Connect gives applications fast and easy access to existing databases on IBM enterprise servers. It has a highly scalable communications infrastructure for connecting Web, Windows, UNIX, and mobile applications to data.

DB2 Connect provides extensive application programming tools for developing client-server and Web applications using industry standard APIs such as ODBC, JDBC, and SQLJ.

Version 7.1 has been enhanced for both the application developer and the database administrator. There is new extended relational support that includes new XML, Spatial and high-speed searching extenders. Application development enhancements have been made along with a new cross-platform stored procedure language. There are new object-relational extensions to the SQL language and an updated DB2 Satellite Edition for the mobile worker. These are just a few of the new enhancements for this release of DB2 Universal Database.

1.5 Performance tools

DB2 has various built-in performance tuning and monitoring tools. These tools are very useful for measuring database server performance and finding any bottlenecks that occur when the PeopleSoft application calls the database for SQL requests. These tools give detailed and snapshot reports for database

activity. Once installed and configured, the PeopleSoft database is likely to grow rapidly. There need to be proper tools and utilities to check database consistency and update database statistics. The utilities available in DB2 help to maintain up-to-date database statistics.

The most important DB2 tools and utilities that benefit PeopleSoft applications are:

- Runstats
- Reorg and Reorgchk
- Visual Explain
- Performance Smart Guide Wizards from Control Center
- Event Monitor
- Snapshot Monitor
- DB2 Governor
- Performance Monitor
- Benchmark Utility.

1.5.1 Workload manager

This is a multi-user object-relational database for applications shared in a workgroup on PC-based LANs, which includes:

- DB2 Universal Database Workgroup Edition
- DB2 Extenders
- DB2 OLAP Starter Kit
- Application Development Client
- Administration Client
- Run-Time Client
- Net.Data
- WebSphere Application Server, Standard Edition
- QMF

1.6 Overview of WebSphere

WebSphere is Internet infrastructure software known as *middleware*. It enables companies to develop and integrate next-generation e-business applications, such as those for business-to-business. e-commerce supports business applications from simple Web publishing through enterprise-scale transaction processes. WebSphere transforms the way businesses manage customer, partner, and employee relationships.

1.7 Customer support

Both IBM and PeopleSoft deliver high-quality Internet-based and phone support. If you encounter an issue, first attempt to discover the source of the problem. If you are unable to resolve the issue, you should contact the appropriate supplier, PeopleSoft or IBM. If you are unable to isolate the problem, a call log should be opened with the PeopleSoft Global Support Center (GSC). Once the problem is analyzed by the PeopleSoft engineer(s), the problem may be directed to PeopleSoft Support or IBM Support until resolution is obtained.

1.7.1 PeopleSoft GSC

In North America, call the Global Support Center (GSC) at 1-800-4 PPLSFT (1-800-477-5738).

Outside of North America, consult the Customer Care directory located under Contact Us on PeopleSoft Customer Connection.

1.7.2 IBM customer support

IBM's commitment to customer support is a known fact. IBM continues to deliver the best support for all UDB products. There is an online customer support site for UDB products at the following URL:

```
http://www-4.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/
index.d 2w/report
```

For customers with a support contract, IBM support with expert assistance is available through a phone call. Find out about 24x7 customer support at the following Web site:

```
http://www-4.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/
help.d2 w/report
```

For general problems and solutions, there is an online knowledge base available where customers can query the knowledge base and get their solutions. The following URL takes you to the DB2 technical knowledge base:

```
http://www-4.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/
techlib.d2w/report
```

There are lot of Frequently Asked Questions (FAQs), Technotes, hints, and tips available on this site. IBM also releases various level of FixPaks. DB2 FixPaks contain updates and fixes for bugs (Authorized Program Analysis Reports, or APARs) found during testing at IBM, as well as fixes for bugs reported by customers. Every FixPak is accompanied by a document, called APARLIST.TXT, which describes the bug fixes it contains. For more information about FixPaks, see Chapter 8, “Installing and configuring the PeopleSoft application server” on page 99. Customers can get information about the latest level of FixPaks from this URL:

```
http://www-4.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/
download.d2w/report
```

2

Overview of PeopleSoft 8 on zSeries

This chapter describes the overall PeopleSoft architecture when the database resides on DB2 UDB for zSeries. We discuss all of the components that make up the design of the PeopleSoft 8 Internet architecture. We cover the following topics:

- Overall architecture
- What runs on each tier

If you know DB2 UDB and zSeries or S/390, but are new to PeopleSoft, this chapter should provide you with enough high-level information to help you understand the important concepts involved in the PeopleSoft Internet Architecture.

2.1 Basic architecture

With PeopleSoft Version 7.5, PeopleSoft introduced the three tier architecture as shown in Figure 2-1. This release added the PeopleSoft Application Server tier between the client and the database server. The 7.5 release allowed better performance for the clients largely due to the SQR compression that occurred on the application server and the added ability to cache on the application server as well as the client. In this architecture, we also made use of the file server which contained the PeopleSoft executables and the report server which ran crystal reports and nvision. The 7.5 version of PeopleSoft still required code on the client.

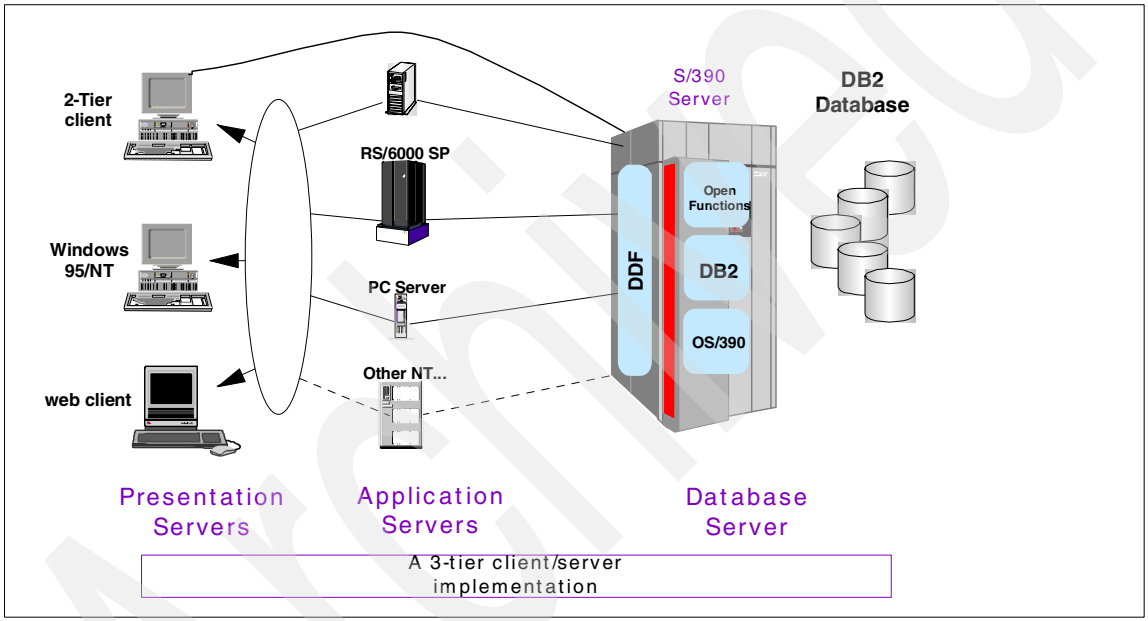

Figure 2-1 PeopleSoft 7.5 architecture

Version 8 of PeopleSoft introduced the PeopleSoft Internet Architecture (also known as PIA) and *no* code on the client, as shown in Figure 2-2. This version is the n- tier version of the architecture that adds the Web server to the basic architecture. Version 8 applications allow any user to access PeopleSoft as long as they have a Web browser on their desktop. We still have the fileserver which now only serves as an installation PC and the report server which has additional requirements. Considerations for this architecture are intranet versus Internet, where and what type of firewall strategy you intend to employ, and availability requirements for this new and fully Web enabled application.

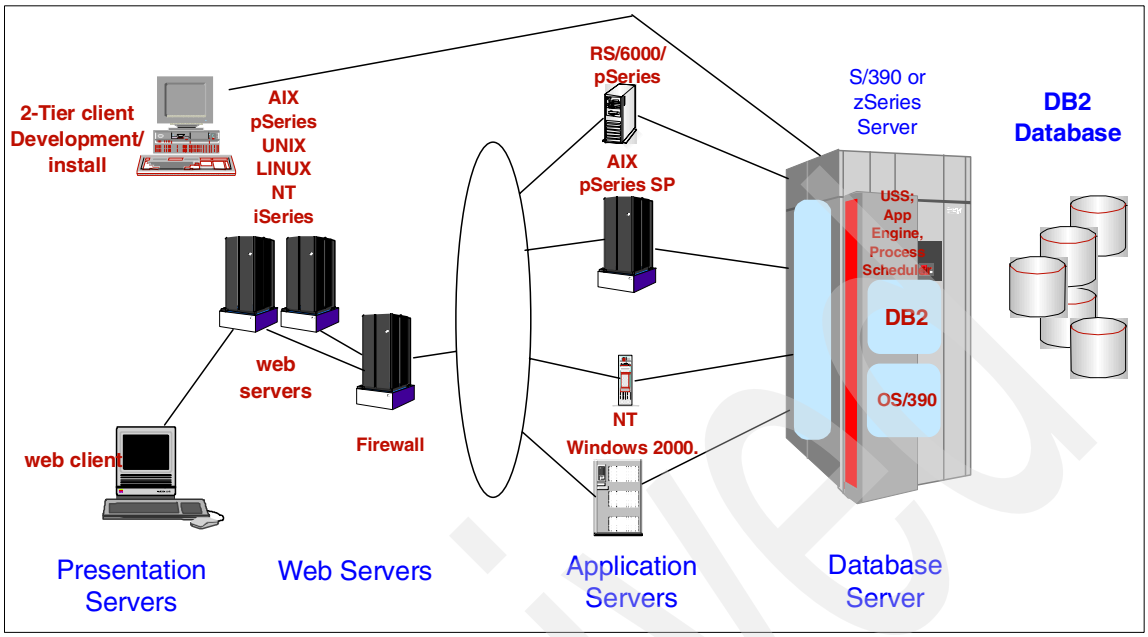

Figure 2-2 PeopleSoft Version 8 architecture

2.2 What runs on each tier?

In the next few paragraphs, we will discuss what runs on each of the various tiers and the required operating system requirements for each tier. PeopleSoft has a formal certification process for each release of their software. Please check PeopleSoft Customer Connection also for the most current requirements. PeopleSoft publishes only the minimum requirements for installation. We may have additional recommendations that may enhance performance or resolve an issue. Prior to installation, please also check the most current recommendations from IBM.

2.2.1 zSeries database server

In version 8, PeopleSoft re-wrote and optimized the process scheduler and application engine into C++. As a result, we have employed UNIX System Services (USS) in the PeopleSoft LPAR to run these products. We also run DB2 for the PeopleSoft Databases. DB2 is the only PeopleSoft certified Database for the PeopleSoft zSeries solution. The current minimum software requirements for installation of PeopleSoft 8.4 on DB2 for zSeries are as follows.

- OS390 V2.10 or z/OS 1.1 with OMVS running
- DB2 V6.1
- Cobol 2.2
- SQR version delivered with your PeopleSoft CDs
- Java 1.3

We highly recommend that you take advantage of the mature tools available on the zSeries platform to manage your PeopleSoft applications as shown in Figure 2-3. You may also refer to the PeopleSoft Customer Connection Web site:

```
http://www.peoplsoft.com
```

where you will find a White Paper entitled *Workload Manager and PeopleSoft*. You can prioritize the PeopleSoft work in your system by the function of the request if you employ WLM in Goal Mode. Additionally, given the implied availability requirement for an Internet application, you may wish to consider using DB2 for zSeries tools such as concurrent backup and reorg for your PeopleSoft applications. Your decision to deploy your PeopleSoft Database server on zSeries gives you the advantage of these mature tools.

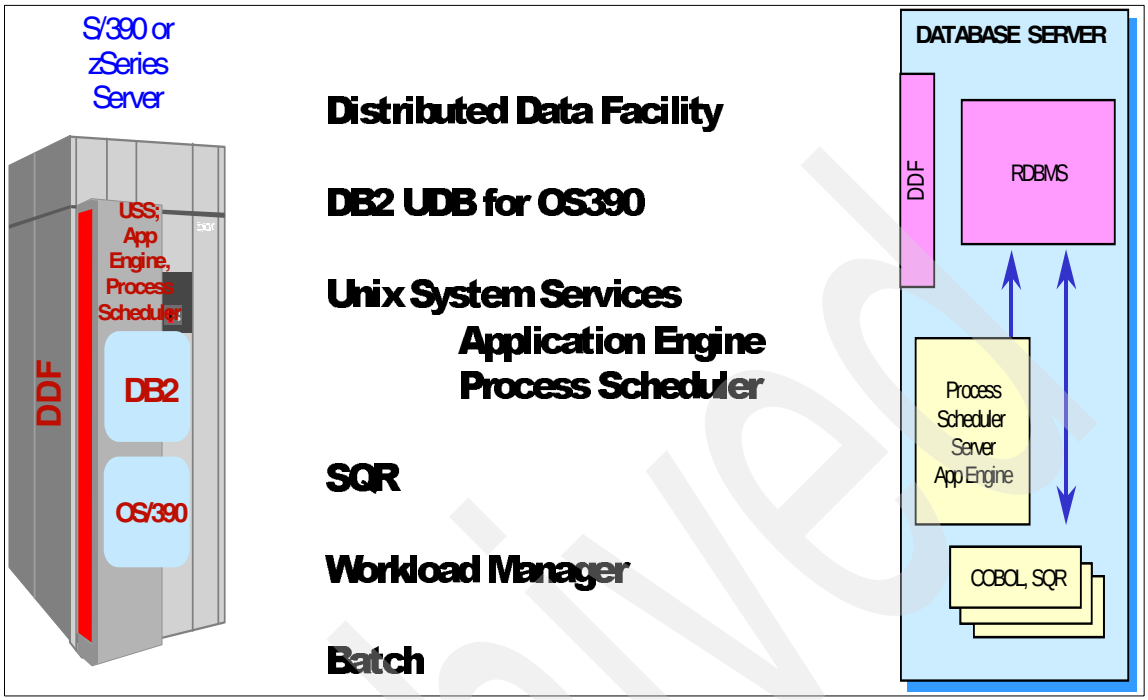

Figure 2-3 PeopleSoft zSeries tools

You will also need to be aware of the recommended UNIX System Services configuration for the best possible performance of application engine and process scheduler. Again, the most recent recommendations for USS configuration can be found on the customer connection Web site at:

```
http://www.peoplesoft.com
```

Additionally, PeopleSoft is written on one code line. This means that whatever database and platform the applications are deployed on will require tuning to give the best possible performance to the PeopleSoft applications. Again, there is a document entitled *Performance and Tuning for PeopleSoft 8.1 on OS/390* that was authored by Brian Holroyd and is available on Customer Connection. We highly recommend that you familiarize yourself with all of this documentation.

2.2.2 PeopleSoft application server in the zSeries architecture

Until very recently, the only certified operating systems for the PeopleSoft application server when using the zSeries database server were AIX or NT. As of this writing, both Sun Solaris 8 and HP/UX 11.11 are also certified application server operating systems as shown in Figure 2-4. We currently have no

performance data when using Sun or HP as application servers. If you are upgrading from PeopleSoft 7.5 to 8.4, you will find that the application server hardware requirement has increased by 2.5. The application server is primarily a transaction processor that compresses and expands and transforms SQL requests to and from the database server and to and from the Web server. The new functionality that this server performs is the translation of requests via Jolt received from the Web server to SQL. Additionally, the application server handles the new message and subscribe functions for PeopleSoft application messaging, business interlinks and business components. In order to communicate to DB2 zSeries, we also need to run connectivity software on the application servers. Currently, the only certified DRDA software is DB2 Connect. The following is a list of software requirements for your application server:

- AIX 5.1 or Windows 2000/NT or HP/UX 11.11 or Solaris 8
- DB2 Connect 7.1
- Cobol Compiler, depending on your operating system and the applications you are using
- The PeopleSoft application server Domains delivered on CDs
- Tuxedo/Jolt which are also delivered on your CDs

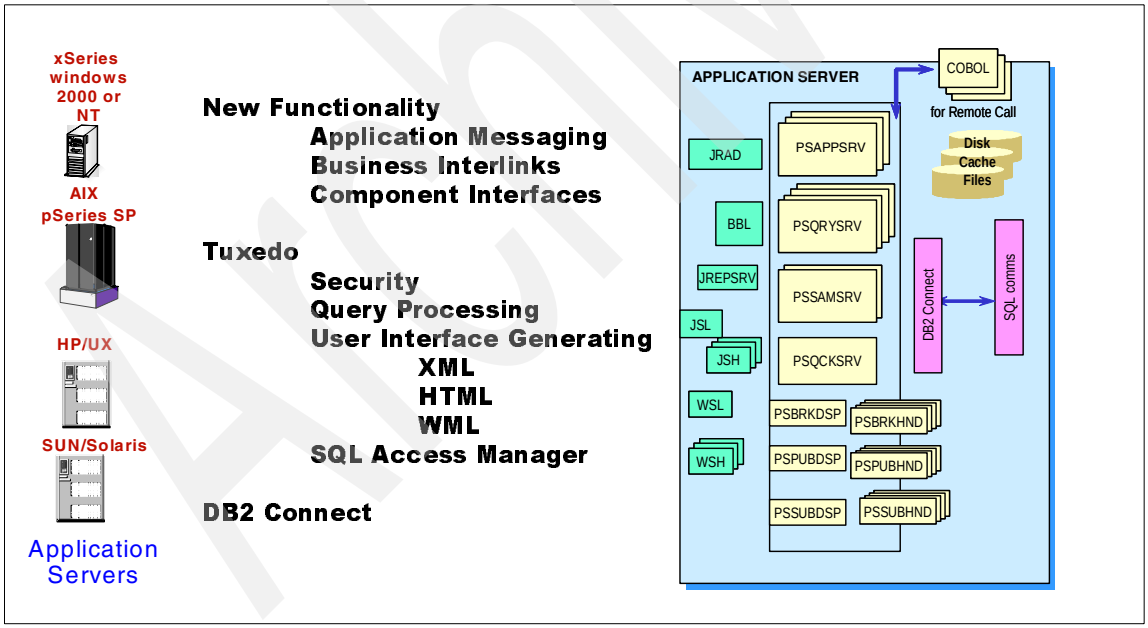

Figure 2-4 PeopleSoft application server zSeries architecture

2.2.3 PeopleSoft Web server in the zSeries architecture

The new PeopleSoft Web server is integral to the PeopleSoft application server and the PeopleSoft Internet Architecture as shown in Figure 2-5. In theory, the Web server can reside on the same hardware as the application server. We have two reasons for depicting this Web server as being a separate box. First, we have found that from a performance standpoint, the Web servers tend to scale better horizontally, and the application servers scale well either vertically or horizontally. In addition, we prefer a hardware and software firewall in our architectures as we believe this to be a more secure configuration. The following is a list of software requirements for your Web server:

- AIX 5.1 or Windows 2000/NT or HP/UX 11.11 or Solaris 8 or Linux
- WebLogic or WebSphere

Your PeopleSoft CDs will give you an option for installation of the Web server. You will be prompted at install to choose WebSphere or WebLogic for your PeopleSoft Web server. Both of these products are OEM and are delivered on the PeopleSoft CD.

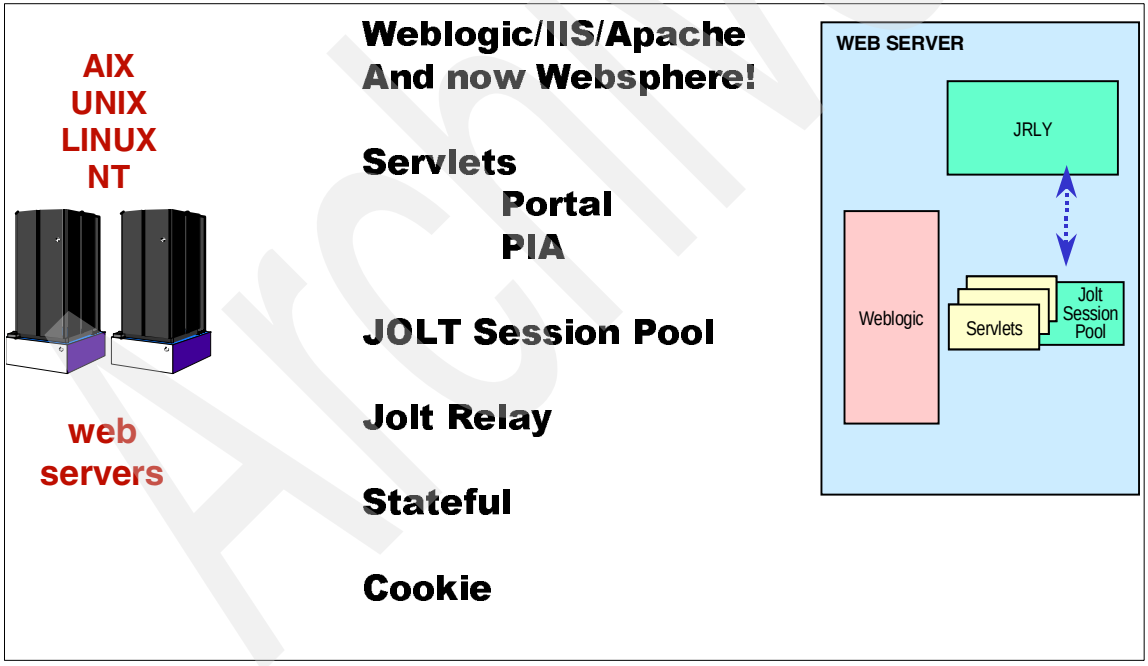

Figure 2-5 PeopleSoft Web server xSeries architecture

2.2.4 PeopleSoft batch servers in the zSeries architecture

For PeopleSoft 8.4, you will define a batch server on zSeries that will be your process scheduler for COBOL, SQR, and application engine jobs. In addition, you will also define a Windows 2000 or NT batch server for your NVision and Crystal Report processing. The difference in prior PeopleSoft 8.4 report servers is that all requested reports that are run by your users are viewed via their Web browser. This makes your report server a report repository also. Therefore, it is critical to establish a report archiving strategy early in your SLA process.

Important: *Please ensure that your company typically installs all of the Windows 2000 or NT executables. If that is not the current practice, please perform a full install of the Windows operating systems for your batch server. We have seen the batch server fail to install if this precaution is not heeded.*

2.2.5 PeopleSoft installation server in the zSeries architecture

You will also need a Windows 2000/NT installation server. This server needs a network drive of about 5 gig. You will use this server initially for installation and later for upgrade and data mover functions. This server also needs an installation of DB2 Connect.

3

Planning for installation

This chapter describes the planning activities you will undertake prior to the installation of the PeopleSoft applications. We discuss DB2 considerations and the environment where the team installed, and provide you with some information on the PeopleSoft ID structure.

3.1 Planning for installation

Immediately after your purchase of PeopleSoft and your zSeries platform decision, the questions will start. You will have questions about software and hardware requirements, how best to configure your infrastructure, and what type of system resources you need to run these applications.

PeopleSoft and IBM have done a great job at providing the answers to these questions in a wealth of documentation that can be found online at:

```
http://www.peoplesoft.com
```

We strongly recommend that you become familiar with all of the documentation pertaining to PeopleSoft and zSeries that is available there.

3.1.1 Prerequisites

The following is a list of minimum hardware and software requirements for PeopleSoft 8.4 installation on DB2 zSeries:

- G6 or better processor
- OS390 v2.10 or z/OS 1.1
- DB2 6.1
- COBOL 2.2
- USS enabled
- Java 1.3
- DB2 Connect 7.1 installed on all of the application servers and the installation server with connection tested
- AIX 4.3.3 or Windows 2000/NT or HP/UX 11.11 or Solaris 8 for the application servers and/or Web servers
- Windows 2000/NT for the file server and report server
- COBOL installed on the application servers and the file server
- 2 GB of space available to the file server

3.1.2 DB2 considerations

This section reviews a collection of DB2 installation and data administration topics from the following sources:

- Our PeopleSoft 8.4 installation experience
- Interviews with IBM consultants who have assisted in implementing or tuning PeopleSoft
- Interviews with PeopleSoft personnel
- Our DB2 experience

Administration of DB2 for PeopleSoft

PeopleSoft provides implementation and tuning information in a variety of documents. Our suggestions for obtaining information are the following:

- Review *PeopleTools 8.4: Installation and Administration of DB2 for OS/390.*

  Do not overlook the appendix in this manual, because it contains some valuable information. The PeopleBooks for the application you are installing may have a section on performance optimizations.
- Become acquainted with the PeopleSoft trace facilities. These tools provide an easy way to pinpoint SQL that may need tuning. Or use your current DB2 monitor to obtain similar information.
- PeopleSoft has one “code line.” That is, it is generically written for a variety of businesses, platforms, and DBMSs. As a result, you will need to customize it for your particular installation. This includes sizing the system, choosing DDL parameters, customizing indexes, and tuning SQL. Document all changes that you make and the reason for making them. These changes may need to be reapplied when you install the next release.

DB2 installation recommendations

We highly recommend that you create a unique DB2 subsystem for PeopleSoft. A list of PTFs recommended for OS/390 PeopleSoft installations may be found on the PeopleSoft Web site. Make sure you have these PTFs installed. Special consideration should be given to certain ZPARMS. Suggested values are the following:

- `CACHEDYN=Y` Enable dynamic caching.
- `CONDBAT=` (number greater than MAXDBAT)
- `DEFIXTP=2`
- `DSMAX =` Value greater than concurrently used datasets.

All tablespaces and indexes that have been defined may not be used. Set DSMAX to a number larger than the number of data sets that are used concurrently. Monitor the dataset open/close activity. If datasets are being opened and closed every few seconds, DSMAX should be increased. In general, for tablespaces you should specify `CLOSE YES`.

- `EDMPOOL` Monitor after installation and tune to maximize dynamic caching.
- `IDTHTOIN=0` This is the timeout value for active threads; setting it to zero means no timeout value.
- `INBUFF=60` This is the maximum size for the input buffer used for reading active and archive logs. The recommended value speeds recovery.
- `MAXDBAT=` (20% greater than the number of concurrent users you expect)
- `MAXKEEPD=5000`
- `NUMLKTS` Start with the default; this parameter may need to be adjusted.
- `NUMLKUS` Start with the default; this parameter may need to be adjusted.
- `CMTSTAT=INACTIVE` Leave thread inactive after commit.
- `RELCURHL=YES` Release data page or row lock after a commit for cursors is specified with hold.
- `TCPALVER=NO` Connection requests require a user ID and password.

Some tables should be placed in a separate tablespace. We recommend that you assign to a separate tablespace any table with a name of the form:

- - Some-name_TEMP
- - Some-name_TMP
- - Some-name_WRK

This gives you the ability to move the data set containing the tablespace to minimize contention during normal operation, plus it will facilitate a recommendation to follow, relative to the use of the Reorg utility. PeopleSoft has grouped tables that will grow or experience high update activity in tablespaces that are named xxLARGE. Consider placing these tables in partitioned tablespaces. Tables that experience high activity rates should be moved to their own tablespace. TABLESPACE Parameter - Locksize Specify LOCKSIZE(PAGE) unless the PeopleSoft documentation recommends LOCKSIZE(ROW). Use row-level locking only when needed for concurrency. Row-level locking will increase CPU overhead. Some PeopleSoft applications have long Units of Work. Thus, specifying row-level locking may require adjustments to the ZPARM parameters.

NUMLKTS and NUMLKUS or tablespace parameter LOCKMAX: the maximum value for NUMLKUS (locks per user) is 100000. `NUMLKUS=0` disables this limit.

Each lock requires 250 bytes of storage. Increasing the number of locks allowed may require that the MAXCSA parameter in the IRLM startup procedure be increased if IRLM parameter `PC=NO` is used. Before increasing MAXCSA, ensure that the OS/390 specifications for ECSA can accommodate the increase.

If you find it necessary to set `NUMLKUS=0`, you may want to consider setting the IRLM startup procedure PC parameter to `PC=YES`. There is some additional CPU overhead in doing this, but it removes the potential of bringing down z/OS if the ECSA fills up.

DB2/PeopleSoft 8.4: EDM pool recommendations

For various PeopleSoft and financial implementations, the EDM pool size range from 50 MB to 120 MB. The major variants are the number of objects included in the database and the amount of storage devoted to caching of dynamic plans.

DDL Model defaults are stored in tables PSDDLMODEL and PSDDLDEFPARMS. If you choose to use native DB2 DDL to define objects different from the way they are shipped with PeopleSoft, then the DB2 catalog will differ from the values in the PSDDLMODEL and PSDDLDEFPARMS table. If you then use PeopleTools to recreate or alter a table, the values from the PSDDLMODEL and PSDDLDEFPARMS table will be used. Therefore, you will need to devise a means of updating the PSDDLMODEL and PSDDLDEFPARMS to reflect the values in the DB2 catalog. SETDBNAM.SQR, SETINDEX.SQR, and SETSPACE.SQR have been developed to assist you with this task.

DB2/PeopleSoft 8.4: buffer pool recommendations

Following the PeopleSoft 8.4 installation procedure will cause all 4K-page tablespaces and indexes to be assigned to buffer pool 0 (BP0). The following are general recommendations for initially configuring the DB2 buffer pools in support of PeopleSoft 8.4. The recommendations assume that the DB2 subsystem is dedicated to PeopleSoft 8.4: and are presented as percentages of the number of buffers you assign to the DB2 buffer pools.

- BP0 - 10%. Restrict BP0 to the DB2 catalog and directory.This facilitates the dynamic SQL that PeopleSoft 8.4 executes.
- BP1 - 15%. BP1 supports large tables. The large tables will generally have a low buffer reuse. Consequently, 15% is moderate relative to the other PeopleSoft buffer pool assignments.
- BP2 - 25%. BP2 is dedicated to large-table indexes. Indexes tend to have a high buffer reuse. Consequently, we recommend a high buffer allocation.
- BP3 - 10%. BP3 is dedicated to the smaller PeopleSoft 8.4 tables.Code tables are usually smaller and tend to experience high access. This large buffer pool allocation facilitates data-in-memory processing of the frequently accessed smaller tables.

- BP4 - 15%. BP4 is dedicated to small-table indexes. Indexes tend to have a high buffer reuse. Consequently, we recommend a high buffer allocation.
- BP5 - 10%. BP5 is dedicated to PeopleSoft tool tables. PeopleSoft tool data is primarily read-only.
- BP6 - 10%. BP6 is dedicated to PeopleSoft tool indexes. Indexes tend to have a high buffer reuse.
- BP7 - 5%. BP7 supports DB2 temporary storage (DSNDB07). The SQL Order By or Group By, for example, have the potential to require a DB2 sort which DB2 supports through the use of DSNDB07.BP32K (minimum). This is support for any 32K buffer pool requirement. Except for BP0, all these are arbitrary buffer pool assignments. Note, however, that DB2 requires BP0 for the catalog and directory tablespaces. Obviously you may, for example, use BP2 to support DSNDB07. The key is the concept of differentiating large tables, small tables, PeopleSoft tool tables, and the associated indexes.

DB2/PeopleSoft 8.4: index usage

In tuning the product for your environment, you will probably add additional indexes and may drop some of the initial indexes. Some of the indexes that install with the product are redundant and are not used. However, the PeopleSoft documentation states that indexes containing an underscore in the name, such as PS_ should not be touched. Also, indexes with a letter, such as PSA or PSB, should not be touched.

DB2/PeopleSoft 8.4: point-in-time recovery recommendations

The usual reason for a point-in-time recovery is an application programming error or a flawed operational procedure. Unfortunately, this exposure is always present, regardless of your hardware/software configuration. Additionally, a point-in-time recovery has the potential to be the most disruptive outage you are likely to encounter. The reason is that in a PeopleSoft 8.4 environment, you may need to recover all objects to a prior point-in-time. Depending on how you have mapped your tables to tablespaces, this could be from twenty to over a thousand tablespaces to recover along with a few thousand indexes. Your usual point-in-time recovery techniques, which you probably regard as conventional at this time, may be insufficient in this environment.

Point-In-Time recovery using DB2 utilities

The scenario for a point-in-time recovery using DB2 utilities is:

1. Determine the set of tables that are logically related.Typically this is a subset of the tables that make up the application database.However, it may be all PeopleSoft 8.4 tablespaces and indexes. Optionally, execute the QUIESCE utility on all of the tables that are candidates to be reset if a point-in-time

recovery is required. This establishes a point of consistency and causes the DB2 buffers for the quiesced tables to be externalized to DASD.

2. Execute the COPY utility on all of the tablespaces that are candidates to be reset if a point-in-time recovery is required. The obvious reason for this step is to back up the data. However, COPY will fail if it cannot externalize the DB2 buffers to DASD. That is the reason we invoked the QUIESCE utility first (to remove one reason why COPY may fail). You may want to place your Image Copy output on DASD and have DFHSM migrate the copies to tape.
3. For the second time, execute the QUIESCE utility on all of the tables that are candidates to be reset if a point-in-time recovery is required. This invocation of QUIESCE will establish a log RBA which will be the point of consistency. When it is necessary to recover to this point of consistency, RECOVER to the RBA established by the second invocation of the QUIESCE utility.

With the combination of the COPY and the second QUIESCE, the RECOVER TORBA will perform as efficiently as a Recover TOCOPY would perform, assuming no logging during the execution of this scenario. When it is necessary to recover to this point of consistency, RECOVER all indexes on all of the tables that have been reset to the prior point of consistency. The indexes must be synchronized with the data in the recovered tablespaces.

3.1.3 Our installation environment

The environment we used to install PeopleSoft consisted of the following:

- An z/OS LPAR in a SYSPLEX with z/OS Version 1.2 and the TCP/IP supplied with that version for the database server as well as OMVS enabled and configured as outlined in the PeopleSoft and USS documentation.
- DB2 V7.1
- DB2 Connect V7.1
- A PC with Windows NT for the application server
- A PC with Windows NT for the Web server
- A PC with Windows 2000 for the installation, file server, and NT batch server.

3.1.4 De-mystifying the PeopleSoft ID

Note: This section was authored by Randall Hicks of PeopleSoft and appears here in its entirety, courtesy of PeopleSoft.

The purpose of this section is to identify each of the IDs that are defined or need to be defined in a PeopleSoft 8 implementation, along with the role each ID plays, where it is relevant, and its relationship to the other IDs that need to be defined to implement a PeopleSoft application.

The following discussion will detail each of the IDs and explain them in the context of Two-Tier Client connections. Following this section is another detailing the log in process in the context of Tee-Tier connections used by Application Server and Process Scheduler.

There are potentially six IDs referred to in running PeopleTools 8.12 based applications on DB2 UDB for OS/390. Three are unique to the PeopleSoft application, two are unique to DB2 and one is unique to the OS/390 Security software. Two must be defined as logon IDs and one as a Group ID in the OS/390 mainframe security software, such as RACF. We will describe each ID, in order of use as you proceed through the logon process, its role in the logon, and subsequent object access processes.

PeopleSoft IDs

The PeopleSoft IDs are as follows:

User ID

The User ID is the ID associated with the individual user logging onto the PeopleSoft application, via PeopleTools. It requires no OS/390 mainframe security definition or table access. Validation of this ID is performed by the PeopleTools themselves at log-in time, by virtue of verifying that a row has been defined in the PSOPRDEFN table for this ID. The row in the PSOPERDEFN table contains information pertinent to PeopleSoft internal security administration, such as the access profile defining what areas of the application this ID will have access to, and the ACCESS ID (defined later) associated with the User ID. PeopleSoft does provide an exit to validate the User ID on the mainframe server, if you choose to invoke an additional layer of security. Multiple User IDs can be defined in the application.

Connect ID

The Connect ID is the actual ID that makes the initial log on connection to the database. This ID does require definition in the OS/390 mainframe security software, and it requires SELECT ONLY access to four tables: PS.PSDBOWNER (one per DB2 subsystem), PSSTATUS, PSOPRDEFN, and

PSACCESSPRFL. Once the row is located in PSOPRDEFN associated with the User ID entered, the associated Access ID is validated within PeopleTools, and the Connect ID is disconnected, and the Access ID is logged on. The Connect ID is defined as a parameter within the PeopleTools Configuration Manager, on the file server where the PeopleSoft executables reside. If multiple servers are established, they may each have their own unique Connect ID, or they may all share the same Connect ID.

Access ID

The Access ID is the ID under which all DML is initiated. This ID requires definition in the OS/390 mainframe security software. It has a very close relationship to the DB2 ID “Owner ID”. The Access ID may be the actual Owner ID (not permitted in most shops), or it may have the Owner ID defined as a Secondary Authorization ID. Because of this relationship to the Owner ID, the Access ID has very liberal database access, and has the potential to do anything the Owner of an object can do, including dropping that object. After the Connect ID is disconnected and the Access ID logged on, the SQL objects referenced in the SQL statements are no longer fully qualified, and rely on the DB2 concept of Current SQLID to identify the owner of the specific object to be accessed. You may define as many, or as few Access IDs as you choose, with the exception that if the Access ID is the DB2 Owner ID, you can only define one

DB2 IDs

The DB2 IDs are as follows:

Owner ID

The Owner ID is the ID that DB2 recognizes as the owner of the DB2 Objects (tables and views). This value populates the CREATOR field in the system catalog table SYSIBM.SYSTABLES. The Owner ID is the common value that groups together all the various objects that comprise a PeopleSoft database within a DB2 subsystem. For PeopleSoft purposes, this value is stored in the PS.PSDBOWNER table, which is the first table queried during the Sign on process.

In order to perform DML activity against an object, the requestor must either be the owner of that object, or be granted rights to perform that activity by the owner. The PeopleSoft applications function on the premise that the requestor *is* the owner, either by direct association, or tough the use of Secondary Authorization processing and the SET CURRENT SQLID function. A DB2 User may create DB2 objects on behalf of another Owner ID, with the proper DB2 authority (usually DBADM or SYSADM). In this scenario, in the SYSIBM.SYSTABLES table, the CREATOR and CREATEDBY field will contain different values. The CREATOR is still considered the object's owner.

Current SQLID

The Current SQLID is the user ID that is identified as “active” within DB2 at the time an SQL action is requested. The function of this ID is to be used in validating access rights to a given object, and to serve as the “default” Owner ID, when the objects in an SQL statement are not fully qualified. As I indicated in two of the definitions above, PeopleSoft does not fully qualify the objects in the DML statements it issues (except at initial sign on), and it assumes that the requestor of the DML activity is the owner of those objects. By not fully qualifying the objects in the DML statements and assuming the requestor is the owner of those objects, PeopleSoft requires that the Current SQLID (that is, the requestor) be the Owner ID of the objects.

The Current SQLID can be the Owner ID of these objects via two scenarios.

The first is where the PeopleSoft Access ID has the same value as the DB2 Owner ID. When the Connect ID is disconnected and the Access ID is logged on, the Current SQL ID becomes the Access ID. At the time an SQL request is submitted, the PeopleSoft Access ID is the DB2 Current SQLID, and because the Current SQLID has also the same value as the Owner ID, the statement is successfully processed. The Owner ID has full access rights to objects it owns, and the objects are identified as belonging to the Owner ID, because each object is qualified, by default, with the Current SQL ID, which is the same as the Owner ID (and the PeopleSoft Access ID). In this scenario, you can only have one Access ID defined in the PeopleSoft application. All PeopleSoft User IDs must be associated with the same Access ID. This scenario would be a very unlikely situation, however, due to EDP audit standards, which seldom permit the definition of a DB2 Owner ID as an ID with the rights to connect to the OS/390 system (access IDs must have these rights, so if the PeopleSoft Access ID and the DB2 Owner ID have the same value, you are, in effect, granting these connection rights to the Owner ID).

More typical is the second scenario, which uses the Secondary Authorization functionality. By issuing a `SET CURRENT SQLID` statement the active Current SQLID within DB2 can be changed to another value. In this scenario, the PeopleSoft Access ID is not the same value as the DB2 Owner ID. When the Connect ID is disconnected, and the Access ID is logged on, the Current SQLID again becomes the Access ID, but the Owner ID of the DB2 objects is not the same as either the Access ID or Current SQLID. The PeopleTools will recognize this (at log on the Owner ID is identified in the query against the PS.PSDBOWNER table) and to synch up the Current SQLID with the Owner ID, they will issue the SQL statement `SET CURRENT SQLID = 'owner id'`, which will change the “active” SQLID value from the PeopleSoft Access ID to the DB2 Owner ID. In this way, the Owner ID does not need to be defined as an ID with log on rights, and the Access ID does not need to be defined as the owner of the database objects, or granted specific access rights to the objects that the owner

would otherwise have. The ability to perform this “transformation” of Current SQLID value does require special configuration in the OS/390 Security software which leads us to a sixth possible ID: the Secondary Authorization Group ID.

Secondary Authorization Group ID

The Secondary Authorization Group ID is nothing more than an ID which is established in the OS/390 Mainframe security software, to which the members defined in the Group can set themselves equal to. The Secondary Authorization Group ID must be the same value as the Owner ID. When the PeopleSoft Access ID is defined within this Secondary Authorization Group, it is said to have the Owner ID as a Secondary Authorization ID, and therefore the permission to issue the SET CURRENT SQLID statement, transforming the Current SQLID value from the PeopleSoft Access ID to the DB2 Owner ID.

Table 3-1 summarizes the IDs.

Table 3-1 ID summary table

ID	Type	Value Characteristics	Defined In PeopleSoft Security	Defined in OS/390 Security
User ID	PeopleSoft	Always unique	Required	Optional
Connect ID	PeopleSoft	Always unique	Required	Required
Access ID	PeopleSoft	May be unique. May be same as DB2 Owner ID	Required	Required
Owner ID	DB2	May be unique. May be same as PeopleSoft Access ID	No	No (If the Access ID and Owner ID are the same value, it could be considered required, but defining this ID in OS/390 security will have already been addressed in the Access ID entry)

ID	Type	Value Characteristics	Defined In PeopleSoft Security	Defined in OS/390 Security
Current SQLID	DB2	Variable value that can be the PeopleSoft Access ID, but must be equal to the DB2 Owner ID to perform DML operations.	No	N/A It has no identity of its own.
Secondary Authorization Group ID	RACF or equivalent	Must be the same value as the DB2 Owner ID	No	Yes, but as a Group ID rather than a Logon ID

3.1.5 Typical PeopleSoft/DB2 logon process

The user will initiate the Log On process by entering the Database Name, and a User ID and Password in the appropriate fields of the PeopleSoft Log On Dialog Box.

A connection request will be made to log on to the database using the Connect ID and Password stored in the PeopleTools Configuration Manager. After the Connect ID and Password are validated by the mainframe security, an SQL Select statement will be issued to select the appropriate DB2 Owner ID value from table PS.PSDBOWNER associated with the database name entered by the User. With the DB2 Owner ID now stored in memory, a fully qualified Select statement is processed against table PSSTATUS, followed by another fully qualified Select statement against table PSOPRDEFN. The Select statement will look for a row in PSOPRDEFN with the User ID value the User entered when initiating the Log In. If the row is found, the User Password is validated, and the associated Access ID and Password will be stored in memory. A fully qualified Select statement is processed against table PSACCESSPRFL to obtain additional parameters. The Connect ID is disconnected, and a new connection request is initiated with the Access ID and password, retrieved from memory. The Access ID and Password are validated by the mainframe security, and the connection is completed. If PeopleTools detects that the Access ID is a value different from the Owner ID value, it will issue a `SET CURRENT SQLID = "Secondary Authorization Group ID"` statement to change the value of the Current SQLID, and processing will continue, assuming that the Current SQLID is now the Owner ID of the database objects.

The following tables detail the various IDs and connection processes for Process Scheduler and Application Server for the DB2 OS/390 platform (Tee-Tier connections).

Application Server - NT

Table 3-2 shows the IDs for the NT Application Server.

Table 3-2 Application Server - NT IDs

ID Type	Validated in PSOPRDEFN	Validated in MainFrame Security	Comments
Connect ID	No	Yes	Initial Database connection is made with this ID. It must have SELECT access to four tables - PS.PSDBOWNER, PSSTATUS, PSOPRDEFN and PSACCESSPRFL
Opr ID (App Server Configuration parameter)	Yes	No	ID is only validated in PSOPRDEFN
Opr ID entered by user to log on to database via Tee Tier	Yes	Possible	Configuration Parameter: Validate Sign-on with Database=? 0 = do not validate (default) 1 = validate

Connection process

When the Application Server is booted, a call to the database is made over an ODBC connection with ConnectID and Password. ConnectID and Password are validated in Mainframe security package. Connection is made by DB2 Connect tough DDF (Distributed Data Facility). The tableowner ID is obtained by a Select from the PS.PSDBOWNER table. The next call is made to PSSTATUS to verify the Tools Version. The next transaction is a Select from PSOPRDEFN to validate that the Opr ID (App Server) has been defined in the Tools security tables. An additional select is made to PSACCESSPRFL to obtain the Access ID. The current connection of the Connect ID is disconnected and a new connection made with the Access ID, which then becomes the current thread owner. A `SET CURRENT SQLID` statement is issued, and the connection process is completed.

Note: For NT Application Servers, the `SET CURRENT SQLID` statement is issued only when it detects that the ACCESSID is not the same as the OWNERID found in PSSTATUS.

When a user attempts to log on to the database via PIA or the traditional three tier connection, the Opr ID (User) they enter is validated in Mainframe Security, if the “Validate Signon with Database” flag is set to `1`, otherwise this step is bypassed, and the Opr ID is only validated in PSOPRDEFN.

Application Server - AIX

Table 3-3 shows the IDs for the AIX Application Server.

Table 3-3 Application Server - AIX IDs

ID Type	Validated in PSOPRDEFN	Validated in MainFrame Security	Comments
UNIX Logon ID	No	No	ID is used to connect to UNIX machine, but must have authority to execute PSADMIN and have access to all App Server directories.
Connect ID	No	Yes	Initial Database connection is made with this ID. It must have SELECT access to four tables - PS.PSDBOWNER, PSSTATUS, PSOPRDEFN and PSACCESSPRFL
Opr ID (App Server configuration parameter)	Yes	No	ID is only validated in PSOPRDEFN
Opr ID entered by user to log on to database via Tee Tier	Yes	Possible	Configuration Parameter: Validate Signon with Database=? 0 = do not validate (default) 1 = validate

Connection process

When the Application Server is booted, a call to the database is made over an ODBC connection with ConnectID and Password. ConnectID and Password are validated in Mainframe security package. Connection is made by DB2 Connect though DDF (Distributed Data Facility). The tableowner ID is obtained by a Select from the PS.PSDBOWNER table. The next call is made to PSSTATUS to verify the Tools Version. The next transaction is a Select from PSOPRDEFN to validate that the Opr ID (App Server) has been defined in the Tools security tables. An additional select is made to PSACCESSPRFL to obtain the Access ID. The current connection of the Connect ID is disconnected and a new connection made with the Access ID, which then becomes the current thread owner. A `SET CURRENT SQLID` statement is issued, and the connection process is completed.

Note: For AIX Application Servers, the `SET CURRENT SQLID` statement is only issued when it detects that the ACCESSID is not the same as the OWNERID found in PSSTATUS.

When a user attempts to log on to the database via PIA or the traditional tee tier connection, the Opr ID (User) they enter is validated in Mainframe Security, if the "Validate Signon with Database" flag is set to `1`, otherwise this step is bypassed, and the Opr ID is only validated in PSOPRDEFN.

Process Scheduler - NT

Table 3-4 shows the IDs for the NT Process Scheduler.

Table 3-4 Process Scheduler - NT IDs

ID TYpe	Validated in PSOPRDEFN	Validated in MainFrame Security	Comments
Connect ID	No	Yes	Initial Database connection is made with this ID. It must have SELECT access to four tables - PS.PSDBOWNER, PSSTATUS, PSOPRDEFN and PSACCESSPRFL

ID TYpe	Validated in PSOPRDEFN	Validated in MainFrame Security	Comments
Opr ID (Process Scheduler configuration parameter)	Yes	No	ID is only validated in PSOPRDEFN

Connection process

When the Process Scheduler is booted, a call to the database is made over an ODBC connection with ConnectID and Password. ConnectID and Password are validated in Mainframe security package. Connection is made by DB2 Connect though DDF (Distributed Data Facility). The tableowner ID is obtained by a Select from the PS.PSDBOWNER table. The next call is made to PSSTATUS to verify the Tools Version. The next transaction is a Select from PSOPRDEFN to validate that the Opr ID (Process Scheduler) has been defined in the Tools security tables. An additional select is made to PSACCESSPRFL to obtain the Access ID. The current connection of the Connect ID is disconnected and a new connection made with the Access ID, which then becomes the current thread owner. A `SET CURRENT SQLID` statement is issued, and the connection process is completed.

Note: For NT Process Schedulers, the `SET CURRENT SQLID` statement is only issued when it detects that the ACCESSID is not the same as the OWNERID found in PSSTATUS.

All processes released by Process Scheduler will go through a similar user authentication.

Process Scheduler - USS

Table 3-5 shows the IDs for the USS Process Scheduler.

Table 3-5 Process Scheduler - USS IDs

ID Type	Validated in PSOPRDEFN	Validated in MainFrame Security	Comments
USS Log On ID	No	Yes	USS Log On ID must have the following authorities: - execute on psconfig.sh - execute on psadmin- Process Scheduler -Select access to PS.PSDBOWNER -Select access to PSSTATUS -Select access to PSOPRDEFN -Select access to PSACCESSPRFL -Authority to issue SET CURRENT SQLID = 'Objowner' statement
Connect ID	No	No	The connect ID serves no functional purpose in a USS configuration and is only provided because Process Scheduler requires some value for this parameter.
Opr ID (Process Scheduler configuration parameter)	Yes	No	ID is only validated in PSOPRDEFN

Connection process

When Process Scheduler is booted, ODBC for OS/390 makes a connection via CAF (Call Attach Facility) with the USS Log On ID. If Process Scheduler is started by submission of a JCL job using the BPXBATCH utility, the DB2

authentication uses the mainframe User ID under which the JCL was submitted, when ODBC for OS/390 make the initial connection. The tableowner ID is obtained by a Select from the PS.PSDBOWNER table. The next call is made to PSSTATUS to verify the Tools Version. The next transaction is a Select from PSOPRDEFN to validate that the Opr ID (Process Scheduler) has been defined in the Tools security tables. An additional select is made to PSACCESSPRFL to obtain the Access ID. A thread disconnect is issued, and a new connection made with the Access ID. In a CAF connection, the SQLID of the connection is still the originating USS LogOn ID. The only way to change this is by issuing a `SET CURRENT SQLID` statement. Tools releases before 8.13 were not issuing this command if the Access ID and tableowner ID were the same. In USS, PeopleSoft's ODBC API now always issues the `SET CURRENT SQLID` statement.

In a batch mode, the owner of processes run within USS, namely AE and the Process Scheduler Job itself, will be the USS Logon ID. The same holds true for COBOL and SQR submitted via process Scheduler, provided the customer has not customized the JCL shell scripts with a hardcoded USER=xxxx card.

Otherwise, jobs submitted via traditional JCL outside of Process Scheduler will be owned by either the USER= card value, or by the User ID under which the job was submitted. In this case, it is critical that the USER of the process either be the same as the USS UserID, or belong to the same UNIX group, to avoid file permission problems, when writing output to the HFS for logs and Report repository.

3.1.6 General installation notes

There are some specifics that your I/T people should verify prior to installation of your PeopleSoft applications and we list some of them here:

- Verify that your IBM sizing has been completed and the hardware resources have been allocated.
- Check the PeopleSoft Hardware and Software Guide to verify that you have all of the prerequisites for installation completed. (available on customer connection).
- Check all PTF and Fix levels. PeopleSoft required PTFs can also be found listed on customer connection.
- If your DB2 datasets are using pooled DASD, please verify that the VTOCS are defined with enough space to handle the 10000s of objects delivered with PeopleSoft.
- Establish and meet weekly with a team comprised of resources form all IT disciplines (systems for all platforms in the architecture, database, network, functional, project team, and integrators).

- Determine what your sustaining support tactics for this new application will consist of (backup, recovery, reporting, change and problem, monitoring, helpdesk, escalation).
- Establish service levels for your PeopleSoft applications and determine the availability requirements for this new 100% Internet and mission critical application.
- Determine your security strategy.
- Read all documentation pertaining to PeopleSoft on zSeries at least twice prior to the install!

3.1.7 Complete installation checklist

For each item, we list three pieces of information where appropriate:

- Transfer Parameters
- Site-Specific Value
- Sample Value

1. OS390 data set high level qualifier

- The high-level qualifier used for PeopleSoft COBOL and SQR datasets
- Suggested Default: HLQ.ppvvv where ppvvv is the PeopleSoft product and release (such as HR840.
- Sample value: `PS.HR840`

2. PeopleSoft file server high level directory

- <PS_HOME>-the directory to which you installed the PeopleSoft CD-ROM, such as N:\HR840.
- Sample value: `N:\HR840`

3. Target directory for generated files

- The workstation directory that will contain a variety of files generated when you run the PeopleSoft Server Transfer program, including file transfers, COBOL compile JCL, and translated SQRs (for example [and] translated to \)
- Suggested Default: <PS_HOME>\STAGE
- Sample value: `N:\HR840\STAGE`

4. File transfer method

- Indicates which file transfer protocol will be used. Supported protocols include 3270/Send and Microsoft FTP
- Suggested Default: Microsoft FTP
- Sample value: Microsoft File Transfer Protocol (FTP)

5. Database server host/node name

- FTP Only: Symbolic IP Name For OS/390 System
- Suggested Default: IP Name of Server
- If you do not use FTP, specify any alphanumeric character in this field

6. Database server login ID

- FTP Only: OS/390 user ID used to connect to OS/390 server and create files and directories on USS
- Note: This ID MUST be in the same UNIX GROUP as any IDs under which subsequent batch processing will be executed
- Suggested Default: LOGONID
- If you do not use FTP, specify any alphanumeric character in this field
- Sample value: `USER1`

7. Job card line 1

- This is the first line of a job card that will be inserted into JCL files by the PeopleSoft Server Transfer program.
- Enter // in first two positions followed by job card information such as job name, keyword JOB, account information, and so on. If the job card extends to two lines, end the first line with a comma and complete Job Card Line 2.
- Note: Any USER= parm coded MUST be in the same UNIX GROUP as the Database Server Logon ID noted in Parameter 6 above.
- Sample value: `//PSHR840 JOB (PSOFT),'J',CLASS=A,MSGCLASS=A`

8. Job card line 2

- This is the second line of a job card that will be inserted into JCL files by the PeopleSoft Server Transfer program. Enter // in first two positions followed by at least one space before continuing to add job card information.
- Suggested Default: //*
- Sample value:
  `//REGION=4096K,MSGLEVEL=(1,1),USER=BATCHID1,PASSWORD=BPSWD1`

9. Job card line 3

- This is the third line of a job card that will be inserted into JCL files by the PeopleSoft Server Transfer program. Enter // in first two positions followed by at least one space before continuing to add job card information.
- Suggested Default: //*
- Sample value: `//  NOTIFY=&SYSUID`

10. OS390/DB2 operator ID

- This parameter is for PeopleSoft internal use.
- Let it default to OPRID
- Sample value: `OPRID`

11. OS390/DB2 table owner ID

- This is the PeopleSoft table owner ID-the high-level qualifier for DB2 tables (aka “CREATOR” in the IBM SYS Catalog tables). If you are using secondary authorization, this will be your secondary authorization ID, otherwise it will be your primary authorization ID.
- Suggested Default: PSOWNER
- Sample value: `PS001`

12. DB2 subsystem name

- This is the DB2 Subsystem used for the PeopleSoft application you are currently installing (DMO or SYS).
- Suggested Default: DDDD
- Sample value: `DSNT`

13. DB2 system data set containing DSN member

- This is the DB2 system dataset that contains member DSN.
- Suggested Default: SYS1.DB2.DDDD.DSNLOAD
- Sample value: `DSN610.SDSNLOAD`

14. DB2 system data set containing DSN3@ATH member

- This is the DB2 system dataset that contains member DSN3@ATH. DSN3@ATH is a sample authorization exit. By implementing the sample authorization exits you can provide group names as secondary authorization IDs.
- Suggested Default: SYS1.DB2.DDDD.EXIT
- Sample value: `DSN610.SDSNEXIT`

15. DB2 system data set containing DSNTEP2 member

- This is the DB2 runtime system dataset containing member DSNTEP2.
- Suggested Default: SYS1.DB2.DDDD.RUNLIB.LOAD
- Sample value: `DSN610.RUNLIB.LOAD`

16. PeopleSoft database name

- Suggested Default: DB
- Sample value: `PSHR840`

17. Plan name for PTPSQLRT via TSO Attach Facility

- This is DB2 Plan used by PTPSQLRT (the COBOL/DB2 API used by COBOL batch and process scheduler jobs).
- Suggested Default: PTPSQLRT
- Sample value: `PTPSQLRT`

18. Plan name for PTPSQLRT via Call Attach facility (for USS)

- Suggested Default: PTPSQLRA
- Sample value: `PTPSQLRA`

19. Language Environment runtime library (i.e., IGY.V2R2M0.SCEERUN)

- Suggested Default: IGY.V2R2M0.SCEERUN
- Sample value: `IGY.V2R2M0.SCEERUN`

20. Language Environment linkedit library (i.e., CEE.SCEELKED)

- Suggested Default: CEE.SCEELKED
- Sample value: `CEE.SCEELKED`

21. COBOL system data set name containing IGY* Members

- This is the COBOL load library containing modules used by COBOL compiler. Its members include IGYCASM1, IGYCINIT, and so on.
- Suggested Default: IGY.V2R2M0.SIGYCOMP
- Sample value: `IGY.V2R2M0.SIGYCOMP`

22. System storage name for temporary data sets

- This is the storage device name used for temporary datasets-used in sorting, passing temporary datasets, and so forth-that are deleted after the job completes.
- Suggested Default: SYSTEMP
- Sample value: `SYSTEMP`

23. System storage name for permanent data sets

- This is the storage device name used for permanent datasets used in dataset allocation, such as those used to store COBOL and SQR files.
- Suggested Default: SYSPERM
- Sample value: `SYSPERM`

24. Assembler system data set containing STIMER

- Sample value: `SYS1.MACLIB`

25. Assembler program name

- Sample value: `ASMA90`

26. SQR high level qualifier

- This is the high-level qualifier used for SQR datasets.Suggested Default: PS.HR840.SQR
- Sample value: `PS.HR840.SQR`

27. SQR program name found in SQR load library

- This is the name of the SQR program contained in the SQR Load Library.
- Suggested Default: SQR
- Sample value: `SQR`

28. SQR plan name

- This is the DB2 Plan name assigned for SQR.
- Suggested Default: DBCALLS
- Sample value: `SQR840`

29. Target server hardware platform

- Sample value: Unix System Services (OS390)

30. PeopleSoft Unix System Services home directory

- Suggested Default: /u/data001/dbname
- Sample value: `/u/data001/PSHR800`

31. Library for DB2 CLI load module (i.e., DSNAOCLI)

- Sample value: `DSN610.SDSNLOAD`

32. Plan name for CLI packages (i.e., DSNACLI)

- Suggested Default: DSNAOCLI
- Sample value: `DSNAOCLI`

33. Attachment type for ODBC to connect to DB2

- Suggested Default: CAF
- Sample value: `CAF` (Call Attach Facility)

34. HFS path to top level of JDK product

- This will provide the value for the JDK_HOME environment variable in the psconfig.sh file.
- Suggested Default: /usr/lpp/java/IBM/J1.3
- Sample value: `/usr/lpp/java/J1.1`

4

Installing and configuring DB2 Connect

This chapter describes the process used to set up the connection between the client and the database or between the application server and the database using DB2 Connect V7.2.

DB2 Connect is mandatory to connect to DB2 Distributed Data Facility (DDF) components on the mainframe.

The two main functions covered by DB2 Connect are:

- Performing EBCDIC/ASCII character conversion. Remember, PeopleSoft is storing data in EBCDIC format, not in ASCII.
- Managing the requests/replies between client servers and DDF via SNA or TCP/IP conversations.

For a more complete description, you can refer to *PeopleSoft 8.4 Installation for DB2 UDB/390*, Appendix F.

4.1 Installing DB2 Connect on the installation server

We installed DB2 Connect V7.2 using the CD-ROM for Enterprise Edition for Windows operating environments. Please verify that you have a complete install of Windows 2000 or Windows NT prior to the installation of DB2 Connect. Some companies choose to license only a portion of the executables from Microsoft. PeopleSoft installation requires a full installation of these operating systems.

4.1.1 Installing DB2 Connect

- Insert the CD-ROM and wait for DB2 to launch as shown in Figure 4-1.

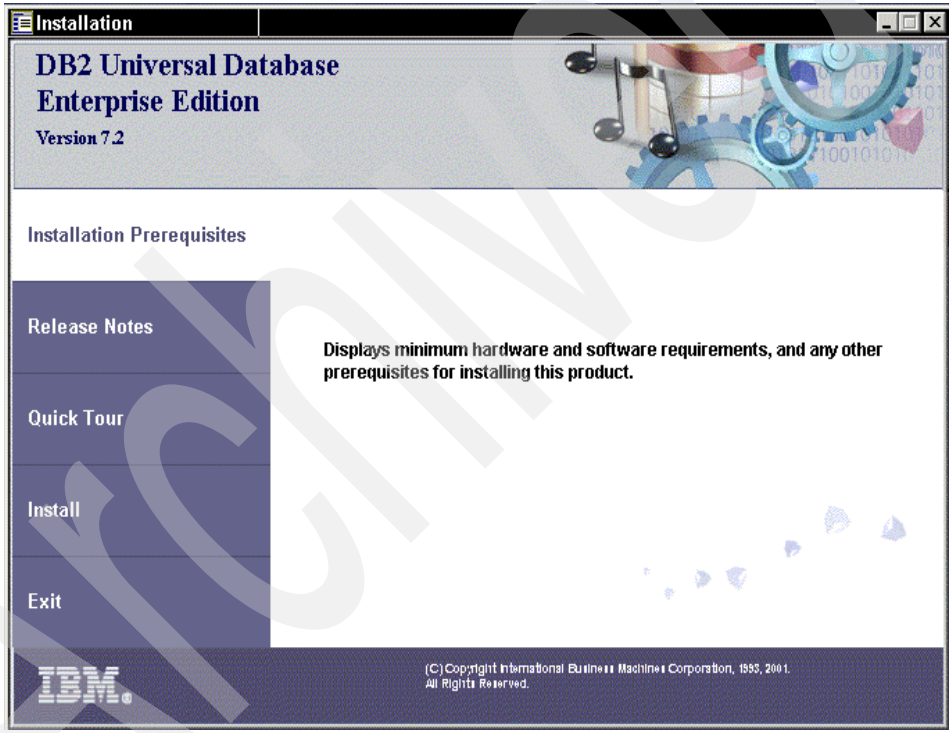

Figure 4-1 DB2 installation window

- Click **Install**.
- We selected **DB2 Connect Enterprise Edition** then clicked **Next** as shown in Figure 4-2.

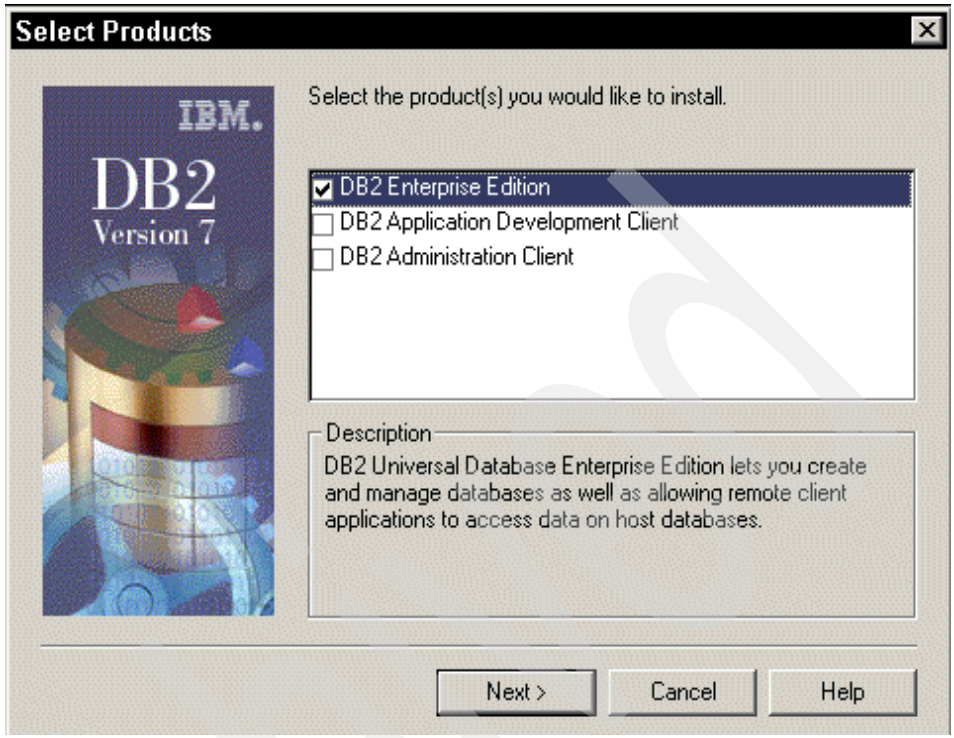

Figure 4-2 DB2 Enterprise Edition selected

- Click **Typical** then **Next** as shown in Figure 4-3.

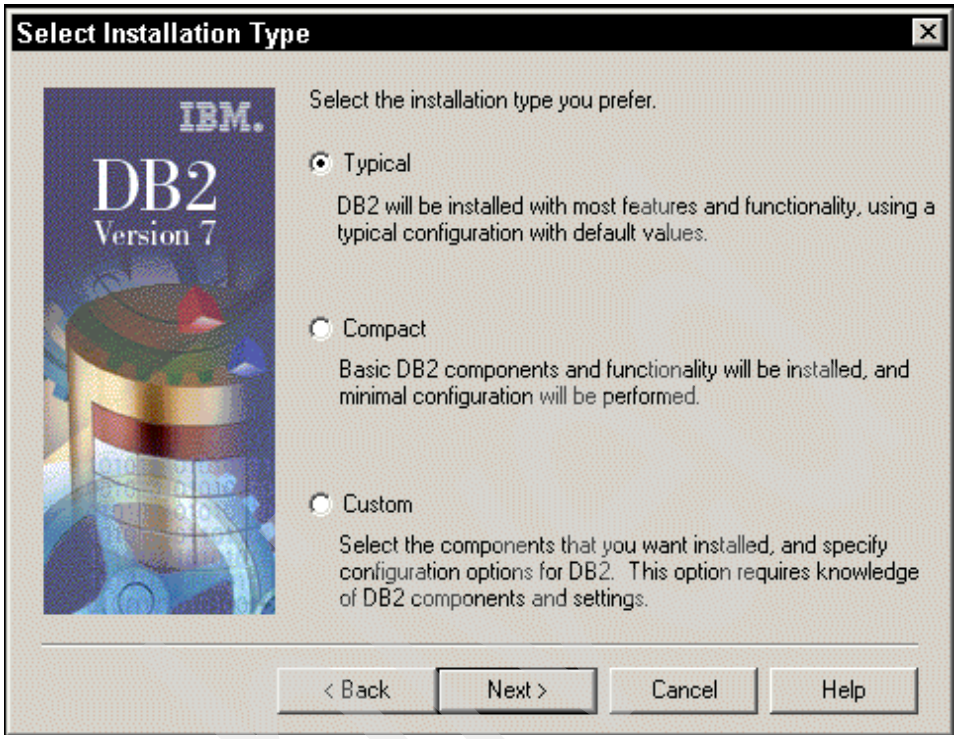

Figure 4-3 Typical selection

- Check that you have the space required and decide which destination folder you want to use, then click **Next** as shown in Figure 4-4.

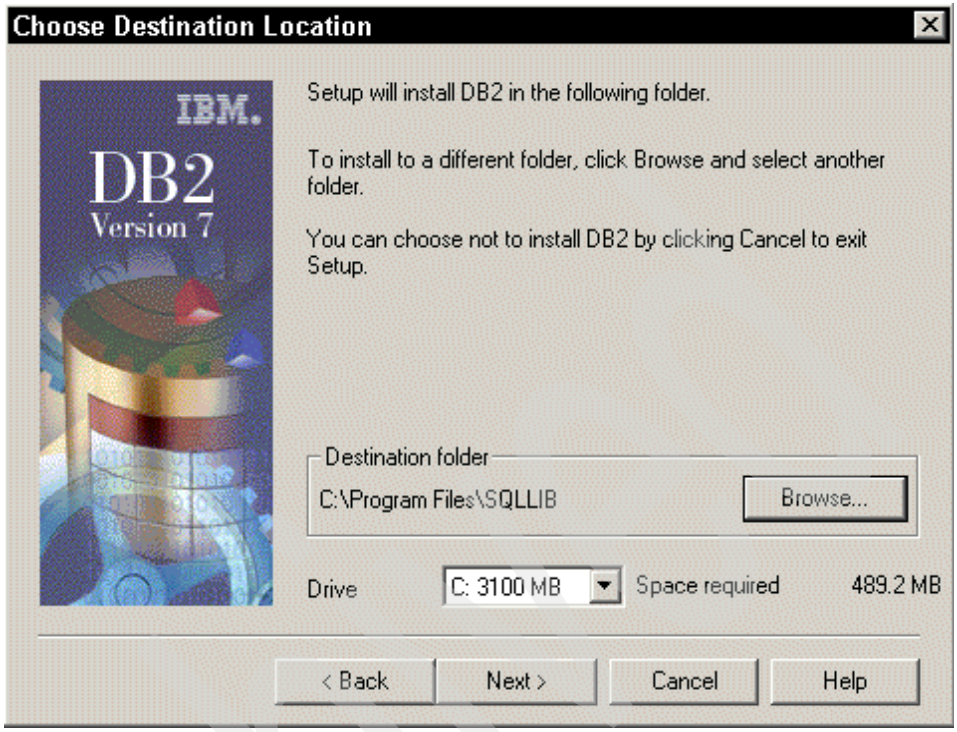

Figure 4-4 Destination location

- Insert the password for DB2Admin then click **Next**. If this is your first installation, you will be prompted in the following step to create it, as shown in Figure 4-5.

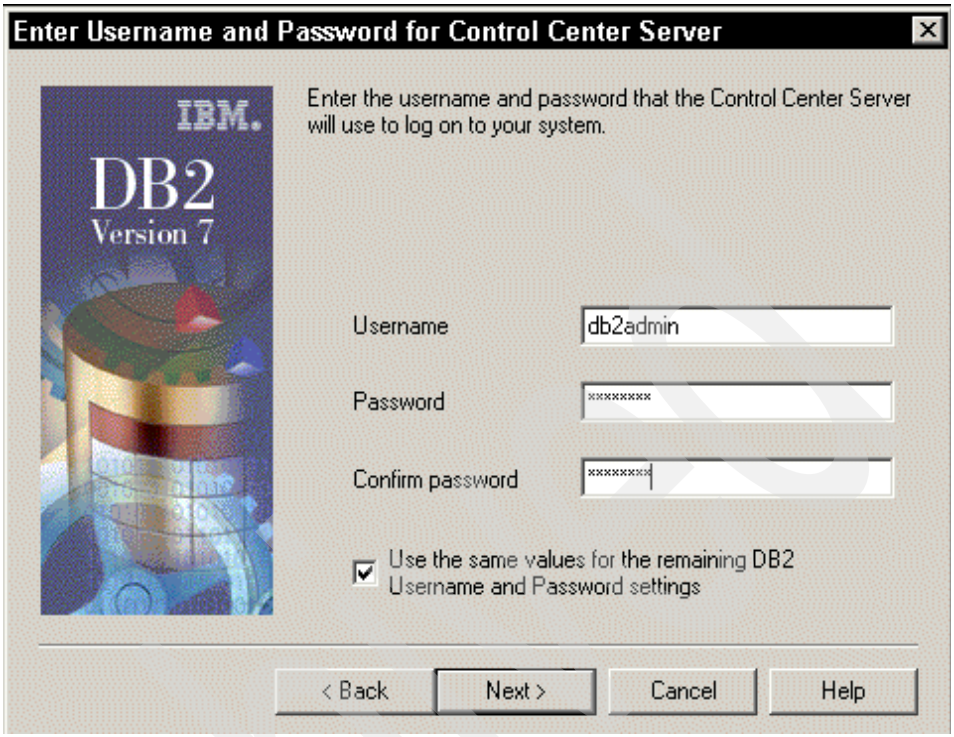

Figure 4-5 DB user name and password

- If this is the first installation, you are asked to create DB2Admin as shown in Figure 4-6. Select **Yes.**

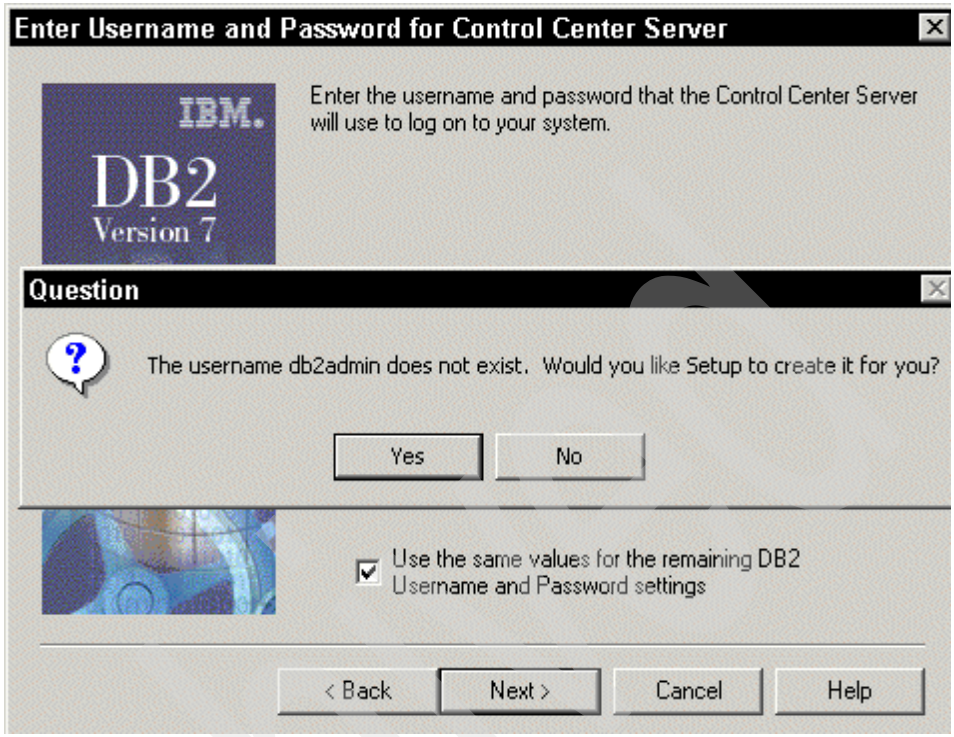

Figure 4-6 User name creation response

- The installation will start when you click **Next** in the following window, as shown in Figure 4-7.

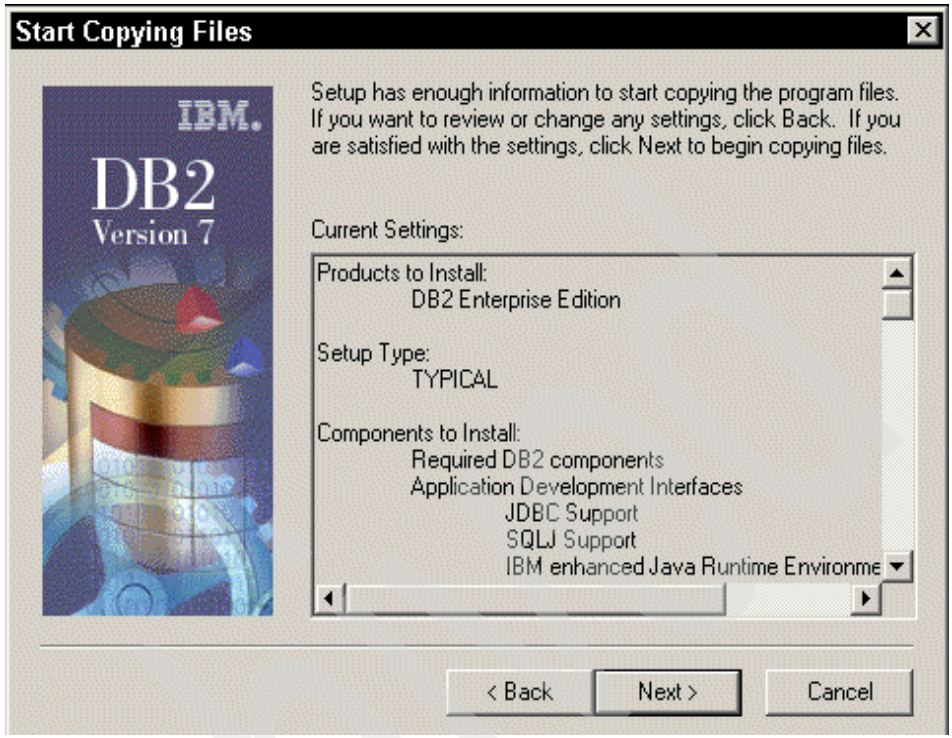

Figure 4-7 Copying files

- We decide not to install the Olap Starter Kit, as shown in Figure 4-8.

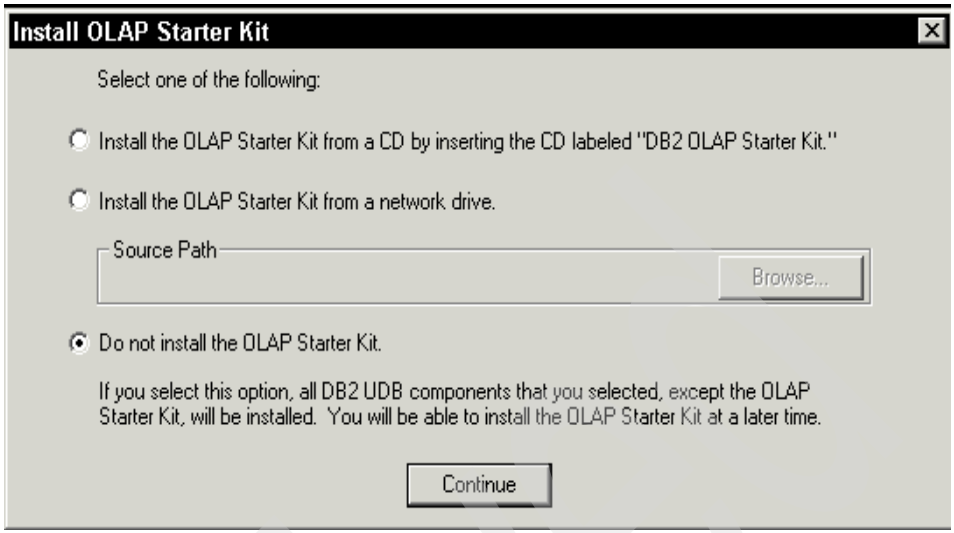

Figure 4-8 OLAP starter kit response

- DB2 Connect is now installed on your server, as shown in Figure 4-9. Click **Finish.**

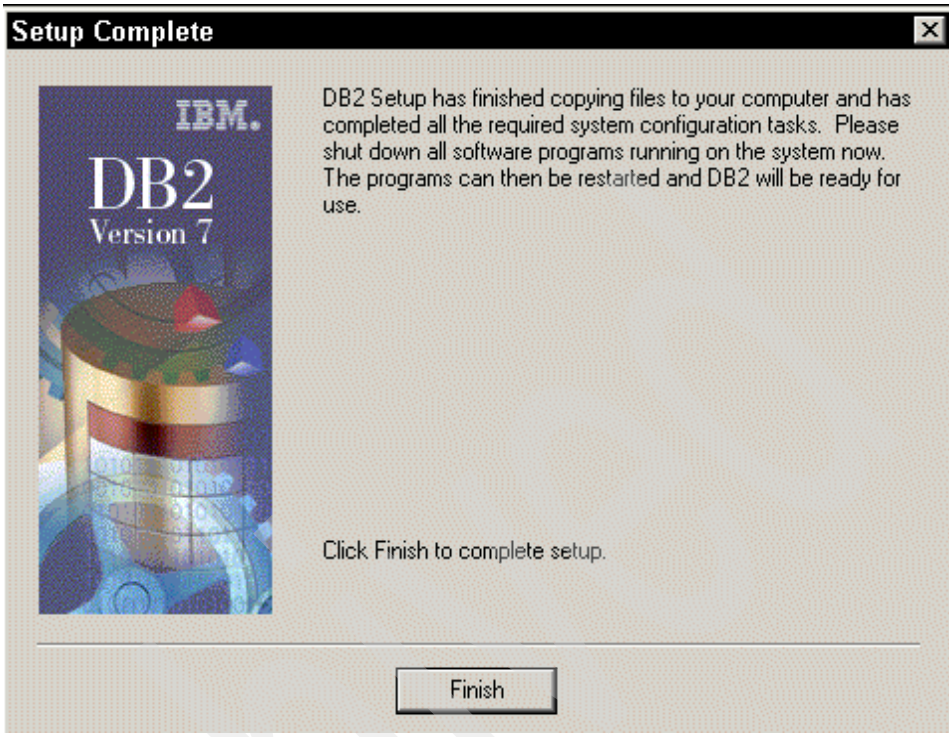

Figure 4-9 Setup complete

- Enter your registration information as shown in Figure 4-10.

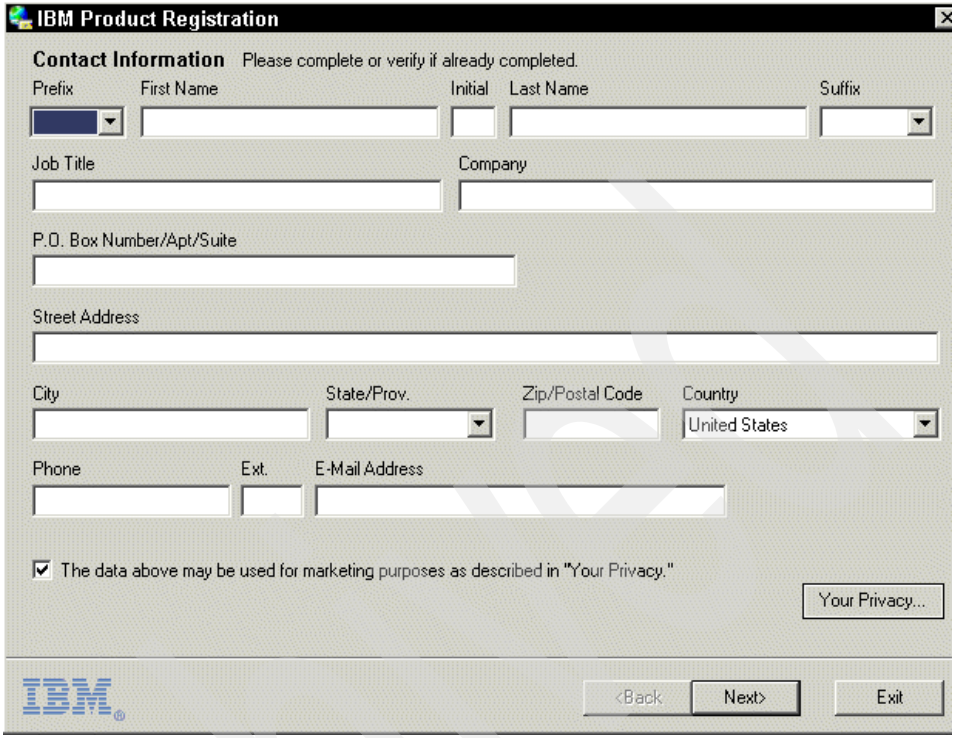

Figure 4-10 Product registration

- Your installation is now complete, as shown in Figure 4-11. Click **Exit**. You now have to configure your DB2 Connect. This is done using the DB2 Client Configuration Assistant.

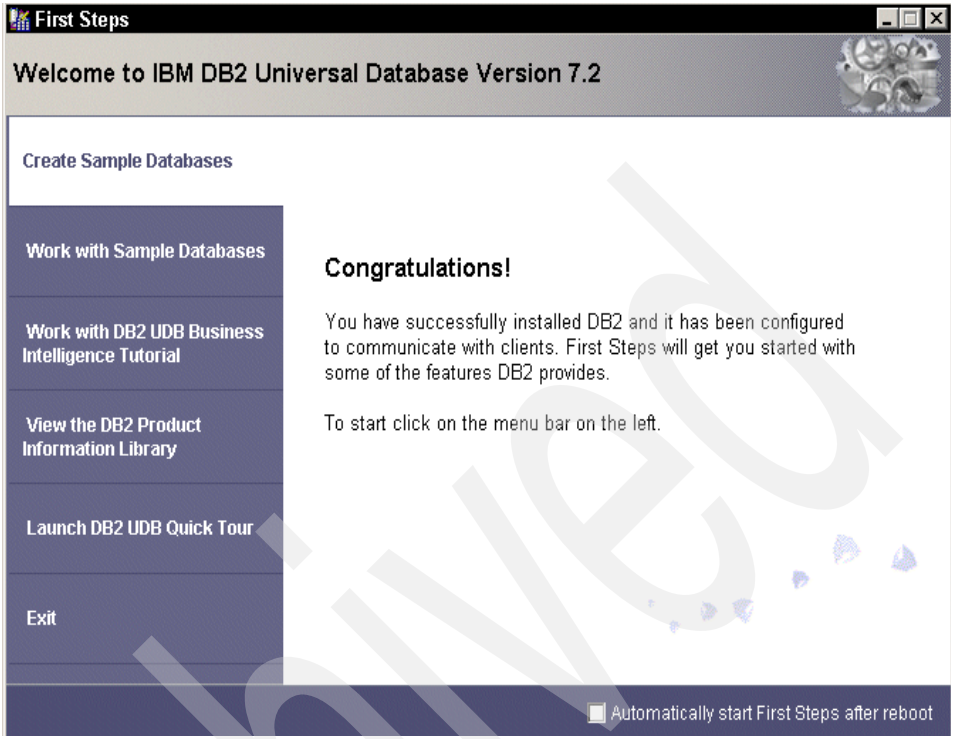

Figure 4-11 Install complete

4.1.2 Configuring DB2 Connect

- Go to **Start-> DB2-> IBM DB2 Client Configuration Assistant** as shown in Figure 4-12.

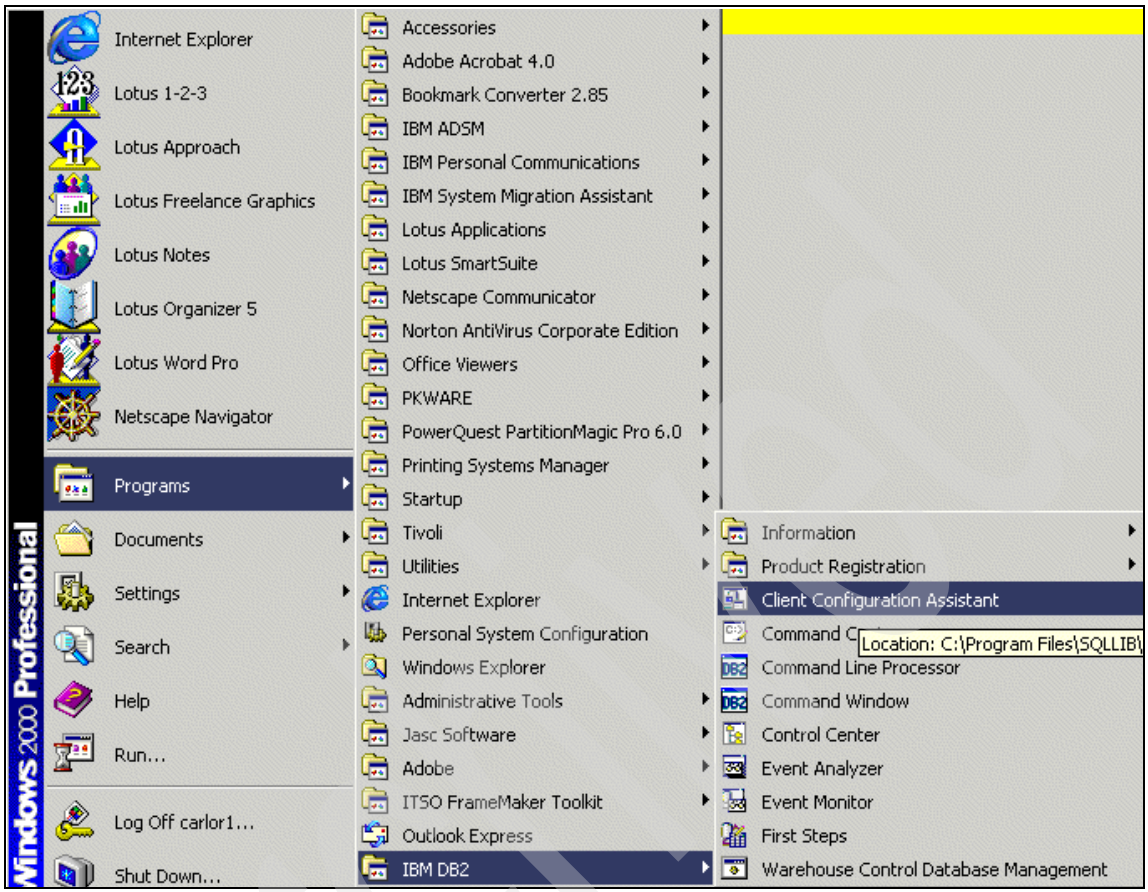

Figure 4-12 Starting DB2 Client Configuration Assistant

- Select **Add a Database**. Some databases are created by default. You can delete them prior to installing yours, if you prefer, as shown in Figure 4-13.

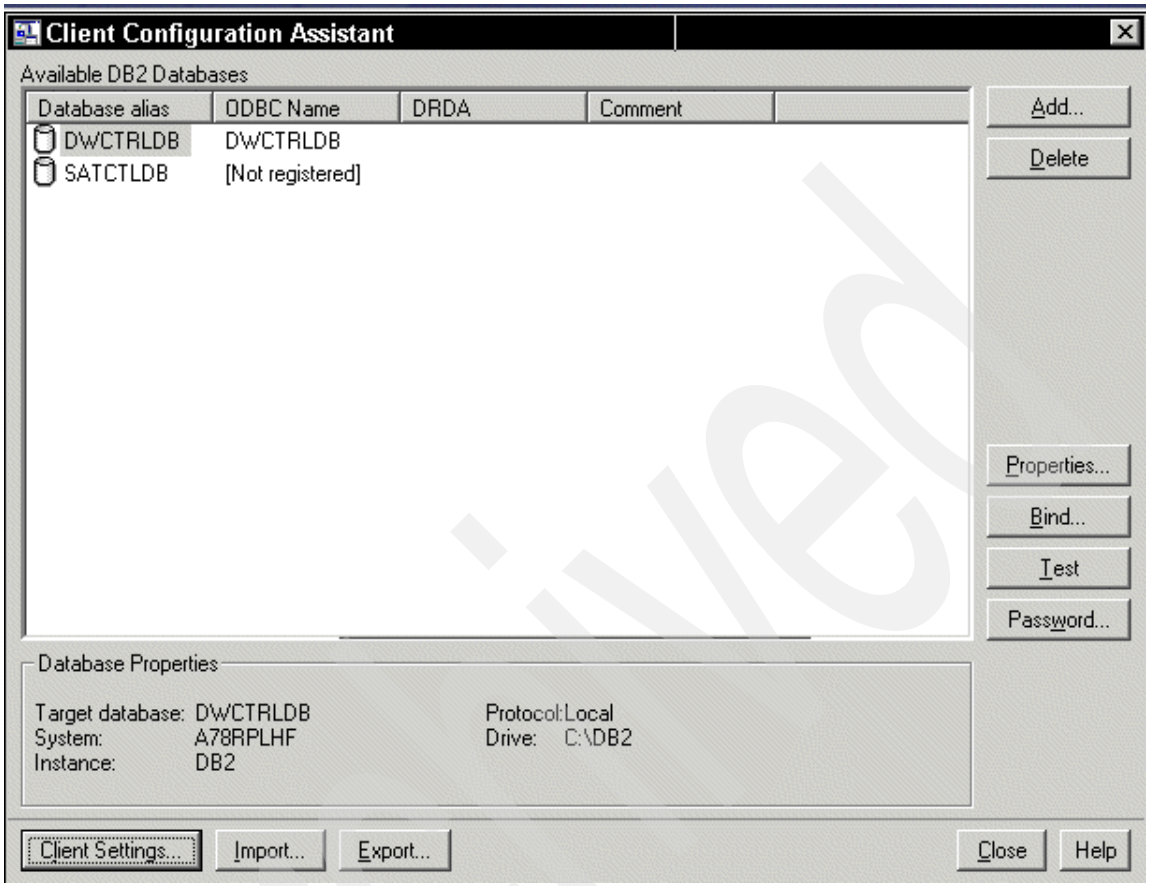

Figure 4-13 Adding a database

- Select **Manually Configure a connection to a DB2 database** then click **Next** as shown in Figure 4-14.

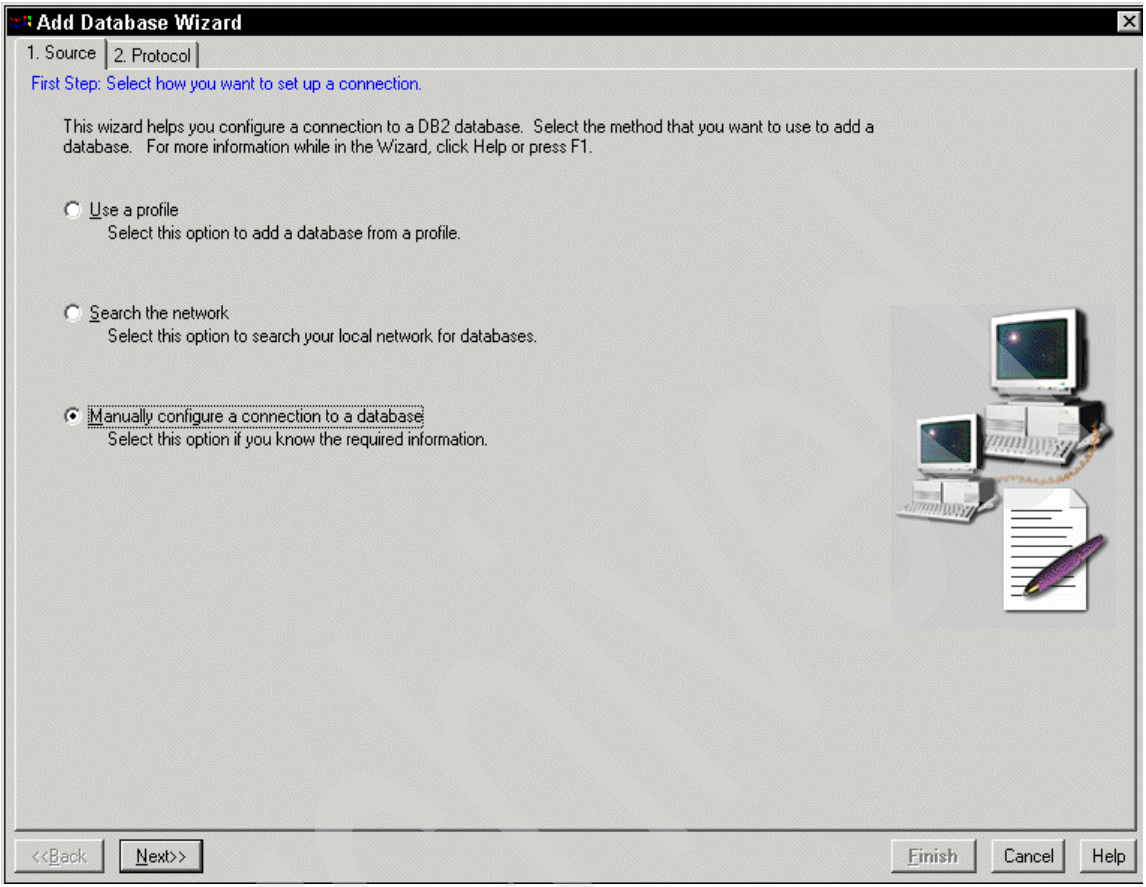

Figure 4-14 Manually configure a database

- Select **TCP/IP** (or **APPC** if you have plan to install via SNA. This has not been our case) then click **The database physically resides on a host or AS/400 system.** Keep the default **Connect directly to the server**. Then click **Next** as shown in Figure 4-15.

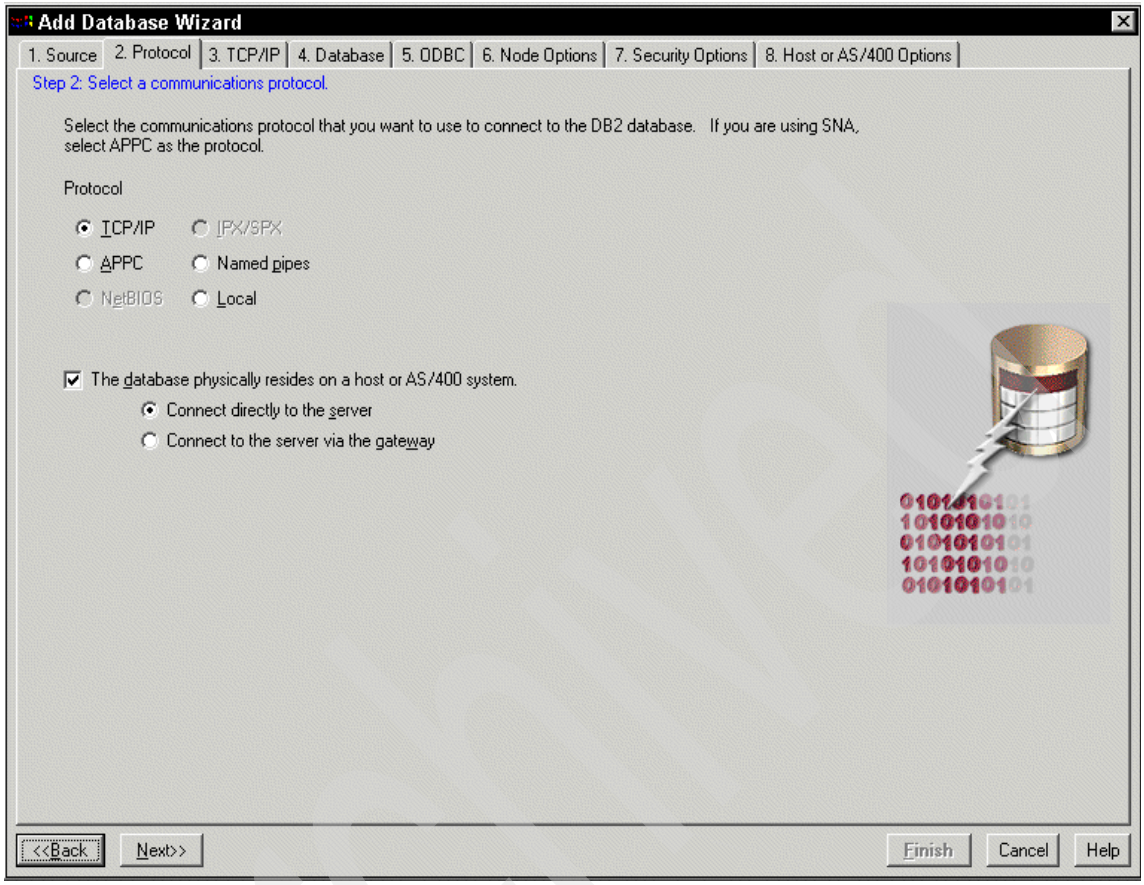

Figure 4-15 Adding database

- Enter your DNS System Host Name or its TCP/IP address in the Host name field.
- Enter the TCP/IP port number used by DDF in the Port number field as shown in Figure 4-16.

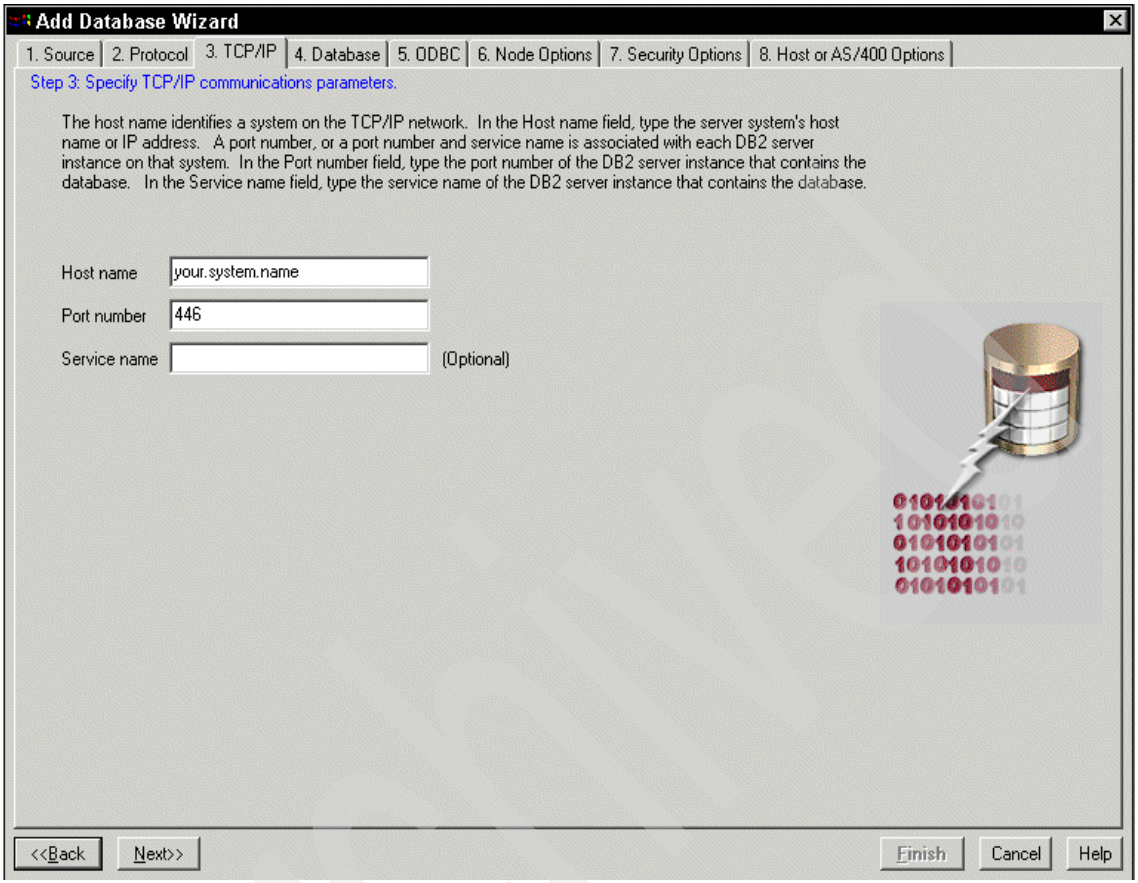

Figure 4-16 Database TCP/IP configuration

- Enter your database name and the alias as defined in your pre-installation checklist then click **Next** as shown in Figure 4-17.

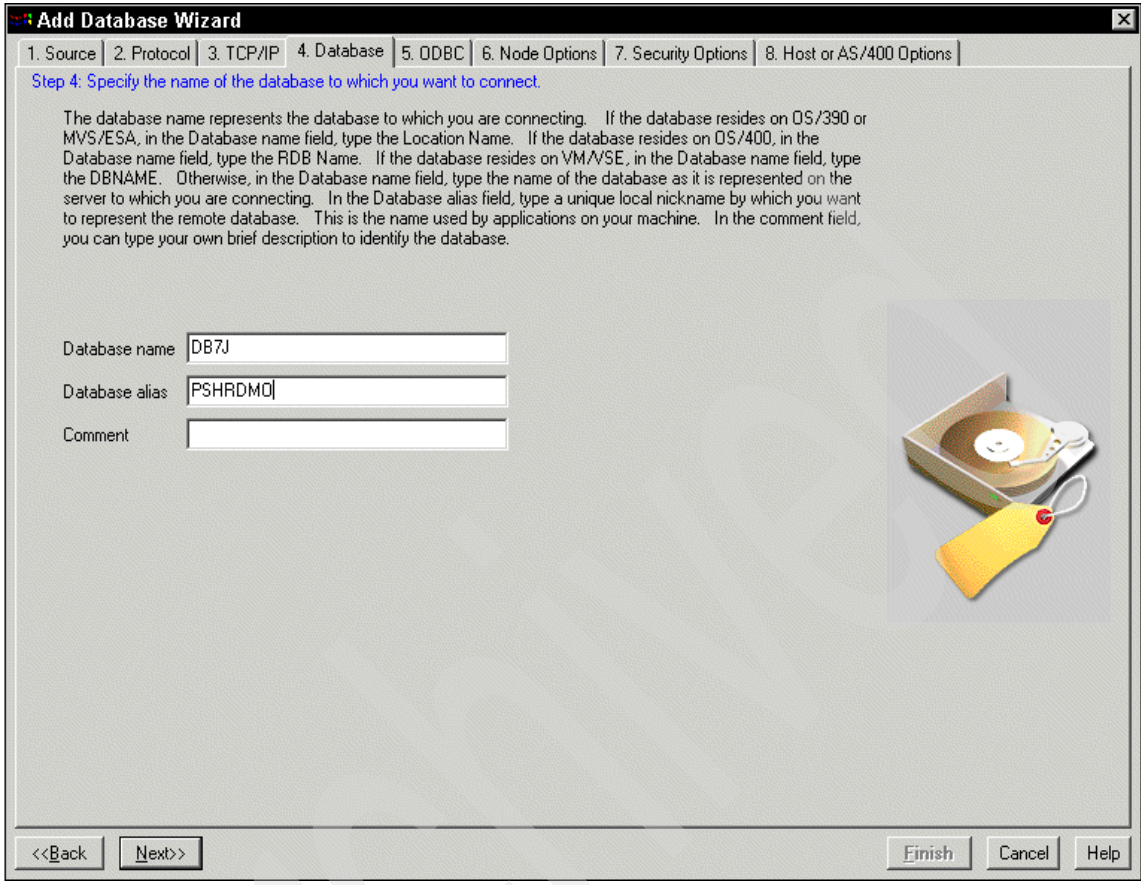

Figure 4-17 Database naming

- Select **Register this database for ODBC** with **As a system data source** then click **Next** as shown in Figure 4-18.

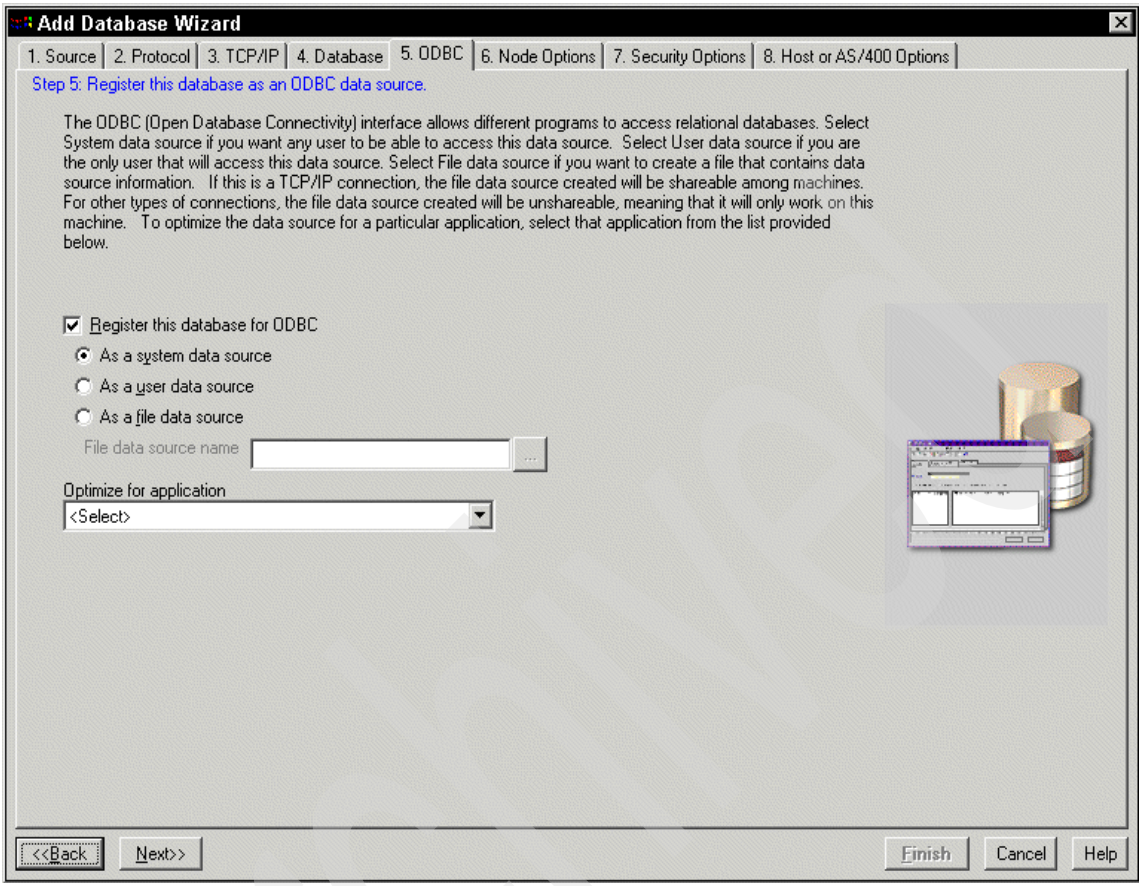

Figure 4-18 Database registering

- The Specify the node options window shows you that **MVS/ESA,OS/390 Operating System** is selected by default as the database is declared on Host. Click **Next** as shown in Figure 4-19.

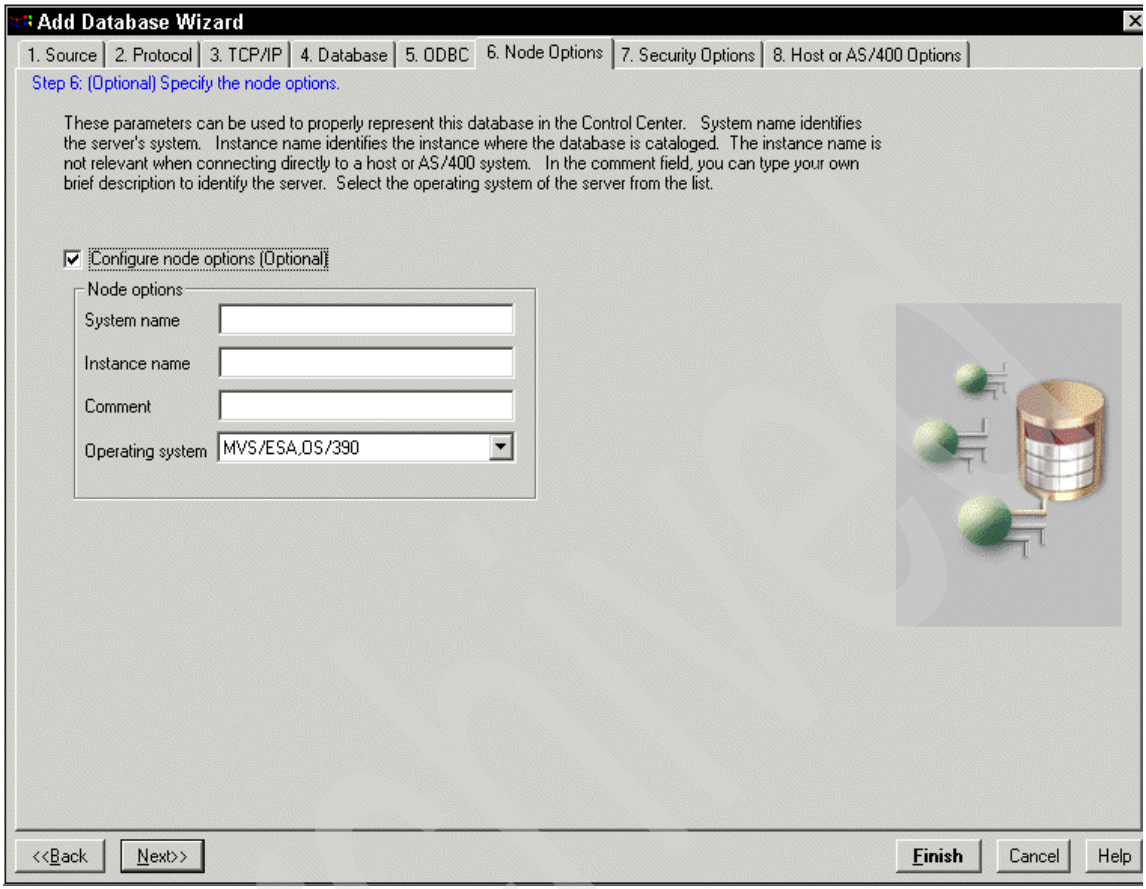

Figure 4-19 Database node options

- In the next window, you will configure security options. We kept the default which is **Host or AS/400 authentication (DCS)** due to previous selection, as shown in Figure 4-20.

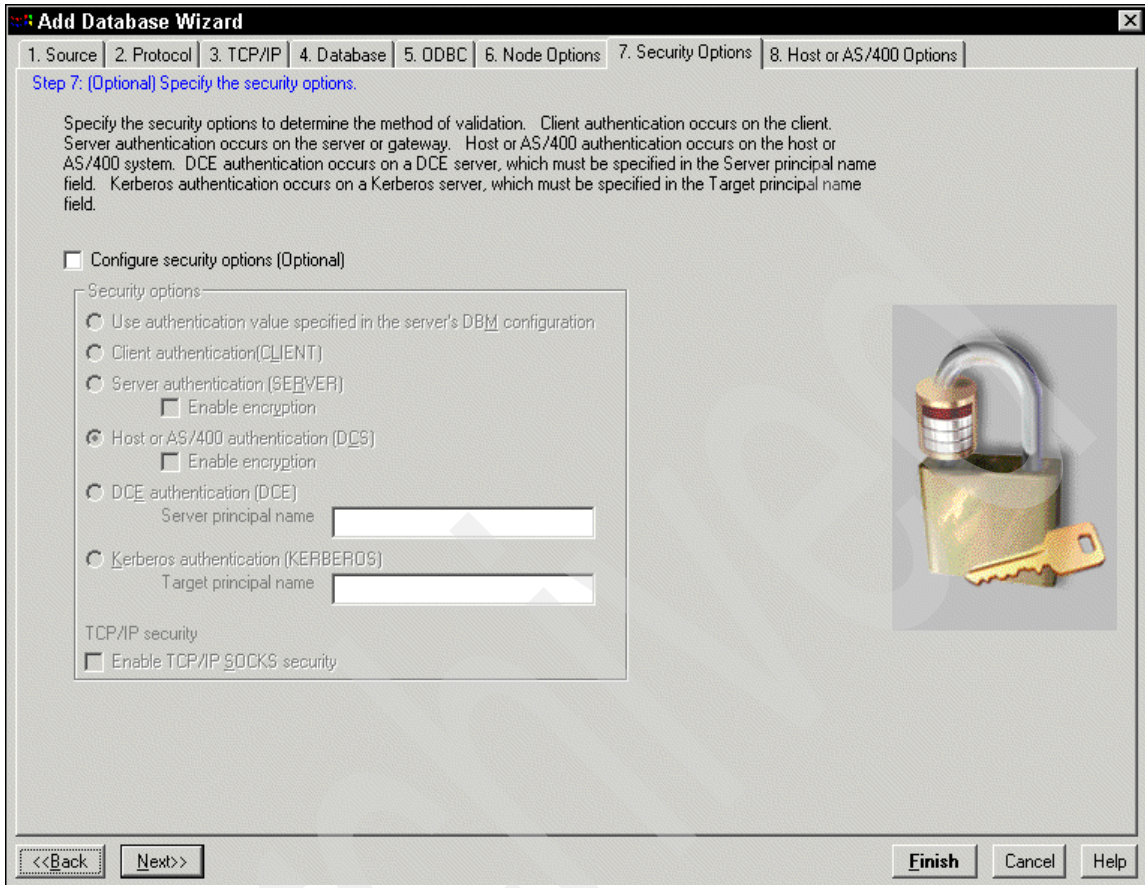

Figure 4-20 Database security options

- Click **Finish** as in Figure 4-20. As you have selected **The database physically resides on a Host or an AS/400**, the options are not relevant (see Figure 4-21).

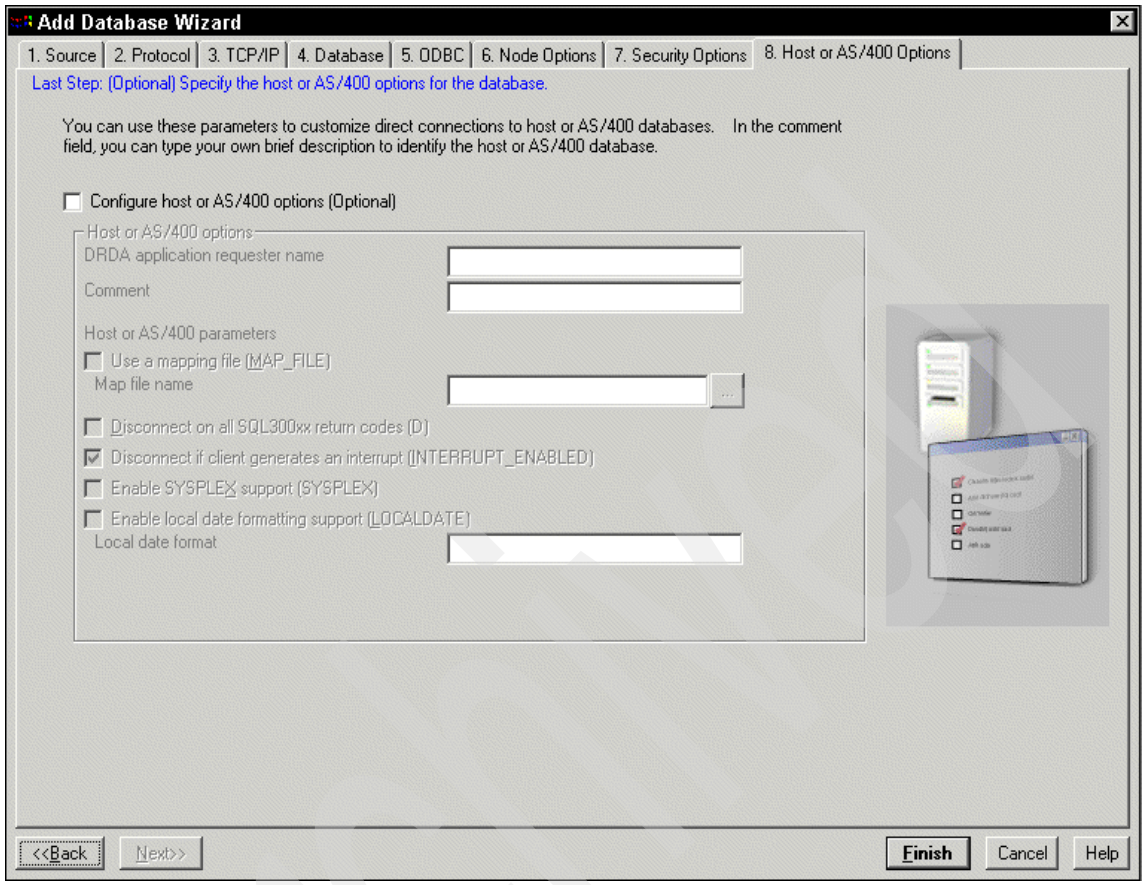

Figure 4-21 Optional AS/400 window

- You receive the confirmation message stating that your database has been successfully created. Click **Test Connection** as shown in Figure 4-22.

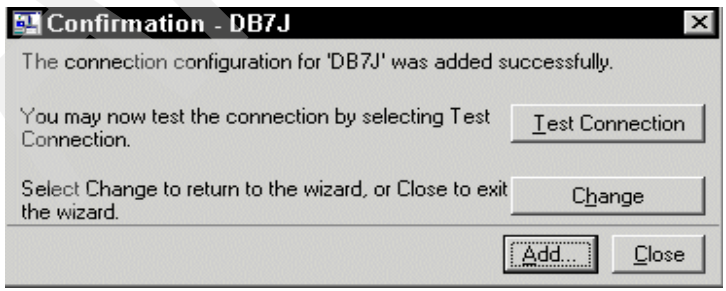

Figure 4-22 Database connection confirmation

- Enter an User ID and Password to test the connection. Keep the default **Share** then click **OK** as shown in Figure 4-23.

Figure 4-23 Connect to database

- The Client Configuration Assistant will confirm that the Connection has been successfully tested, as shown in Figure 4-24.

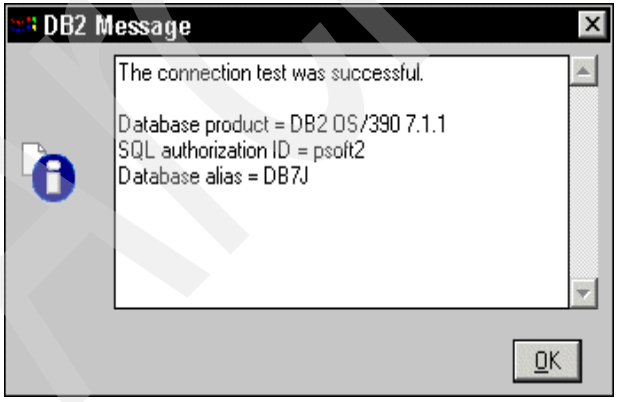

Figure 4-24 Success confirmation

- This completes the installation and configuration of DB2 Connect on the installation server.

4.2 Installing DB2 Connect on the application server

You must also install DB2 Connect on your PeopleSoft application servers. We used V7.2. Please repeat the above directions for installation on each of the application servers. Below, we outline some hints and tips for traces and performance for DB2 Connect on the application servers.

4.2.1 Setting logs to capture dynamic SQLs

If you have troubleshooting issues later on, you might decide to capture the DB2 Connect activity to check that everything is working normally at this level. You can do this using the following process:

1. Go to Client Configuration Assistant.
2. Click **Properties.**
3. Click **Settings**.
4. Respond **Yes** to `Would you like to connect to the data source?`
5. Enter your User ID and password.
6. When finished, click **OK.**
7. Click **Advanced...** as shown in Figure 4-25.

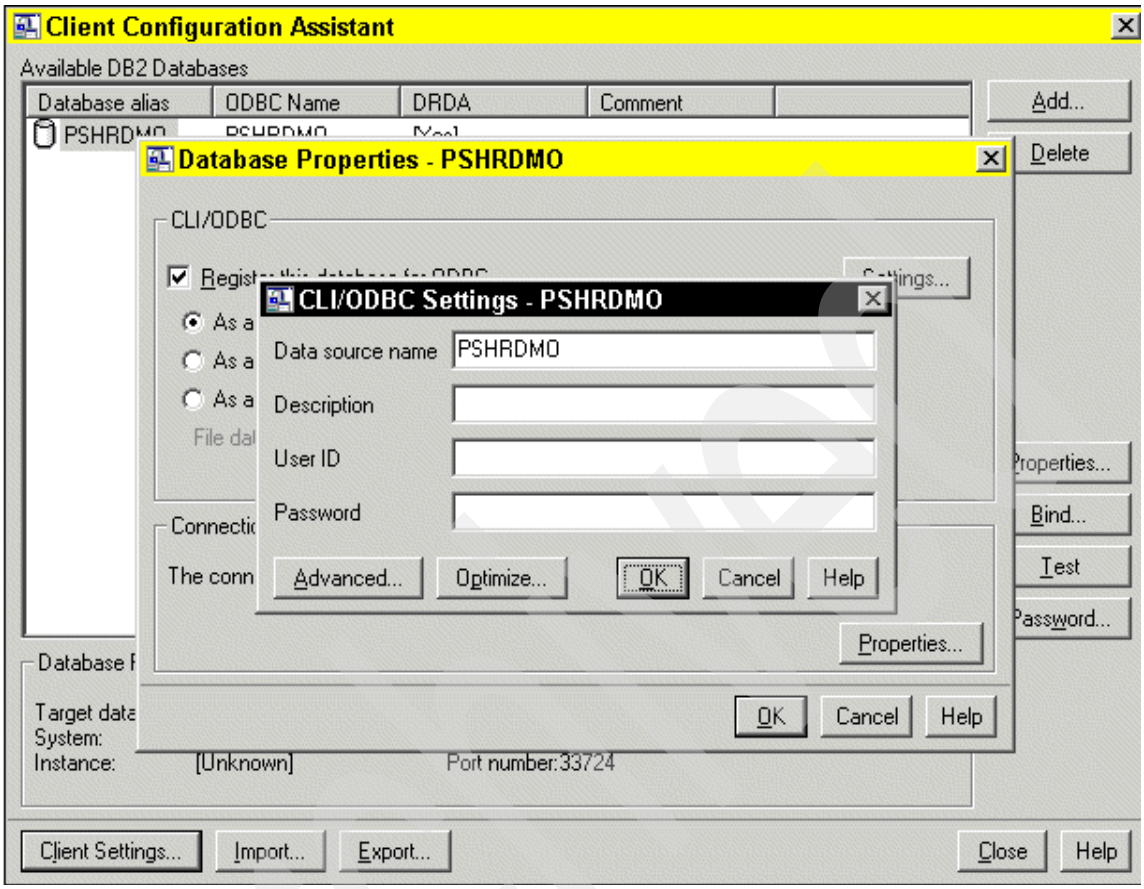

Figure 4-25 Database properties

8. You will see the Advanced CLI/ODBC settings (see Figure 4-26).

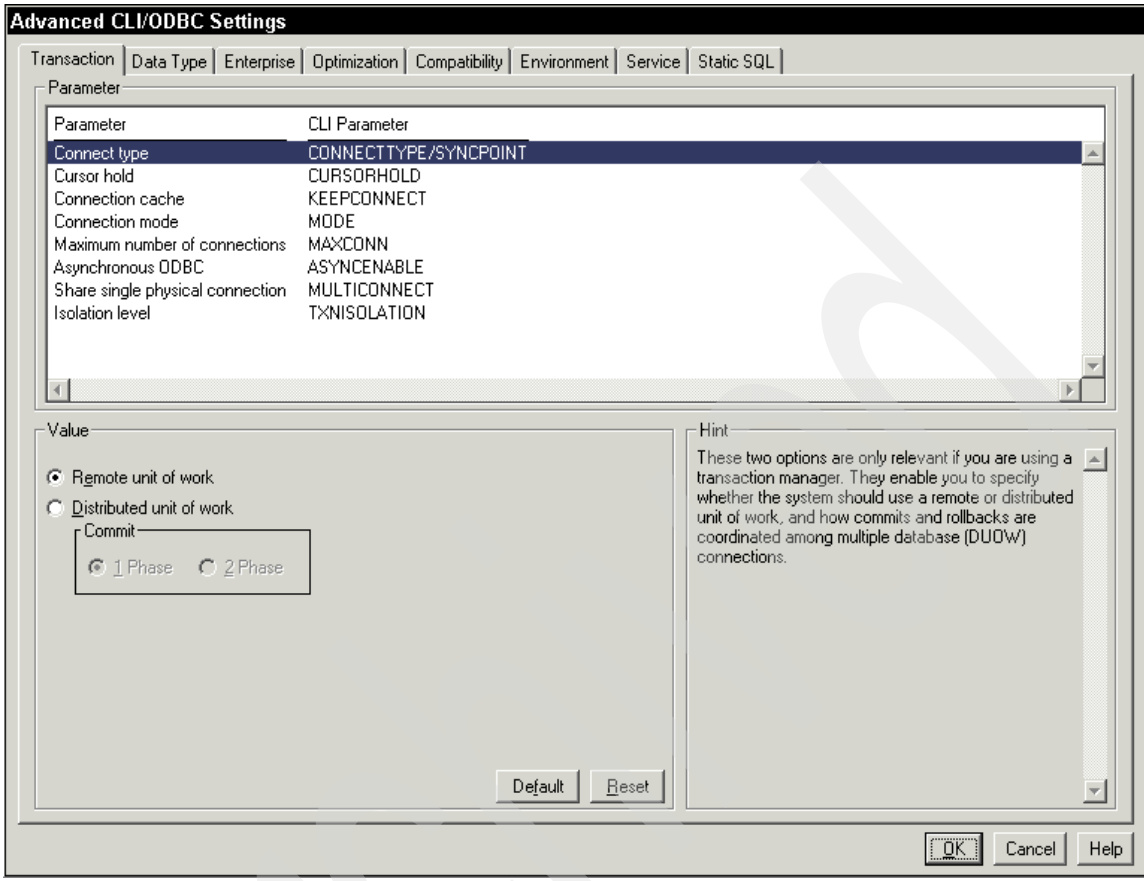

Figure 4-26 Advanced CLI/ODBC settings

9. Go to Service.
10. Select **Trace(Common).**
11. Select **Enable trace** and enter your file name and folder destination choices.

    Each SQL executed using DB2 Connect will be logged in this file, as shown in Figure 4-27.

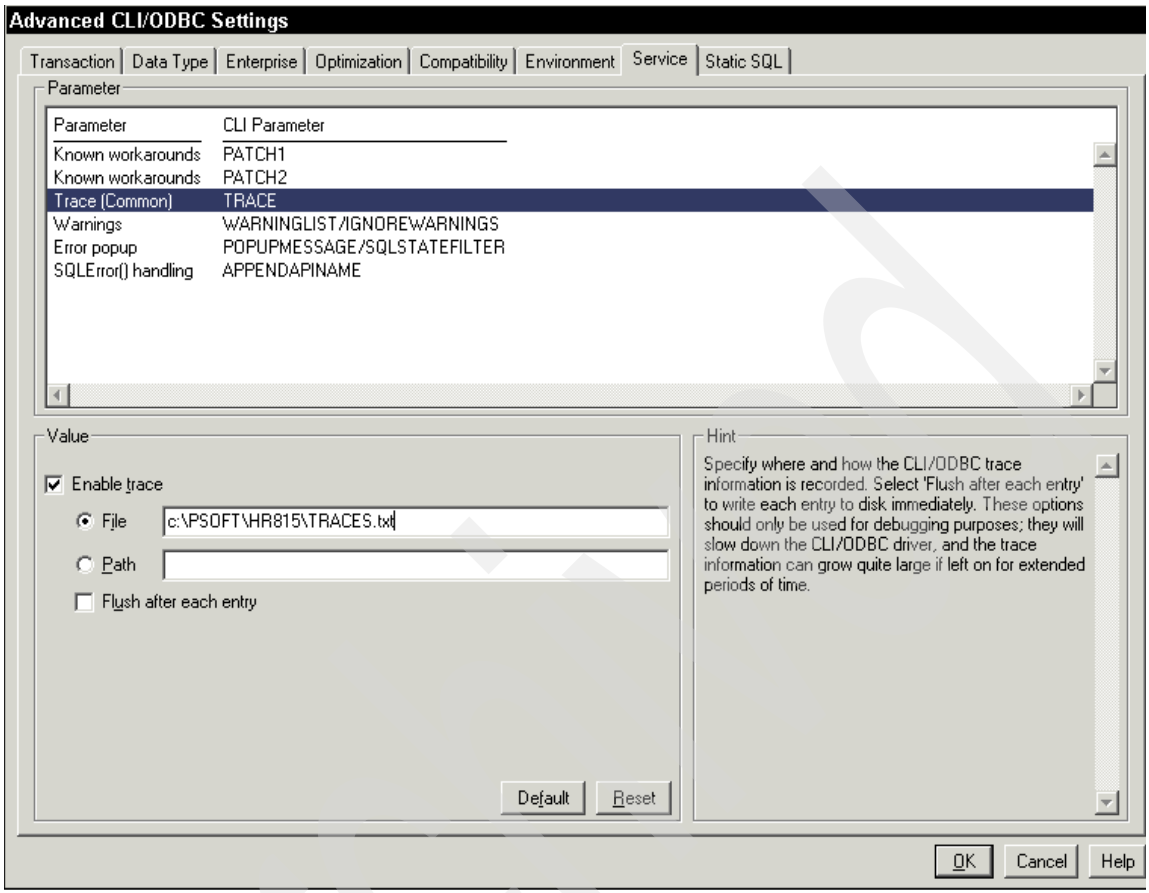

Figure 4-27 Tracing

4.2.2 DB2 Connect performance recommendations

Here we describe the changes made to the DB2 Connect default options. Indeed, PeopleSoft and IBM recommend the following changes for PeopleSoft:

- `DeferredPrepare = 1` in DB2CLI.INI for Performance Boost
- `CursorHold = 1` to allow Persistent Cursors. This is now the default, but check if you are upgrading for the V7.5 or before PeopleSoft release and from DB2 Connect V5.2.
- `Autocommit = 0`. The default used to be Autocommit=1 so an override was required. This is not the case anymore with DB2 Connect V7.2.
- `Rqrioblk=32767` In Client Setting. This is the default with DB2 Connect V7.2 can could be checked.

a. To check, in the first window of Client Configuration Assistant, click **Client Settings...**.

b. Go to Performance and look at the Client I/O block size

- `Disableunicode=1`
- `Disablekeysetcursor=1`

5

Installing PeopleSoft and third party software

This chapter describes the different steps for installing PeopleSoft and third party products on your file server. We describe our experience while installing the following:

- PeopleSoft Tools and Applications
- Crystal Report
- SQR on z/OS

5.1 Installing PeopleSoft Tools and applications

Once you have established connectivity using DB2 Connect between the file server, the database server, and the installation client, you are ready to start the installation process. The first step is to set up the file server. It is important to maintain your PeopleSoft Code on the file server. This is where you will apply fixes and updates prior to sending the fixes to your other servers.

If you have an installation client with a large disk drive, you can install directly from the client. This saves space on the client.

5.1.1 Creating PeopleSoft directories and install applications

The steps we used were as follows:

1. Step 2-1: create a folder or a logical drive for the PeopleSoft directories.

   We used one server running Windows 2000 as our file server. We created a folder C:\PSOFT\HR8xx. Be sure you have enough space, as it will require at least 2GB of disk for the application and the PeopleTools.

   Create one folder (or Directory Map) per application you plan to install.

2. Step 2-2: map the drive.

   We did not need to map since we installed directly on the file server.

3. Step 2-3: install PeopleSoft PeopleTools CD-ROMs.

   As of PeopleTools 8.40, PeopleTools and PeopleSoft applications are installed with InstallShield Wizard.

To start the install process, insert the PeopleTools CD into the CD drive. Using Windows Explorer to see the contents of the CD, execute the SETUP.EXE file as shown in Figure 5-1.

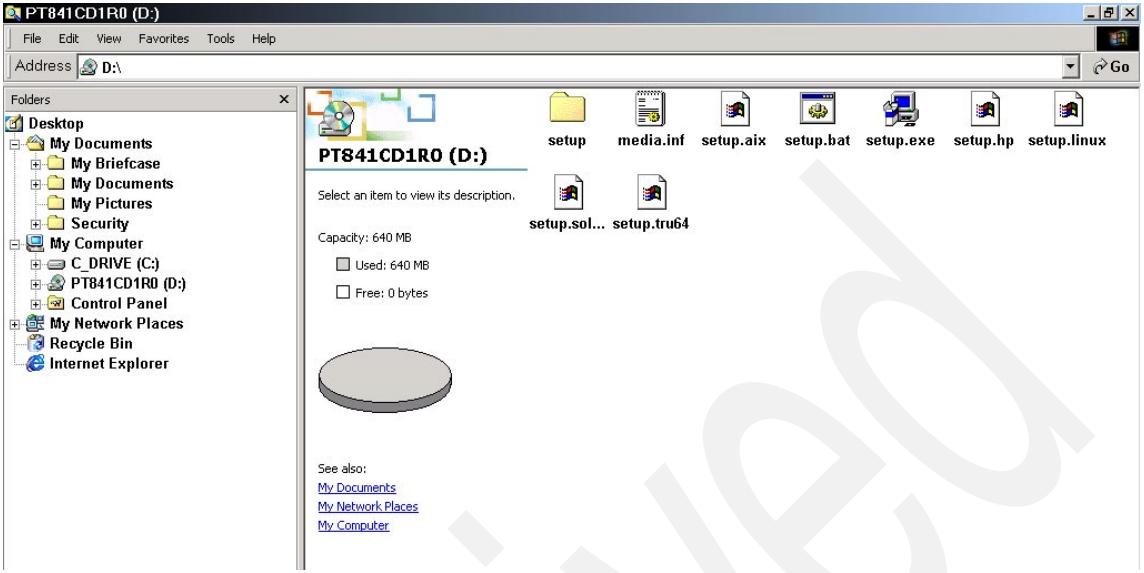

Figure 5-1 Selecting the setup.exe file

After SETUP.EXE runs, you will see the first screen of the Wizard as shown in Figure 5-2.

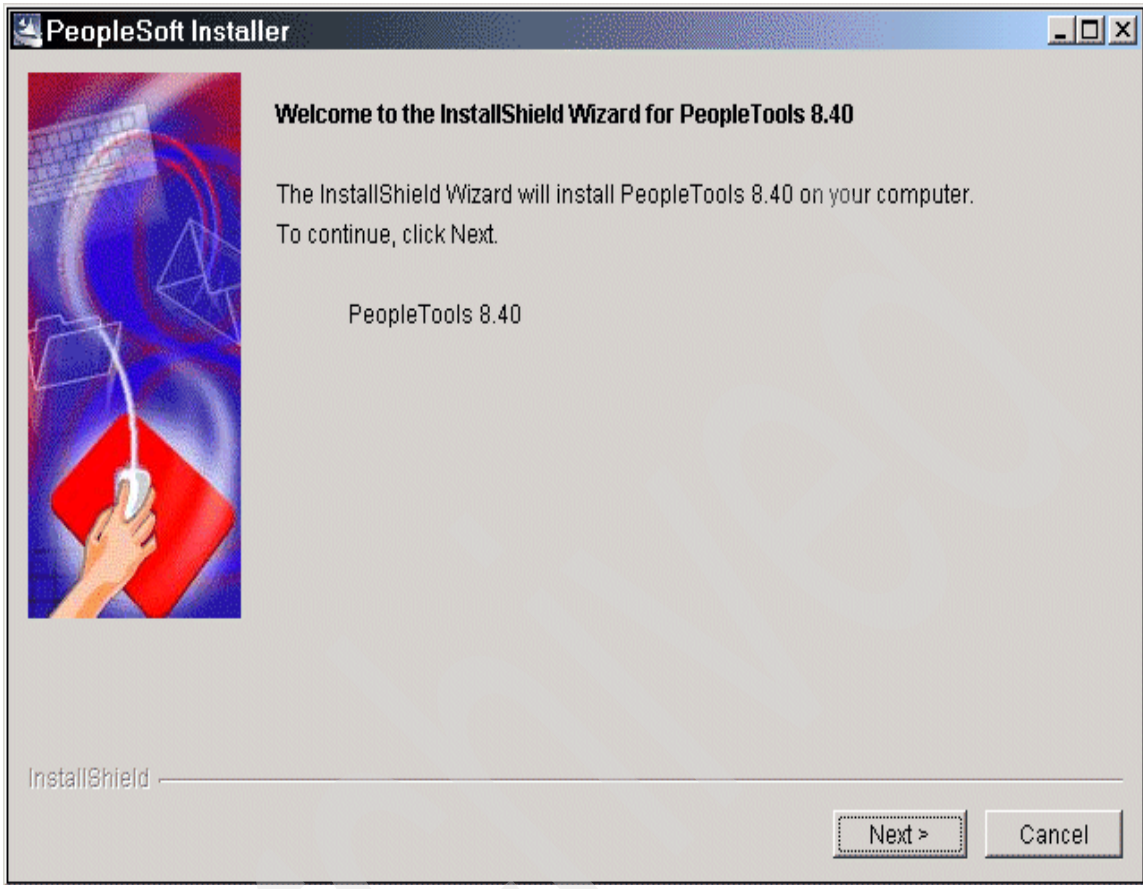

Figure 5-2 Welcome screen for the install wizard

c. Next you will need to read the License Agreement and agree to it as shown in Figure 5-3.

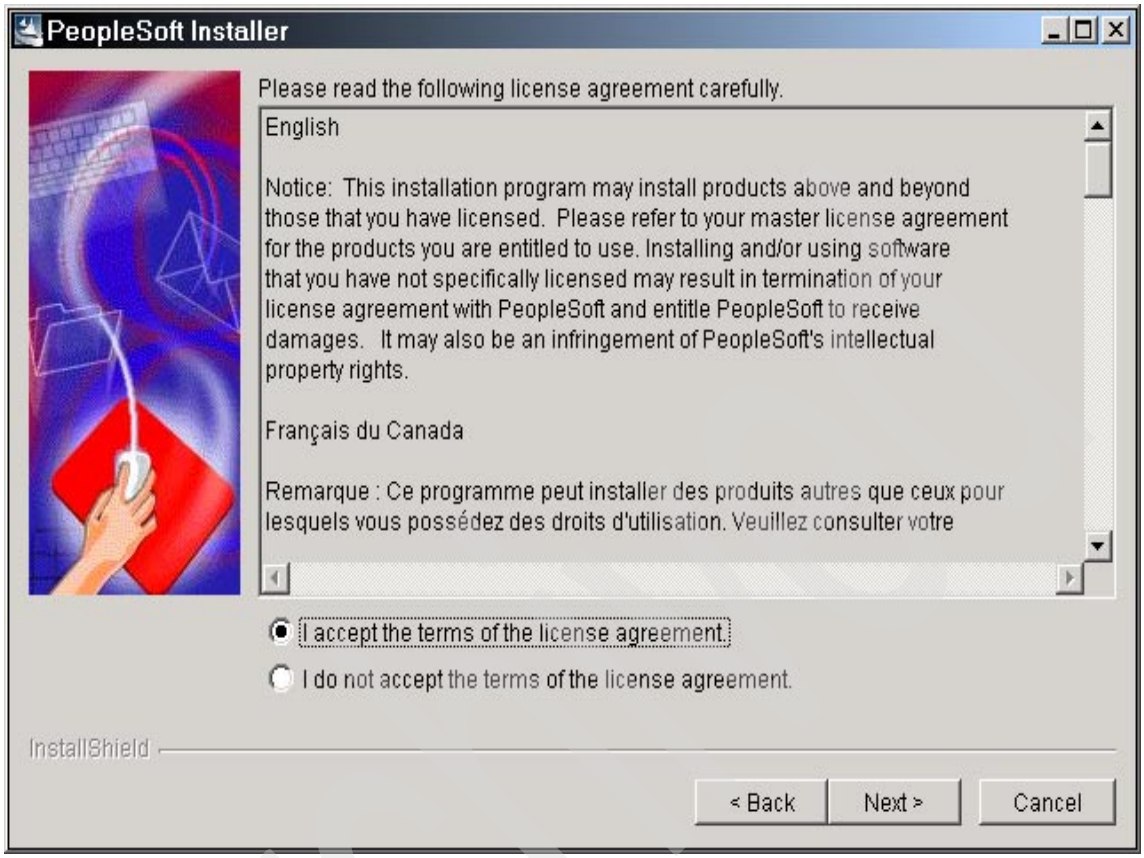

Figure 5-3 Reading the licensing agreement

d. Enter your license code. Make sure the product license code you are entering is the one for z/OS, if you have ordered code for a different platform. Code for all platforms is on the same CD as shown in Figure 5-4.

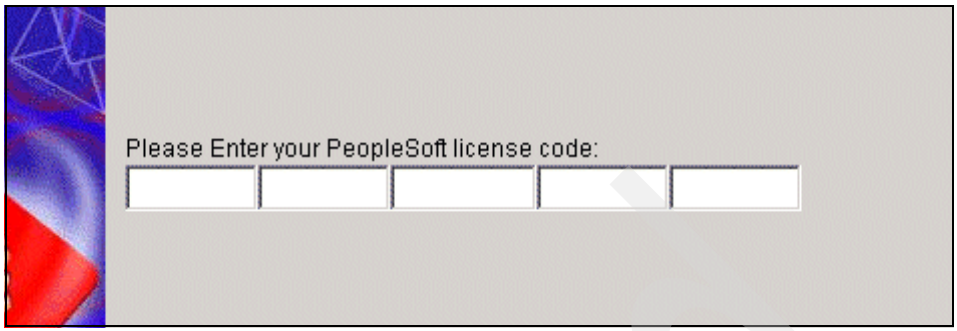

Figure 5-4 Entering your license code

e. Next, the Wizard will allow you to pick the Servers you wish to install, as shown in Figure 5-5.

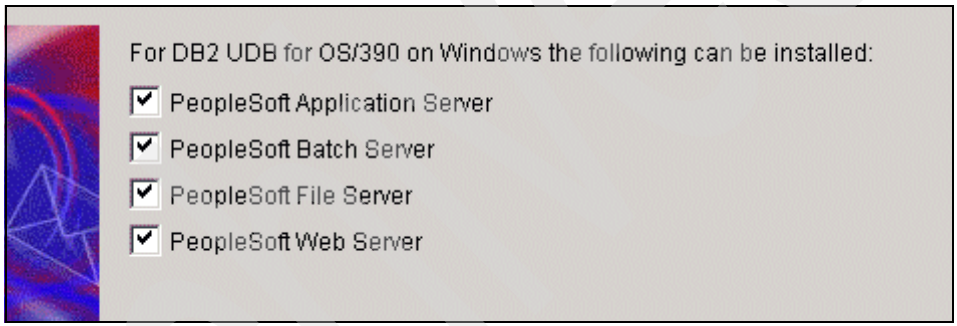

Figure 5-5 Selecting the servers to install

f. Next, install PeopleTools into the directory defined in step 2-1 of 5.1.1, “Creating PeopleSoft directories and install applications” on page 72 and shown in Figure 5-6.

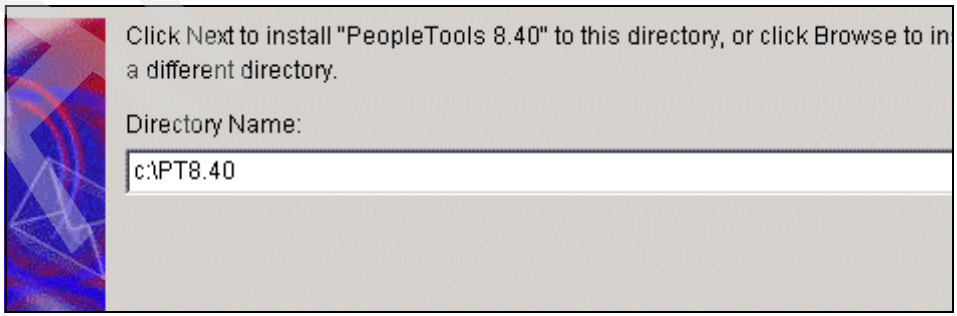

Figure 5-6 Selecting the directory for the code

g. Next, select the location of the Connectivity Program Directory as shown in Figure 5-7.

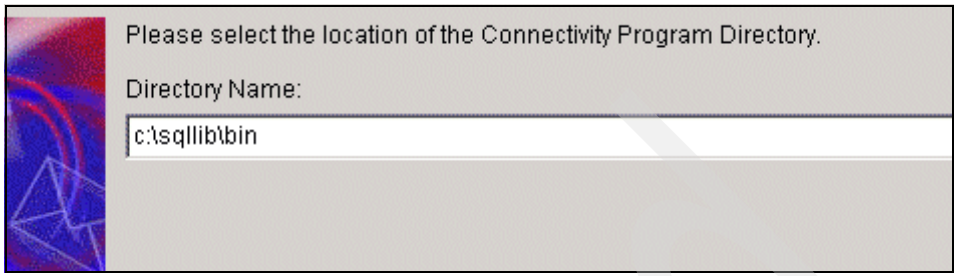

Figure 5-7 The connectivity program directory

h. The next two screens have to do with ICONS as shown in Figure 5-8 and Figure 5-9.

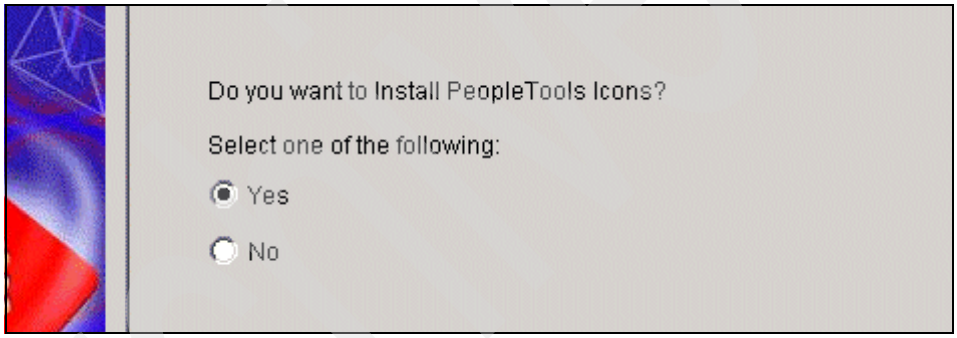

Figure 5-8 Choosing to install icons

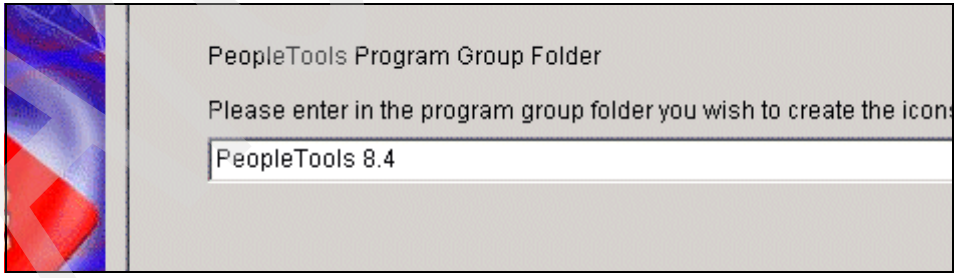

Figure 5-9 Icon group folder

i. Next you will determine the PeopleTools features to install as shown in Figure 5-10.

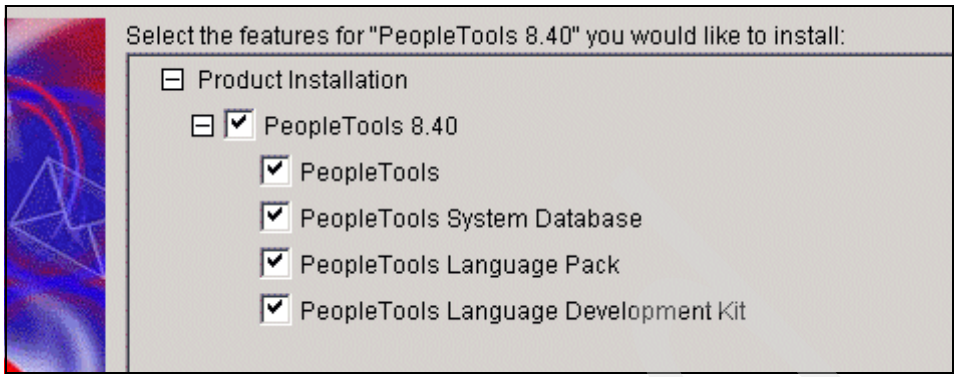

Figure 5-10 Selecting PeopleTools features

j. The next screen will confirm the features to be installed and the directory to be used as shown in Figure 5-11.

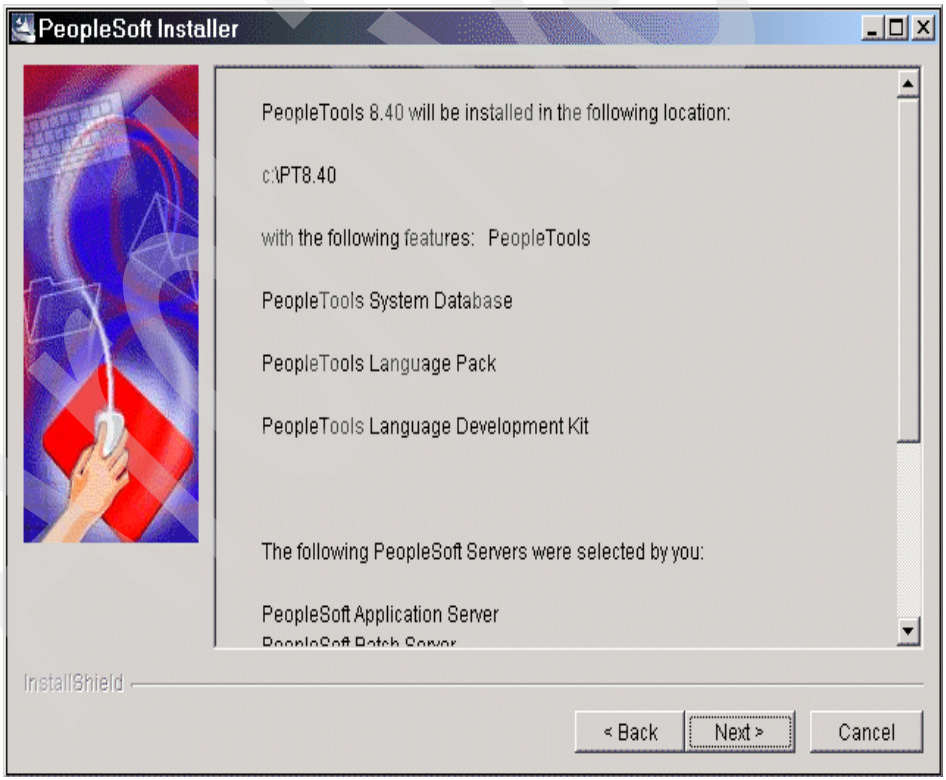

Figure 5-11 Confirming your selection of servers

k. The installation now begins and you can see the progress as shown in Figure 5-12.

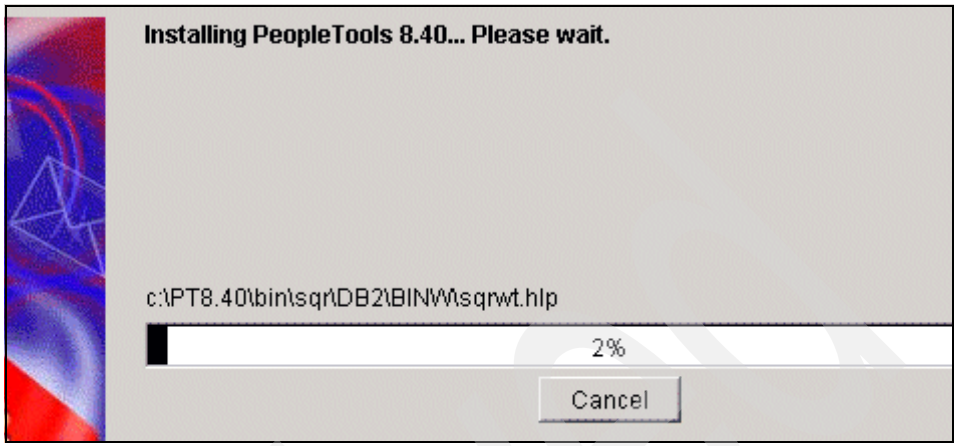

Figure 5-12 Installing PeopleTools -- progress bar

l. PeopleTools 8.4x spans two CDs, so you will be prompted to insert CD #2 as shown in Figure 5-13.

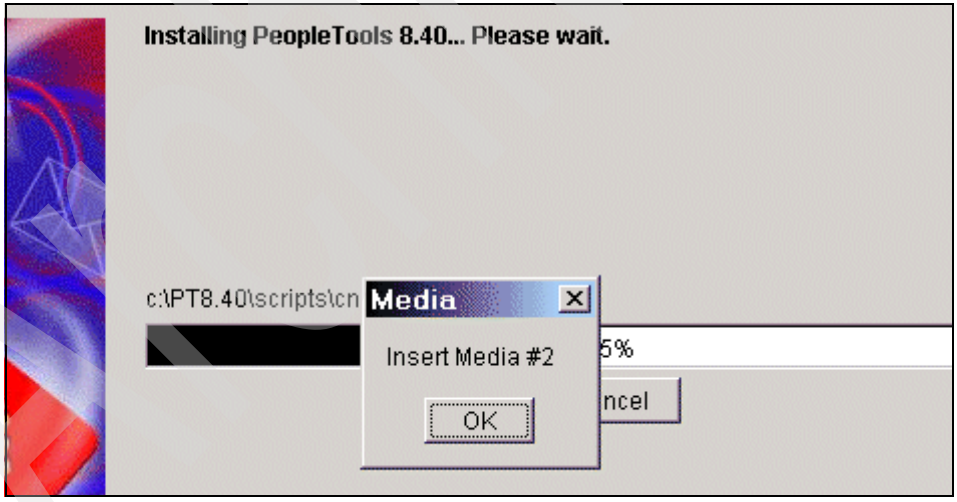

Figure 5-13 Inserting CD-ROM #2

m. Finally, you see confirmation that the Wizard has successfully completed the PeopleTools 8.4x install as shown in Figure 5-14.

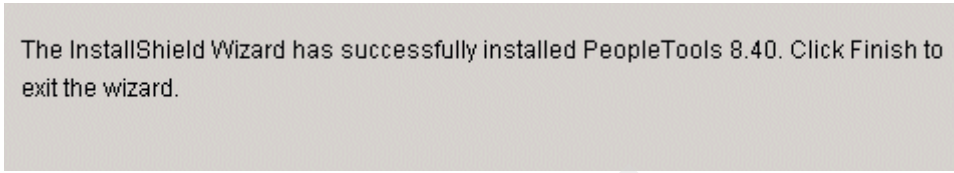

Figure 5-14 Exiting the wizard.

n. PeopleTools is now installed.

To install the applications, perform the same steps. First insert the application CD and with Windows Explorer look at the CD. Execute setup.exe and follow the same steps as with PeopleTools.

Note: One question you will have to answer is whether or not you want to have a Unicode Database. For DB2 for z/OS and OS/390 you will need to say no since DB2 for the mainframe doesn't support Unicode Databases until DB2 V8.

5.1.2 Installing Crystal Reports

Crystal Report, from Seagate, cannot be run on z/OS and therefore needs to be installed on the client workstation and on the file server.

Just insert the CD-ROM. If not launched automatically, go to the CD-ROM root and double-click **setup.exe**. Select the destination folder. For the file server, use the map directory. Choose an HR8xx folder, where xx is the PeopleSoft release as shown in Figure 5-15.

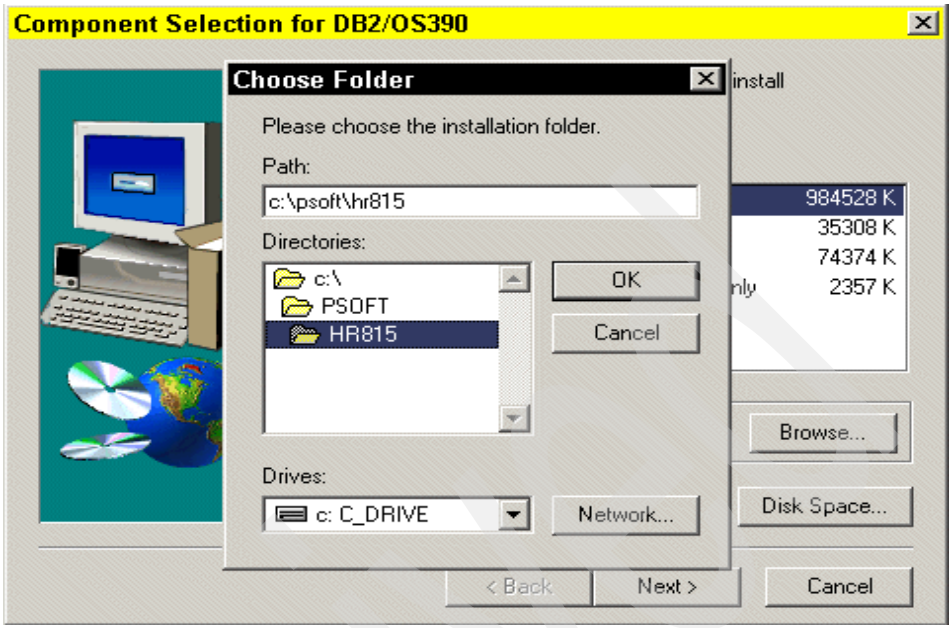

Figure 5-15 Map directory for Crystal Reports installation

5.1.3 Installing SQR for z/OS

Next is the process used to install the SQR product (on z/OS in our case) as described in task 2.10 of the PeopleSoft Installation Document.

1. Go to folder C:\PSOFT\HR8xx\BIN\SQR\z/OS and double-click **SQR8xx/OS.EXE**, where **xx** is the release of the product. This will extract all the files and folders required for the installation. We kept the default destination folder.
2. Log on to TSO on your system.
3. Allocate a PDS file which will receive the Installation JCL with regards to the naming conventions you have chosen in the pre-installation checklist. Use PeopleSoft recommendations or your own system guidances. We used PSOFT.SQR8xx.CTNL and PeopleSoft proposed DCB.
4. Transfer the Allocation JCL to this PDS using FTP:
   a. Logon to a Prompt Command session.
   b. Enter (under c:\) root the following command:

      ```
      C:\ftp YourSystemIPAddress then Enter.
      ```

   c. Enter your User ID then select **Enter.**

d. Enter your Password then select **Enter.**

e. The root is now FTP>. Enter the following command:

```
put c:\sqr8xxz/OS\install\sqrlib.jcl 'psoft.sqr.cntl(sqrlib)'
```

5. Log on to TSO. The JCL is on the library ready for changes:
   – Change the HLQ Qualifier, Allocations and Disk.

Note: PeopleSoft asks you to change the $SQRHLZ$ to your High Level Qualifier (Pre-installation Checklist Parameter). In our case, we chose **PSOFT.SQR8xx**. Be careful when you later run the Server Transfer Program Task 4C-1, as the parameter 26 'SQR High Level Qualifier' will have to be PSOFT.SQR8xx.SQR.

   – Create a JOB member to store the job card you will use during the installation. We have added few parameters to the PeopleSoft proposal:
     - `REGION=0M`: Some programs like IEBCOPY will require a Region Cards, so add the parameter for all JCL.
     - `/*JOBPARM SYSAFF=*`: This is required in case you are running in a SYSPLEX environment to be sure your JCL will be executed in the right subsystem.
   – Submit the JCL and check the return code using SDSF.
6. Open a new command window and run the FTP transfer as indicated in the installation document. The <> is just to highlight and must not be entered.
7. Use TSO to change the PSOFT.SQR8xx.INSTALL.JCL as described.

Note: Do not forget to change the DSN610 DB2 High Level qualifier to the appropriate value, in our case, `DB7J7`. For example, DSN610.SDSNLOAD was changed to DB7J7.SDSNLOAD.

8. Submit the several JCL jobs specified in the Installation document: LNKUSQR, LNKUSRQT, LNKUSRQQP.
   a. Change SQR HLQ and the job card
   b. Change DSN610.SDNSLOAD in OBJLIB2
   c. Change DSN610.SDSNLOAD in LKED.SYSLIB
9. To validate your installation, PeopleSoft suggests a list of Member Count per PDS created. We did the check and found differences. We have not been able to explain those differences so far, as follows:

SQR.SQR.ASSEMBLE with one member DBCALL1 instead of 0, SQR.GIF with 13 members instead of 25, and SQR.SQRINI with zero member instead of one.

With CRM 8 and PeopleTools 8.4, we had other values.

10. Bind your SQR:
    a. Change SQR HLQ
    b. Change DSN610.SDNSLOAD
    c. Change Plan DSNTIA61 and Library Names for program DSNTIAD (DSNTIA71 and DB7JU.RNLIB.LOAD in our case)
    d. Eventually, change the SQR Plan Name
11. Link your SQR executables. We did not, as we used the delivered version.
12. After the database creation, you will be able to test your SQR using the two famous SQR DDDAUDIT and SYSAUDIT jobs. PeopleSoft provides a SAMPLE JCL in PSOFT.HR8xx.JCLLIB(SQRSAMP) to submit SQR.

6

Installing the zSeries batch environment

In this chapter, we describe our experience while installing the PeopleSoft batch environment on zSeries. This chapter is not intended to reproduce the PeopleSoft installation manual, but act as a supplemental resource. As a general approach, we follow the manual and comment on our findings.

6.1 Installation worksheet

Locate and complete the installation worksheet that you may have filled out in Chapter 3, “Planning for installation” on page 19. You will use the information that you filled in throughout this process.

6.1.1 Allocating zSeries Partitioned Datasets

We followed the installation instructions and allocated one PDS. We used basic ISPF utilities to allocate the dataset and named it PSOFT.PS840.CNTL also as indicated with no problems.

We then FTPed the files as indicated from the file server to our new zSeries PDF, and edited the members as instructed in the installation guide. When we submitted PSLIBCBL and PSLIBSQR, we received jobcard errors. Once those errors were corrected, the jobs ran properly.

6.1.2 Running PeopleSoft server transfer

We became very good at running the PeopleSoft server transfer, because of our typos. If you make such an error, we highly recommend that you delete what you had created in the zSeries PDS files and re-run the server transfer. Your original entries will still be in the PSXFR.CFG file on the file server so you will not need to re-enter everything. A sample PeopleSoft Server Transfer is shown in Figure 6-1.

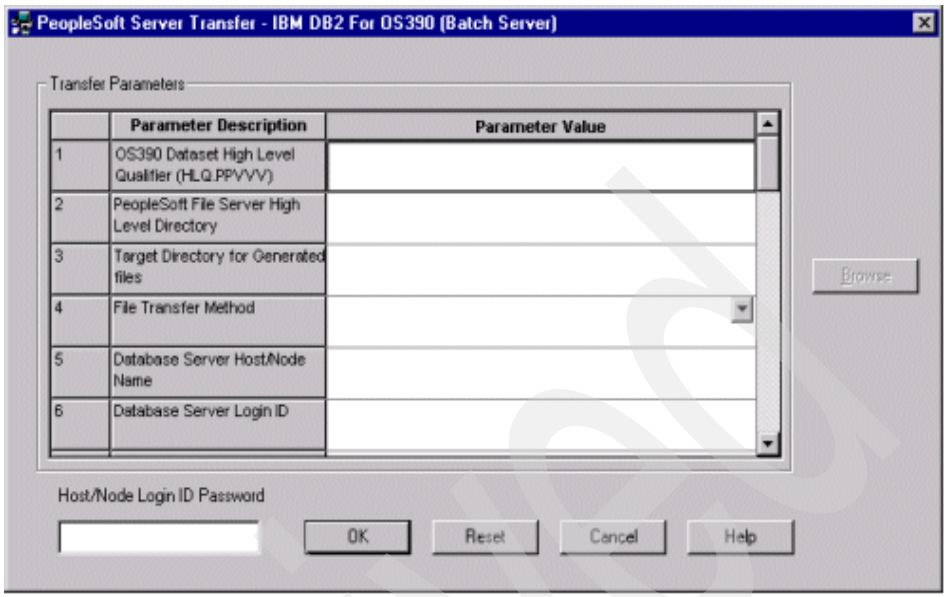

Figure 6-1 PeopleSoft server transfer

6.1.3 Granting access to USS files

When we first logged on to OMVS to run the psmv.sh script, we were unable to run it successfully as we were not in the proper working directory to run this script. Once we went to pshome, we were successfully in granting access.

6.1.4 Installing and binding SQR for zSeries

To install the SQR members into the appropriate z/OS PDS, you will log on to OMVS and run the shell script as indicated in the installation manual. Running these scripts will populate the PDS by submission of a JES batch job. Check this job to ensure that `CC=00`.

This was successful for us and a wonderful improvement over previous versions which required that we load a tape.

We then successfully ran the DB2 bind as indicated in the installation manual.

6.2 PeopleTools programs

Assemble PeopleTools programs and compile and link edit COBOL. We followed the remaining instructions as indicated in the installation manual and received `CC < 4` for all batch compiles.

7

Creating PeopleSoft databases

In this chapter, we discuss installation planning as well as the tasks required to create the PeopleSoft databases. We do not intend to reproduce the *Peoplesoft Implementation Guide* here, but to comment and give additional direction that relates to our own experience installing these products at the Almaden Research Center.

7.1 Installation planning

When planning for your PeopleSoft databases, you must make some decisions prior to the install. For instance, you will have over 10,000 objects to manage with this application. Based on the number of objects, you may wish to consider installing PeopleSoft into its own DB2 Subsystem. In addition, PeopleSoft financial applications require decimal 31 precision. We have not seen that any other application running in the same subsystem as PeopleSoft had any problems with the DEC=31 zparm setting. However, this could be a consideration for your installation.

7.1.1 Table and tablespace strategy

The PeopleSoft installer will install the databases, tablespaces, and tables as delivered by PeopleSoft. If your company has a standard for installation such as table per tablespace strategy, please let the installer know that you wish to modify the DDL prior to installation. PeopleSoft has done a good job at parsing the tables that will become very large into their own tablespaces. You may wish to consider partitioning some of the large tables when they are defined.

Additionally, we recommend putting all of the temp tables into their own tablespace.

7.2 Tasks to be performed

The following tasks need to be performed:

1. Task 7-1: transfer DDL to OS/390.

   From this point, you can find in your target library the DDL PeopleSoft will use to create the DB2 objects: Stogroups, Database, Tablespaces, Tables. Even if you now have the two options, depending on whether or not UpdateStatistics is used, it is almost sure that you will have to change your database organization later to improve performances. This step is creating a standard database. You might decide to create more bufferpools, more stogroups, and more tablespaces in order to put some tables in their own space or to regroup unused tables in a single tablespace for easier management. A lot can be done, but this is perhaps not the right moment. It is probably best to change the defaults later, unless you have a clear view of what you want.

2. Task 7-2: create PS.PSDBOWNER as shown in Figure 7-1.

```
 File  Edit  Edit_Settings  Menu  Utilities  Compilers  Test  Help
EDIT       PSOFT.HR815.DDLLIB(PSDDL) - 01.10               Columns 00001 00072
Command ===> _____________________________________________ Scroll ===> PAGE
000039    CREATE DATABASE PSOWNRDB STOGROUP PSSGTSPT  ;
000040    COMMIT;
000041 -- CREATE PEOPLETOOLS TABLESPACES
000042    CREATE TABLESPACE PTOWNER    IN  PSOWNRDB
000043           USING STOGROUP PSSGTSPT PRIQTY 48    SECQTY  48
000044           FREEPAGE  0 PCTFREE 20
000045           SEGSIZE    4 BUFFERPOOL BP0   LOCKSIZE ANY  CLOSE NO  ;
000046
000047    GRANT USE OF TABLESPACE PSOWNRDB.PTOWNER TO PS ;
000048    GRANT CREATETAB ON DATABASE PSOWNRDB TO PS ;
000049
000050    SET CURRENT SQLID='PS';
000051    CREATE TABLE PSDBOWNER
000052    (DBNAME CHAR(8),
000053     OWNERID CHAR(8),
000054     DEDICATE_THREADS CHAR(1))
000055    IN PSOWNRDB.PTOWNER;
000056
 F1=Help      F2=Split     F3=Exit      F5=Rfind     F6=Rchange   F7=Up
```

Figure 7-1 Create PSDBOWNER

3. Task 7-3: grant privileges on PS.PSDBOWNER.
4. Task 7-4: grant privileges to OwnerID.
5. Task 7-5: create Database, Stogroup and Tablespaces.

   For Stogroup, you have to change at least the list of disks serial to be included and the VCAT name, and of course, the OWNER of the objects.
6. Task 7-6: create Tables.
7. Task 7-7: configure the DB2 Connect Gateway.

   This should have been done already as detailed in Chapter 4, “Installing and configuring DB2 Connect” on page 43.
8. Task 7-8: create the DataMover Imports Script.

   You have to select the Database. Select **DB2 UDB** for OS/390 as shown in Figure 7-2.

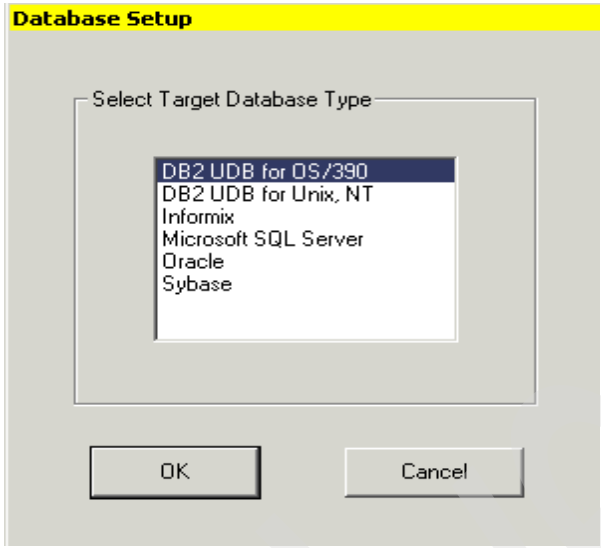

Figure 7-2 Database setup

Specify the correct values for Stogroup, owner ID, and Access ID as shown in Figure 7-3.

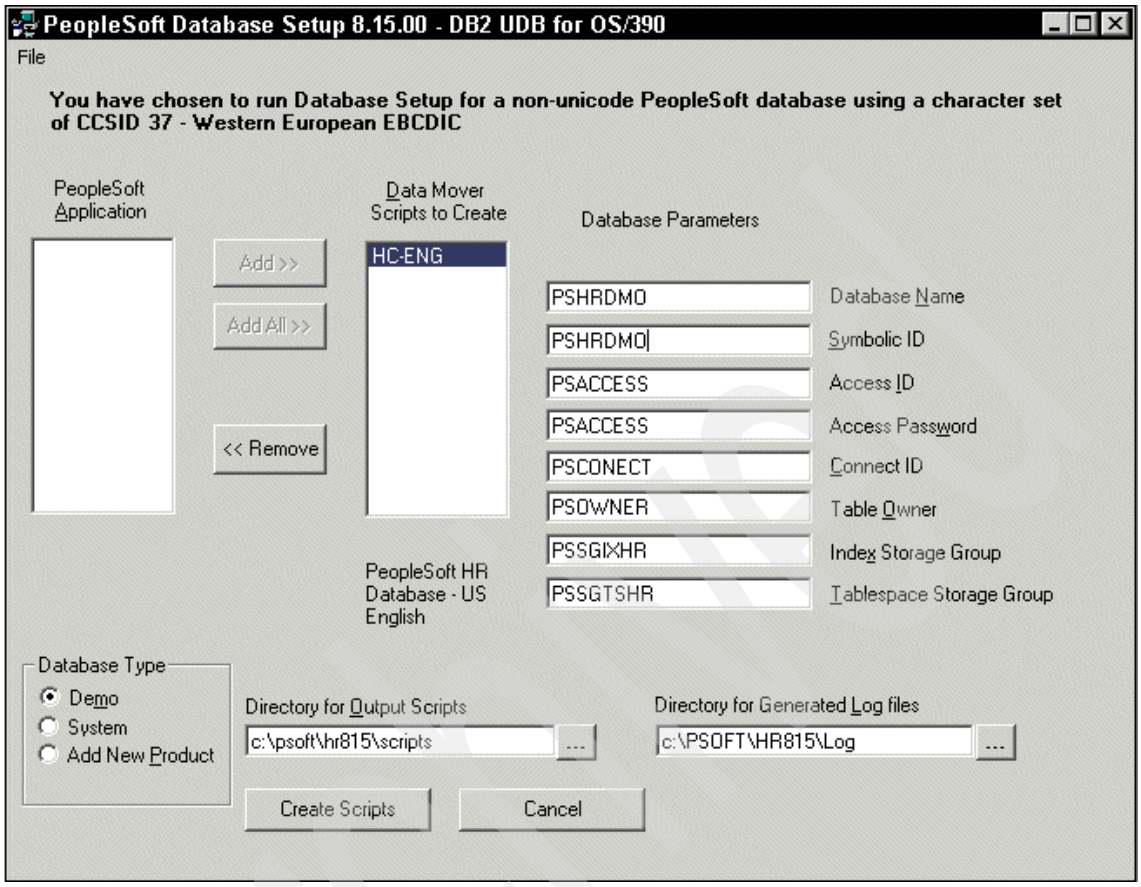

Figure 7-3 Database ID setups

9. Task 7-9: run Data Mover Import Scripts.

   When connecting to DataMover on the Boot Strap mode, you may or may not experience an error message, as shown in Figure 7-4.

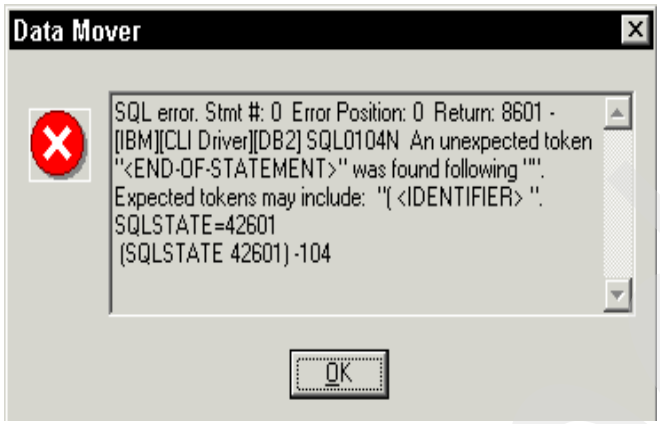

Figure 7-4 Data Mover error window

If you have it, go to your Environment Variables:

a. Click **Start** -> **Settings** -> **Control Panel**.

b. Then go to **System** -> **Advanced** -> **Environment Variables**.

c. Enter **New...**:

- Variable Name: `DB2DBMSADDR`
- Variable Value: `30000000`

Results should be as shown in Figure 7-5.

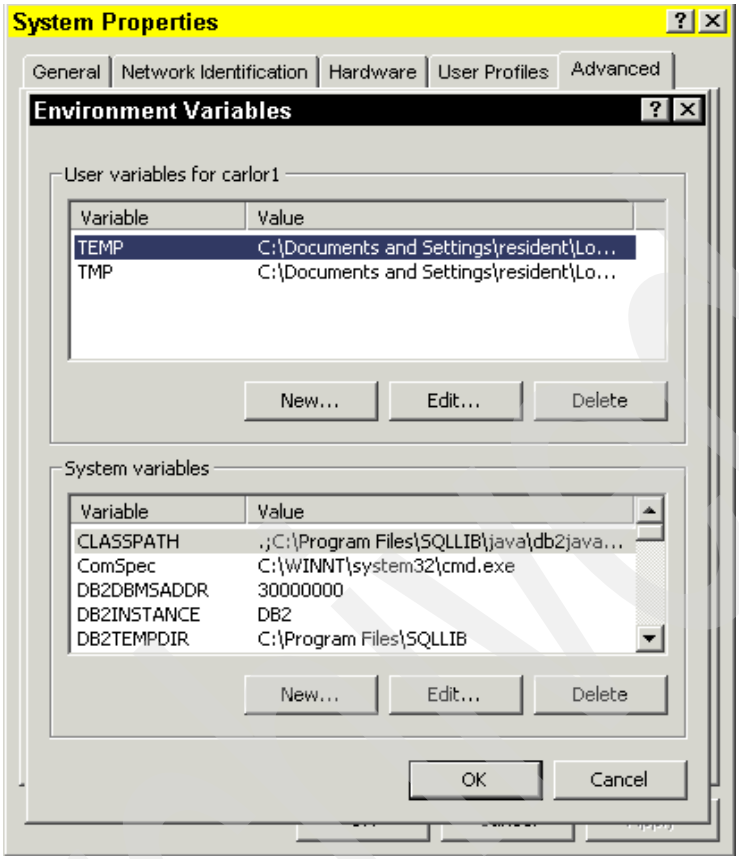

Figure 7-5 System properties

You should connect with the user ID you have defined in ConnectID while implementing the Client using Configuration Manager.

PeopleSoft put a note to highlight that the DataMover transfer is generating a lot of Logging activities. If you are in your own DB2 Subsystem without other activities, you can remove the logging activity during this task.

This is a long process, as announced. You should perhaps take that time to prepare your Runstats cards (not provided by PeopleSoft). A Select in the SYSIBM.SYSTABLESPACE will require a little bit a reformatting. Runstat cards will then be ready to run. Prepare your I/C jobs too, as you will need them soon.

10. Task 7-10: create Indexes.

11. Task 7-11: update the Database to the latest PeopleTools release.

    We had to do this during one of our tests. It is a quite long process. Don't forget to back up your environment beforehand, in case you encounter any problems later.

    First, you have to run the SQLs, after FTPing them to TSO in order to run them. You need to change the owner, the database and Stogroup names.

    Second, you have to copy the project in the PeopleTools database to align the metadata. In Copy Project from Files, be aware that you must target the directory **projects** and not the project folder itself.

    The document ask us to copy two projects (step 3) in the Copy Project From File section and then click the **COPY** button of the Copy Dialog box (step 4). It seems to be impossible. You have to do both steps 3 and 4 in sequence for each chapter.

    You also have to Update PeopleTools System Data by running the Scripts indicated using Data Mover.

12. Task 7-12: restore the Informatica Repository. This is not applicable in our case.
13. Task 7-13: prepare for Initial Sign-On.
14. Task 7-14: preload PSRECFIELDDB table.
15. Task 7-15: run the DB2 RUNSTATS utility.

    Hopefully, everything has been ready since Step 9 on page 93.

16. Task 7-16: create PeopleSoft Views as shown in Figure 7-6.

    Using DataMover to create several thousands of views is not very quick, so be ready for a long process (several hours).

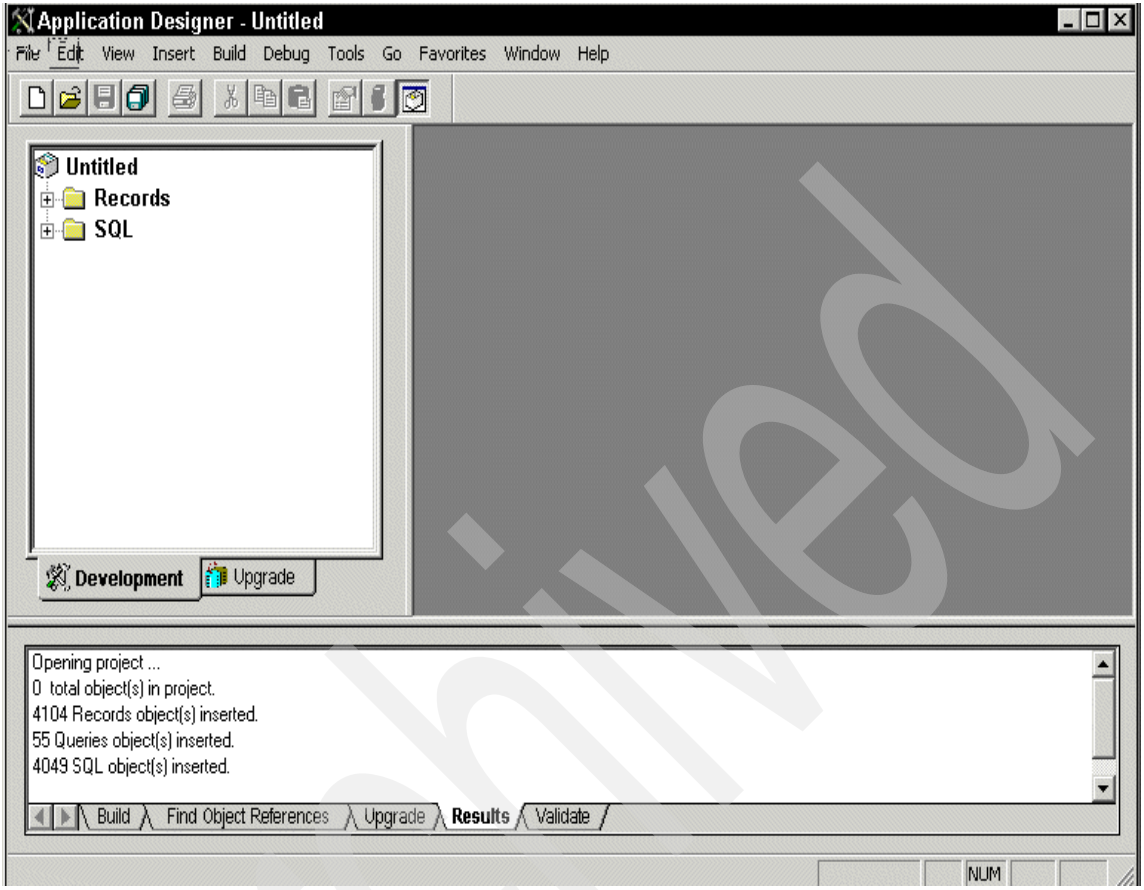

Figure 7-6 Creating PeopleSoft views

17. Task 7-17: update PSRECFDEFN for %UpdateStats. This is not applicable in our case.
18. Task 7-18: build Temporary Tables.
19. Task 7-19: run Additional Data Mover Scripts. You can change later the number of temporary tables you will use but if you can anticipate, do it now. Each set of temporary tables will allow you to run a process in parallel in order to exploit the capacities of z/OS and will require you to find criteria to split the data.
20. Task 7-20: load Stored SQL Statements.
21. Task 7-21: prepare to run SQR.

    Since we were planning to run SQR on OS/390, we did not perform this task.

22. Task 7-22: update PeopleSoft System Tables.
23. Task 7-23: bind DB2 Plans.
24. Task 7-24: run the Version Application Engine Program.

    We tried, and received a warning: The window opened by PSAE closed automatically before we could check what was in it.
25. Task 7-25: change the Base language.

    As we were running only English, we did not perform this task.
26. Task 7-26: run Currency Scripts against the SYS Database.

    We have installed a DMO database, so this task was not relevant in our case, although it might be in yours.
27. Task 7-27: check the Database.

    To validate that the installation is correct, PeopleSoft delivers two SQRs named DDDAUDIT and SYSAUDIT. To run them on OS/390, you need to submit them through JCL.

    a. Go to your JCLLIB library (PSSOFT.HR8xx.JCLLIB in our case).
    b. Select member **SQRSAMP**. You have sample JCL ready to run.
    c. Validate the different variables and change the %PRCSNAME% to `DDDAUDIT` and `SYSAUDIT` for the second run.
    d. Submit that JCL. Results are available on the SQRLIST dataset (PSOFT.HR8xx.SQRLIST in our case). Be careful, it seems that even if the output is declared as SHR, the job will try to access in OLD mode.
28. Task 7-28: disable %UpdateStats.

    You will have to do this twice, once for Process Scheduler, and once for Client Configuration.

    a. Go to Client Configuration.
    b. Edit the Profile.
    c. Select **Disable DB Stats** in the Process Scheduler part.

    You will also have to check and set `Dbflags =1` (to disable) while configuring the Process Scheduler on OS/390.

8

Installing and configuring the PeopleSoft application server

This chapter describes the steps we followed to install and configure an application server on a server running Windows 2000. The PeopleSoft application server is based on BEA products TUXEDO and JOLT.

8.1 Installing Tuxedo and Jolt

1. Task 6A-1: uninstall any previous version of Tuxedo from your server.
2. Task 6A-2: design an Administrator.
   a. On Windows 2000, you must click **Start -> Settings -> Control Panel** as shown in Figure 8-1.
   b. Select **Users and Passwords**.

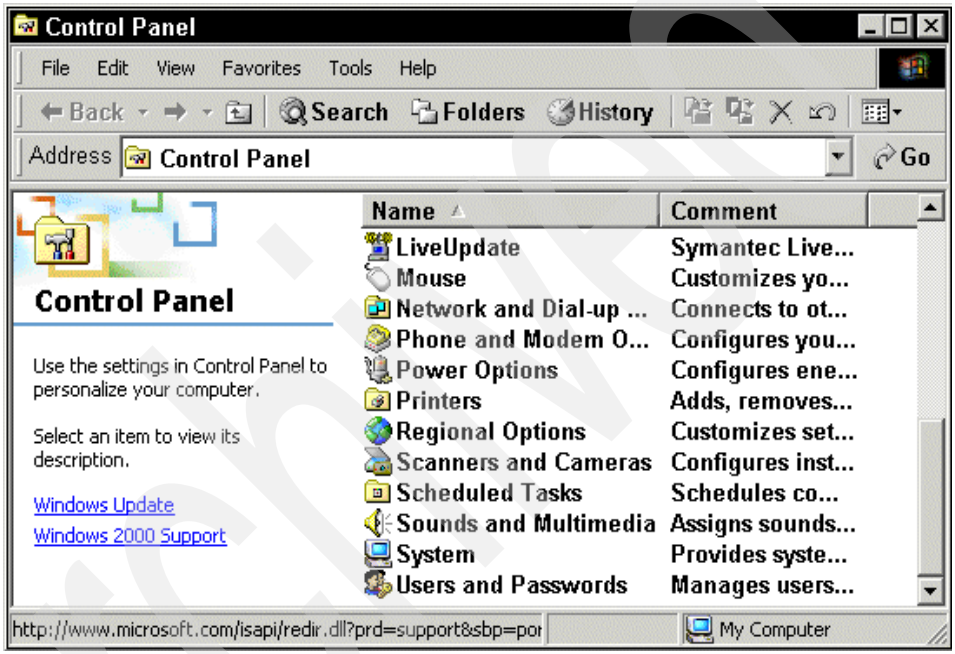

Figure 8-1 Control panel

   c. Then you should add a user called TUXADM, as shown in Figure 8-2.

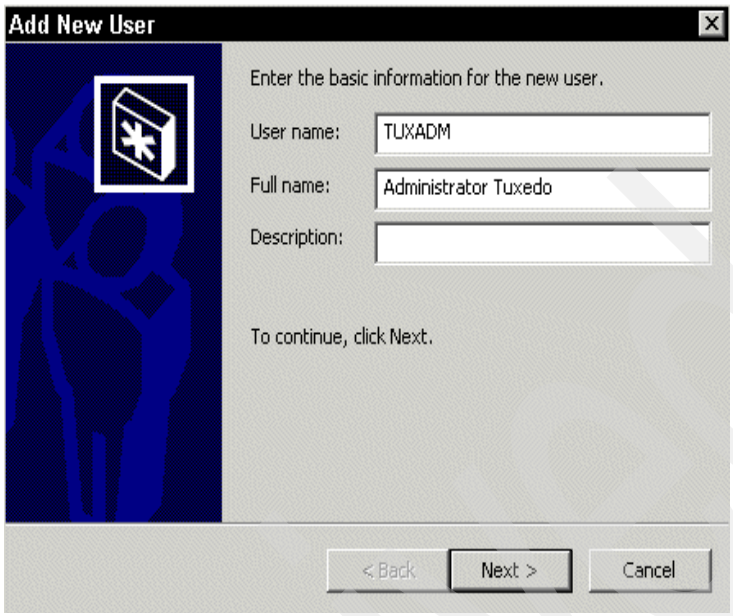

Figure 8-2 Add new user

d. Give the user ADMINISTRATOR Group authority, as shown in Figure 8-3 and Figure 8-4.

Figure 8-3 Type of user

e. Then click **Finish**. Your Administrator is declared. You may be asked to reboot.

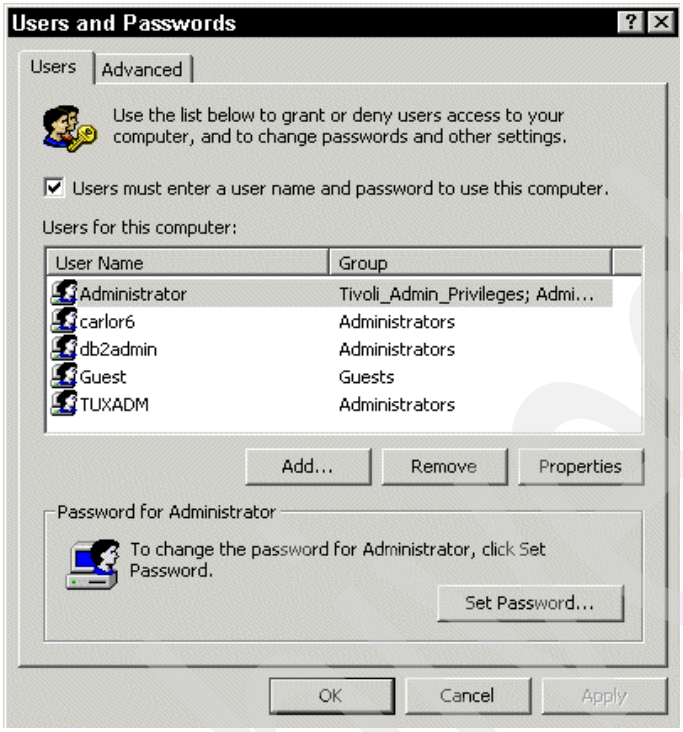

Figure 8-4 Users and passwords

3. Task 6A-3: set TUXDIR in Environment Variables as shown in Figure 8-5.
   a. Go to **Start->Settings->Control Panel**.
   b. Then click **System->Advanced-Environment Variables**.

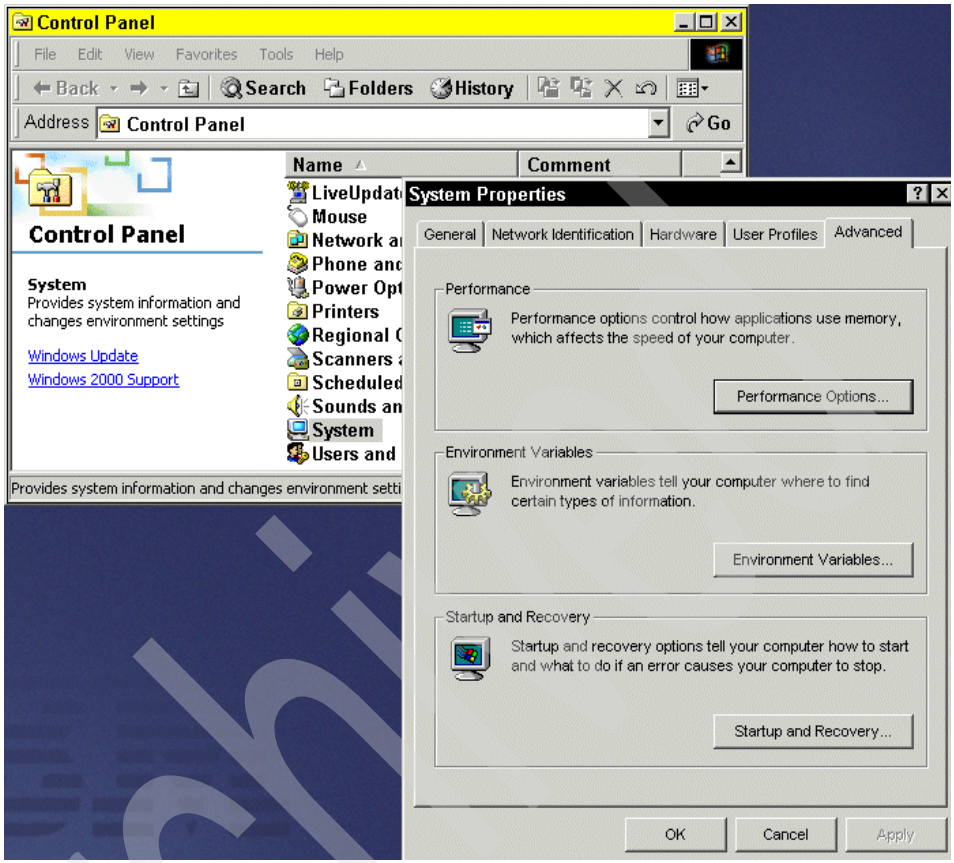

Figure 8-5 System properties

c. In System Variables, select **New** and enter the Variable name `TUXDIR` and its value (we kept c:\tuxedo), then click **OK** and **OK** again to apply your changes, as shown in Figure 8-6.

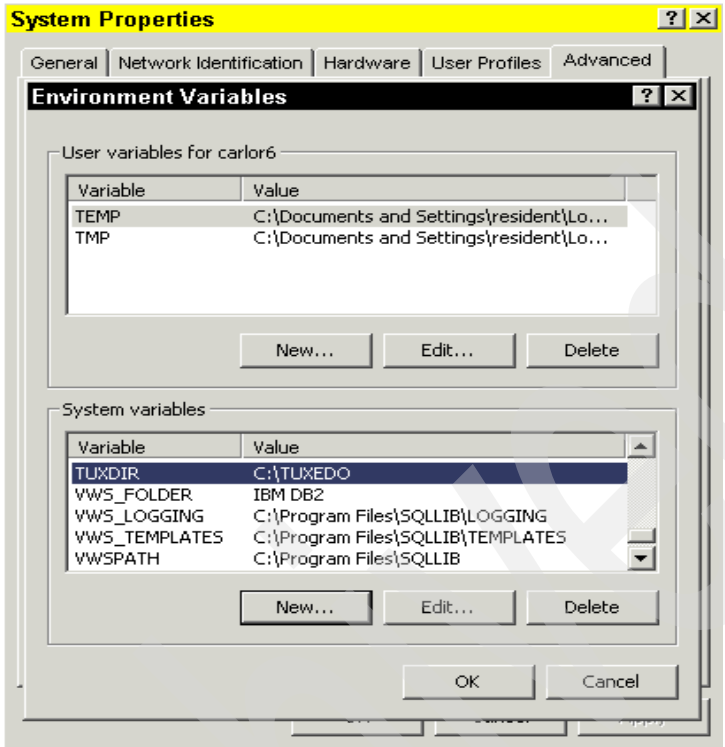

Figure 8-6 Environment variables

4. Task 6A-4: install TUXEDO.

   Do not forget to match the folder name you use here with the destination folder used to declare tuxdir in the previous step.

5. Task 6A-5: change the Service Account.

   a. Go to **Start -> Programs -> Administrative Tools -> Component Services**.

   b. Select **Services** and drill down to TUXEDO IPC Helper. Check if it is started. If so, stop it as shown in Figure 8-7.

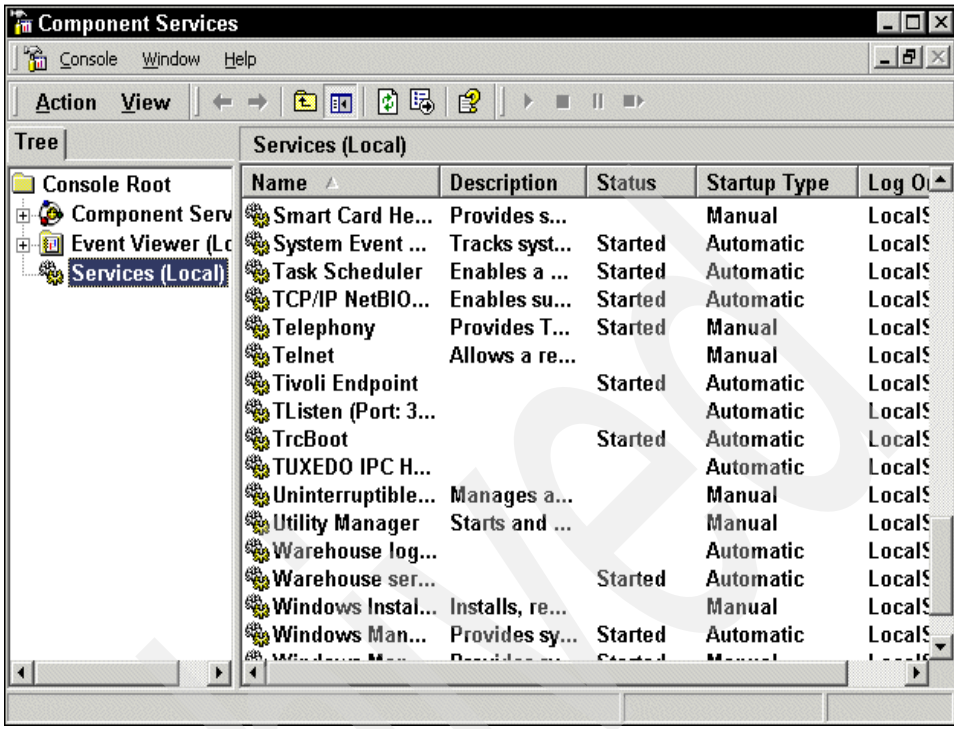

Figure 8-7 Component services

c. Double-click **TUXEDO IPC Helper** as shown in Figure 8-8.

d. Make the changes you need and click **Start**.

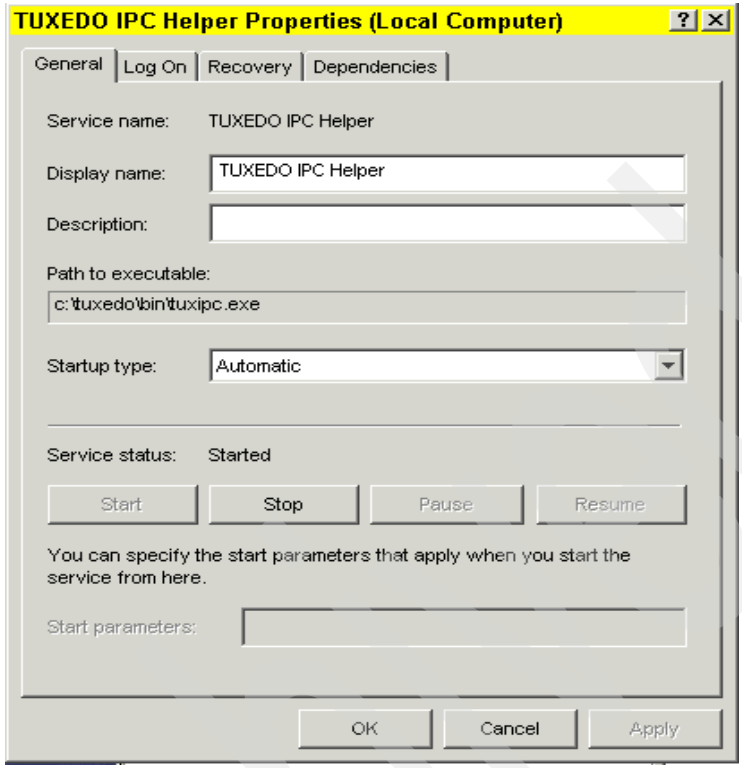

Figure 8-8 TUXEDO IPC helper

6. .Task 6A-6: check your Environment Variables.
7. Task 6A-7: verify the Server Installation: Tuxdir refers to your installation directory.

8.2 Configuring the application server on Windows 2000

1. Task 8A-1: set the environment and path variables.
2. Check Authorization.You should have one of your **Permission Lists with Can Start Application Server** selected, as shown in Figure 8-9.

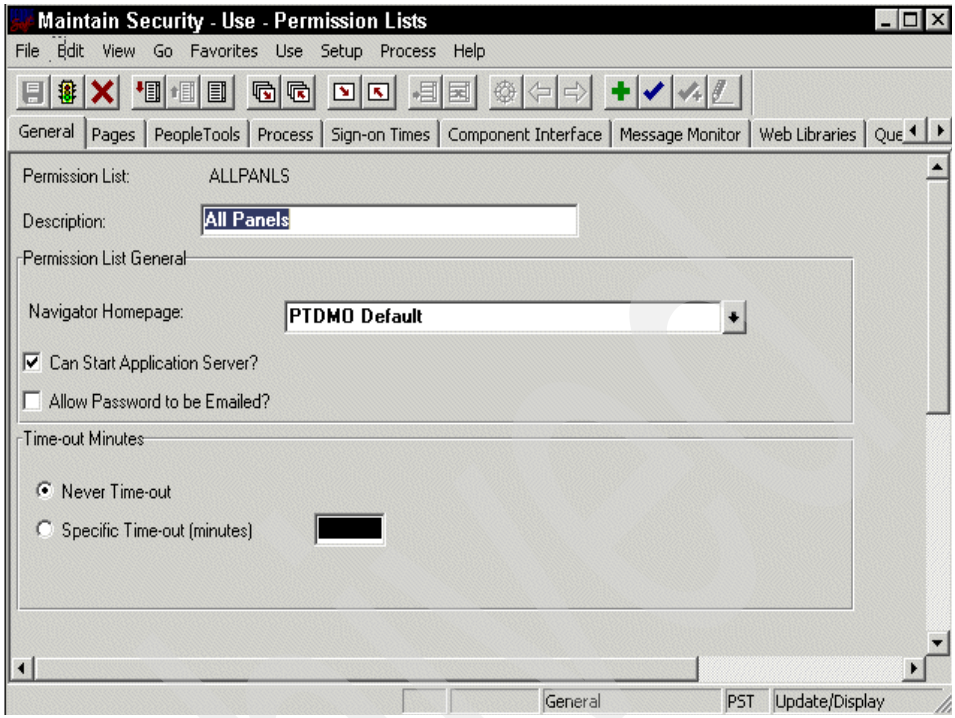

Figure 8-9 Permission lists

Enter the values required for the components transfer as shown in Figure 8-10.

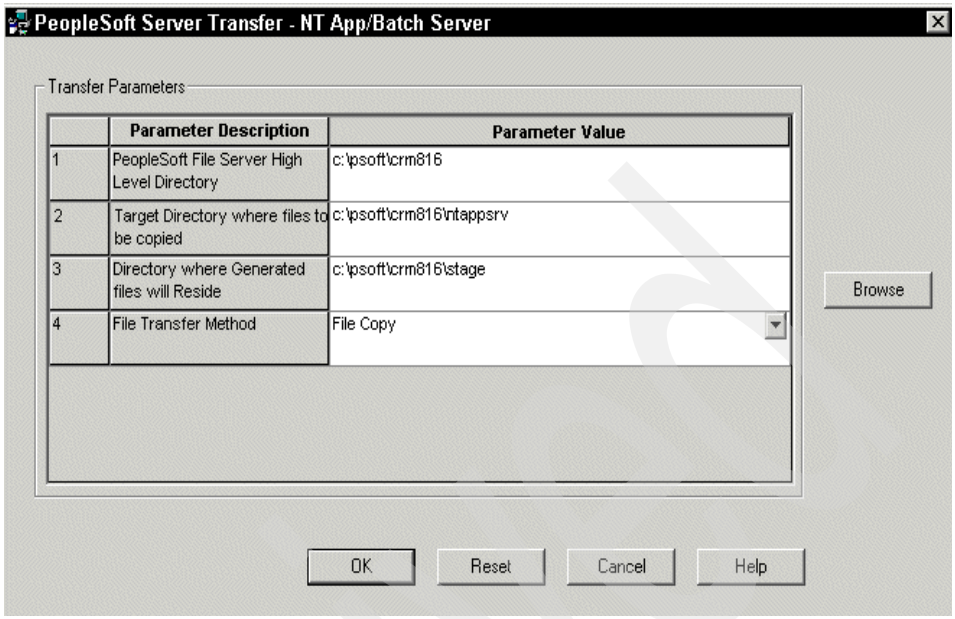

Figure 8-10 Component transfer

3. Task 8A-1: set Environment and Path Variables.
4. Task 8A-2: set up Cobol for remote call.
5. Task 8A-3: set up the initial application server domain. Go to your <ps_home> directory->appserv and double-click **psadmin.exe**. Go to option 2. Enter the name of the domain. By default, we used the database name.
6. Task 8A-4: specify the Application Server Domain Parameters as shown in Figure 8-11.
7. You cannot use the quick configuration as it is referring to a Microsoft Database type.

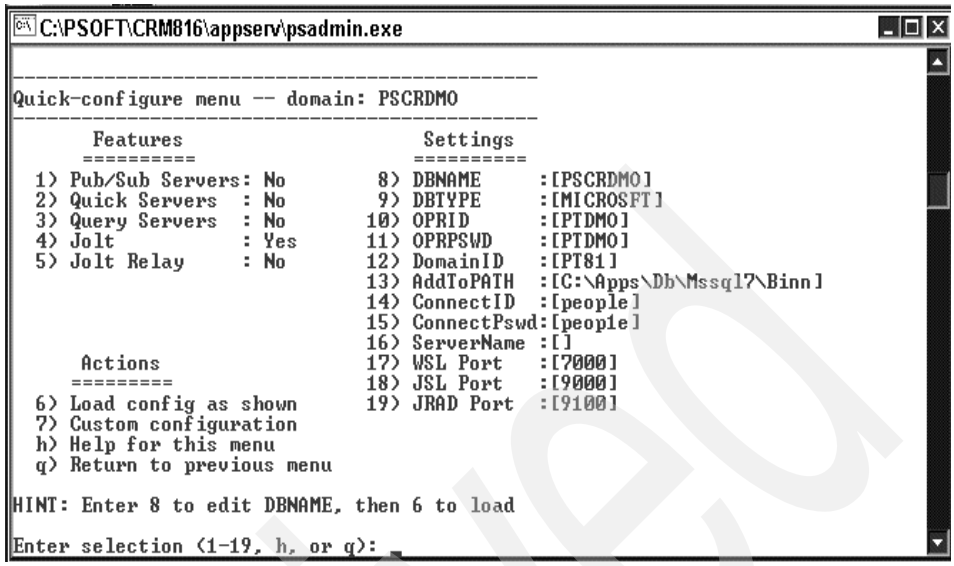

Figure 8-11 Server domain parameters

Therefore, you must go through the normal configuration process. The parameters we had to change are listed in Table 8-1.

Table 8-1 Parameters changed

Parameter	Value we used
DBName	PSCRDMO
DBType	DB2ODBC
Userid	VP1(any valid id)
UserPwd	VP1(any valid id)
Encrypt Psw	Yes
Connectid	PSCONN (the ConnectID you have chosen)
Connect PSW	PSOTF8 (and its Application Server sword)
DomainID	Application Server 1 (PeopleSoft Application Server1)
AddtoPath	C:\Program Files\sqllib\bin (where the DB2 Connect bin file is
All Others as Default	

8. Task 8A-5: verify Database Connectivity.

   We did this previously via DB2 Client Configuration Assistant.

9. Task 8A-6: start the application server.

   You should have nine processes or more started, depending on your configuration.

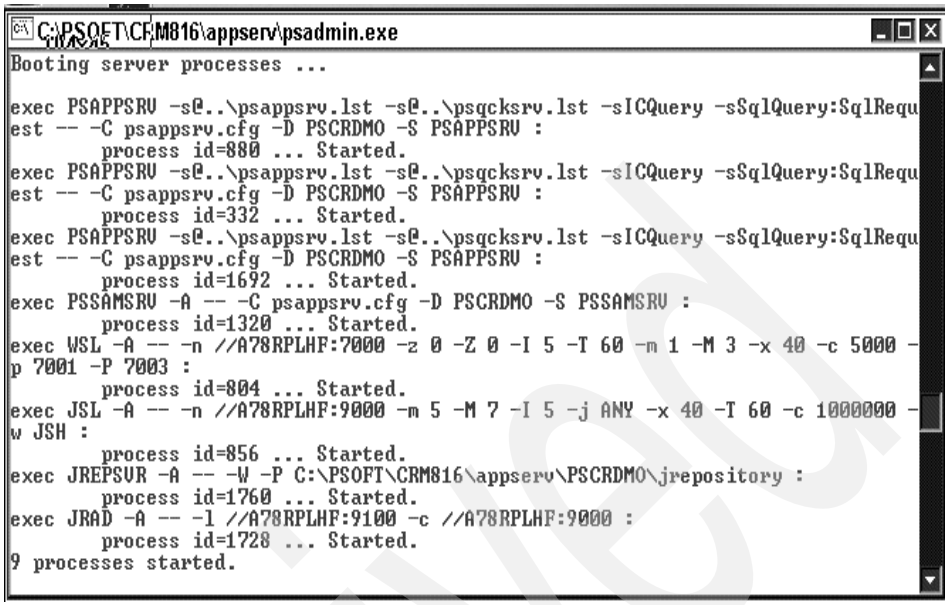

Figure 8-12 Application server started

10. Task 8A -7: shut down the application server.
11. Task 8A -8: examine log files.
12. Task 8A -9: set up Multiple Domains: you need to change the values as indicated by PeopleSoft to keep each domain independent. A new Domain name and new Port number for each process are required.

9

Installing and configuring the Web server

This chapter describes the different steps we followed to set up the PeopleSoft Internet Architecture on Windows 2000. With PeopleSoft 8, the PeopleSoft Web Server has become a key component of the architecture.

The Web Servers supported by PeopleSoft 8.4 are:

- BEA WebLogic Server 5.1
- Apache HTTP Server 1.3.12
- IBM WebSphere Application Server 4.0

We installed the imbedded Apache HTTP Server 1.3.12. However, the process to upgrade to WebSphere Application Server 4.0 is available on PeopleSoft Customer Connection and IBM WebSphere Application Server 4.0 is embedded with PeopleSoft release 8.4.

9.1 Installing Apache 1.3.12 on Windows 2000

Installation of Apache Web Server is a four-step task. The PeopleSoft Installation Document is self-explanatory in those four steps:

1. Install the Java Runtime Environment.
2. Install the Java Servlet Development Kit.
3. Install the Apache Web Server itself.
4. Install the Apache Jserv (Java Server) as shown in Figure 9-1.

   In the version we have installed, the word `Original` was not part of the jserv configuration, so don't worry if you don't find it. You can still perform the changes required.

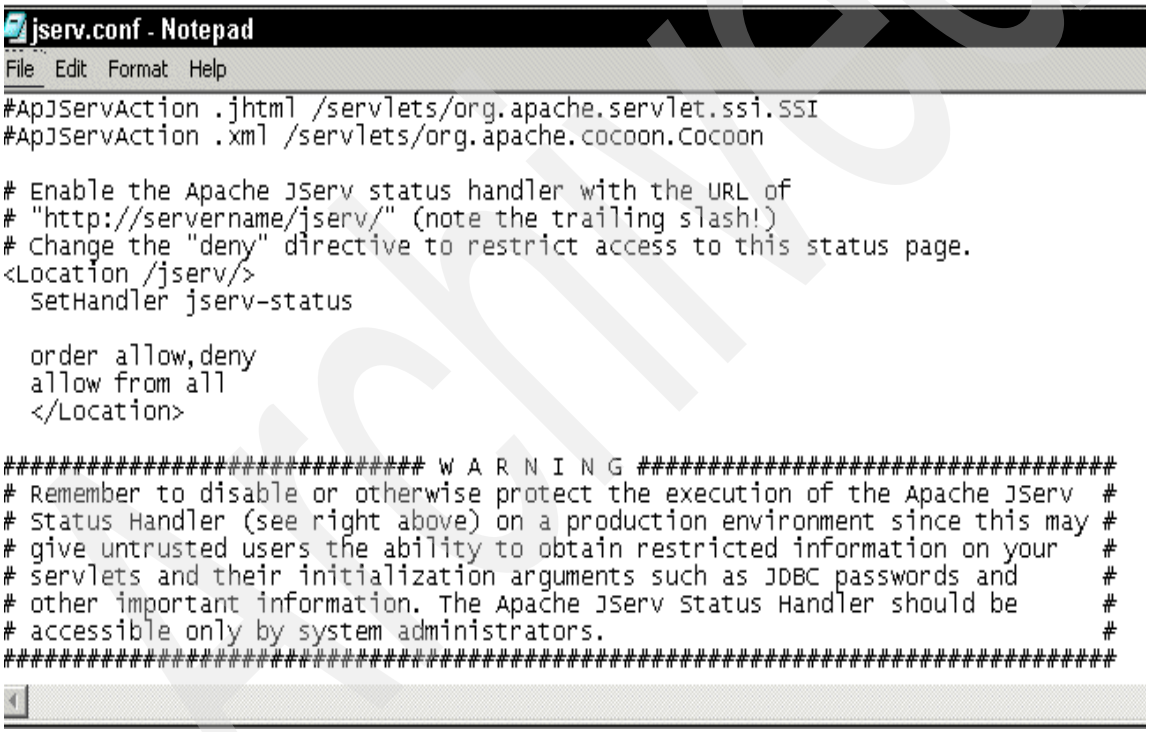

Figure 9-1 jserv.conf file

5. You can now test your Web server as indicated.
6. You have to Install the PeopleSoft Internet Architecture as indicated in installation guide task 9A-4. Just follow instructions from PeopleSoft.

7. You are now ready to test your PeopleSoft.
   a. Open your Web Browser.
   b. Enter in the address field: `<Host name>/servlets/iclientservlet/peoplesoft8/?cmd=login`.
   c. You should see the following error message: `You do not have access to this page! Use Maintain Security to add access to this page`.
   d. In that case, you need to:
      i. Connect to Application Designer as VP1
      ii. Select **PeopleTools -> Maintain Security -> Use -> Use Permission Lists**.
      iii. Select **ALLPANLL -> Web Libraries**. Select **WEBLIB_MENU->Edit** and make sure all have full access.
      iv. Save and retry. You should see Figure 9-2.

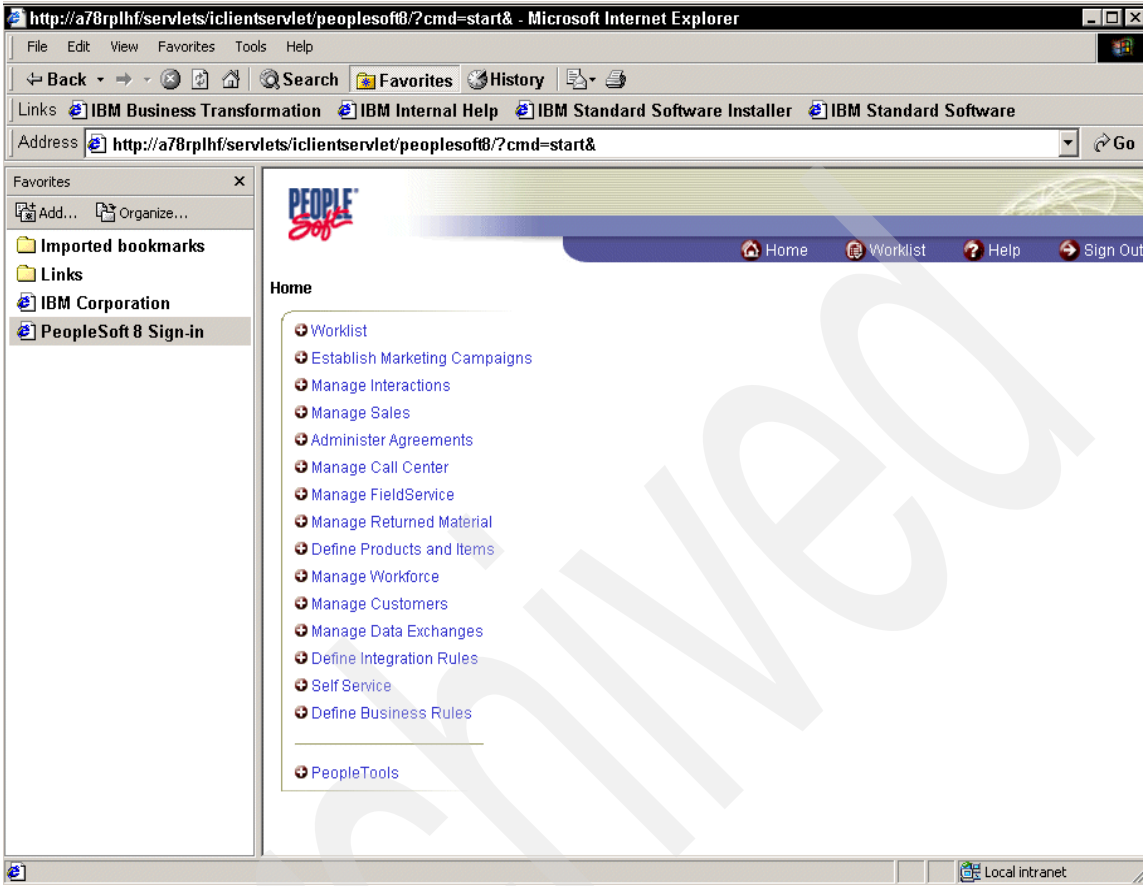

Figure 9-2 PeopleSoft main window

10

Configuring Process Scheduler on Windows 2000 and z/OS

This chapter describes the different steps you will go through to install PeopleSoft Process Scheduler. We discuss the Windows 2000 setup for products that run only on Windows NT/2000, such as Cube Manager or Crystal Reports. Then we discuss the z/OS setup for products such as SQL, COBOL, or Application Engine programs.

10.1 Installing Process Scheduler on Windows 2000

Note: We did not do this on Windows NT but the same should apply.

Some of the PeopleSoft Tools, Products or third parties are not supported on operating systems other than Microsoft Windows NT 4.0 or Windows 2000. This is why we recommend that you have a dedicated Intel-based server to act as a Report Server in PeopleSoft on zSeries architecture.

PeopleSoft Installation Guide refers to this chapter as 10A.

1. Start PSADMIN by double-clicking **<ps_home>\appserv\psadmin.exe**.
2. Select **Option 2 -> Process Scheduler**.
3. Select **Option 4-> Create a Process Scheduler Server Configuration**.
4. Enter your Database Name. It will be used to identify which Process Scheduler Profile you want to configure, start or stop later on.
5. Select the **NT Template** as indicated by PeopleSoft and check your messages.
6. Enter `Q` to quit.
7. Select **Option 3 -> Configure a Process Scheduler**.
8. Select **Option 1** which is the configuration you have created in a previous step.
9. Answer `Y` to the question `Do you want to change any values (y/n)?`.

For each block of parameters, you will be ask if you want to change one of the values. Then each value with the current setting will be proposed for change. Selecting **Enter** implies no change.

In this chapter, we will indicate only the blocks where we made change versus the PeopleSoft Documentation.

1. Block 1: Startup
   - `Dabbing = PSCRDMO` as defined in your configuration
   - `DBType = DB2ODBC`. The PeopleSoft Documentation is misleading. Your database is on DB2 UDB for z/OS, so your type is DB2ODBC.
   - Enter your user ID and password and request encryption.
   - Enter your ConnectID and associated password and request encryption.
   - Keep Server Name unchanged. This is not applicable.

   The results are shown in Figure 10-1.

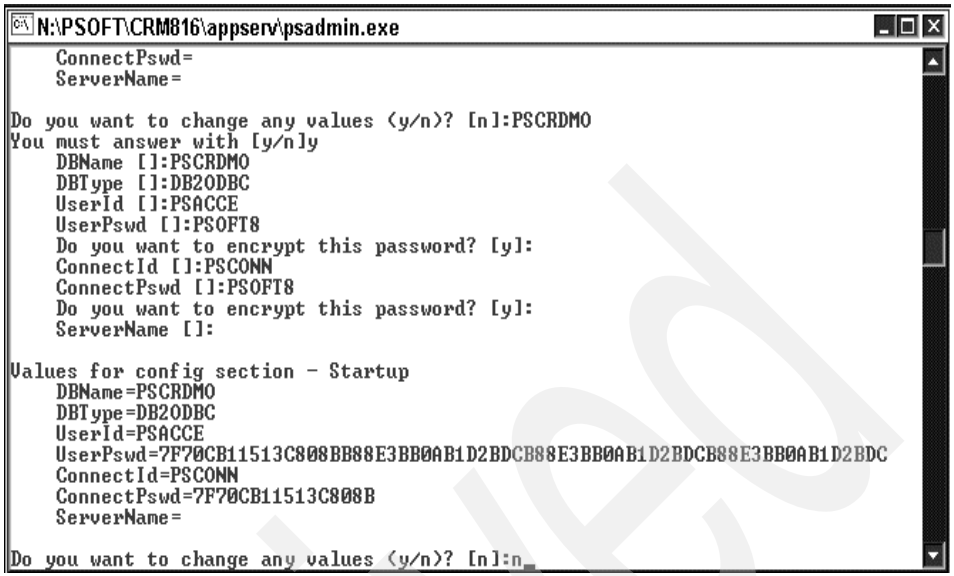

Figure 10-1 Application server startup

2. Block 2: Process Scheduler
   - DBBIN must include the path to your DB2 Connect binaries. For example, `c:\Program Files\SQLLIB\bin`.
   - CBLBIN should indicate your <ps_home>/cblbin library.
   - Other elements remain unchanged. Verify that `DbFlags=1` to disable %UpdateStat.

   The results are shown in Figure 10-2.

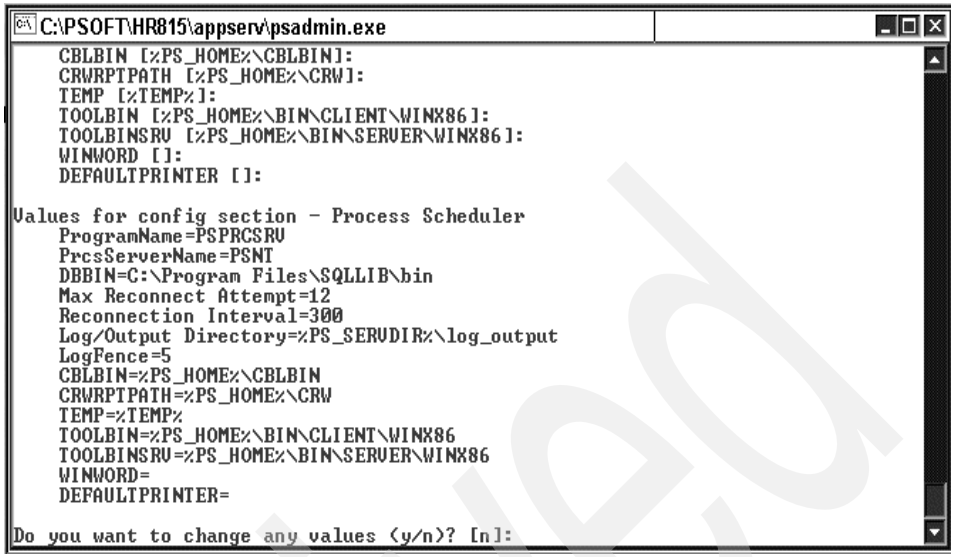

Figure 10-2 Process scheduler

You can now test your process schedule:

a. Go to Start a Process Scheduler Server.

b. Select your database in the Database List.

c. A screen will pop up and, after a few seconds, it should indicate that your process scheduler is started, as shown in Figure 10-3.

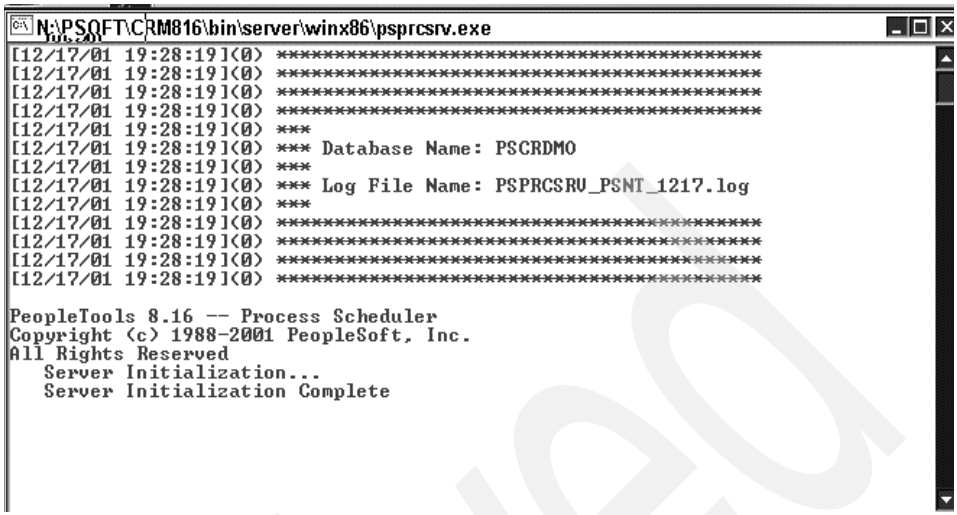

Figure 10-3 Process scheduler initialization

10.2 Installing a Process Scheduler on z/OS

Running a Process Scheduler on z/OS is required if you plan to run your Batch Server on z/OS (except for the Windows-only product).

The configuration of a Process Scheduler is done under the UNIX System Services of the z/OS.

1. Connect to USS via Telnet.
2. Enter `cd $PS_HOME` to change the directory to your PS_HOME value (see Chapter 4, “Installing and configuring DB2 Connect” on page 43).
3. Enter `cd appserv` to go to this directory.
4. Enter `psadmin`.

   PSADMIN will be started and the configuration screen will appear (similar to the one for Windows) as shown in Figure 10-4.

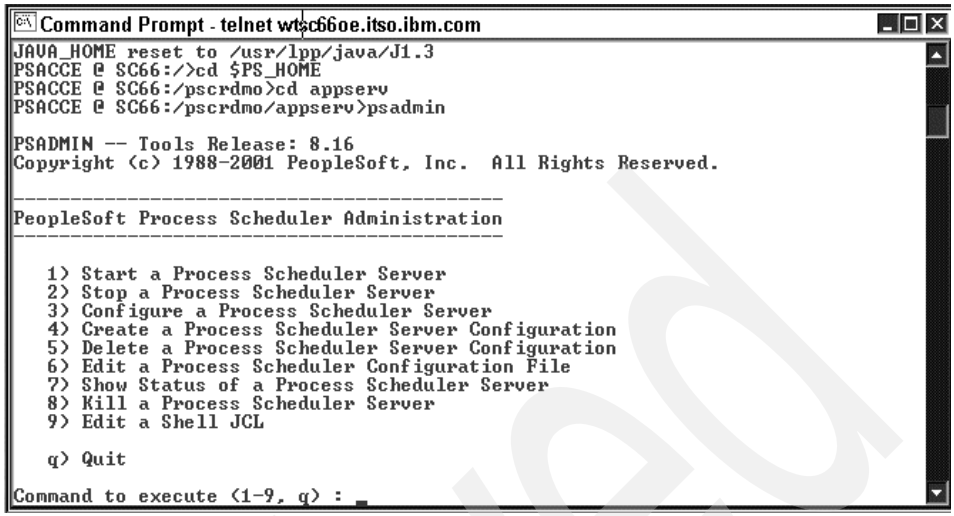

Figure 10-4 Process scheduler administration

5. Go to Create a Process Scheduler Server Configuration.
6. Enter the database you want to connect to.

   OS390 is automatically selected as this is the only possible choice under USS, as shown in Figure 10-5.

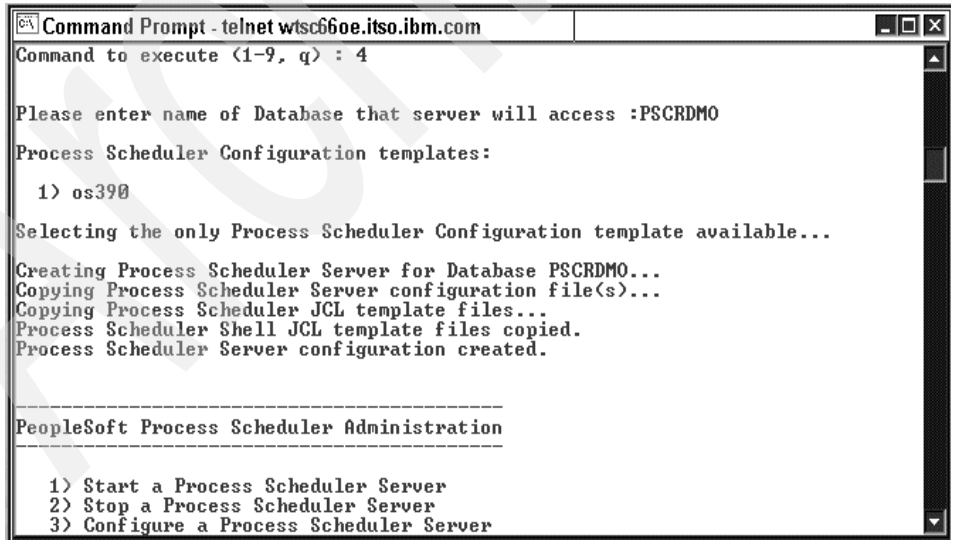

Figure 10-5 Database connect

7. Then go to Option 3: Configure a Process Scheduler Server.
8. Select the database you want to work with.
9. Respond `Y` to `Do you want to change any of those values (y/n)?`.

   In the first block, enter the database name and the different user IDs, just as for Windows, as shown in Figure 10-6.

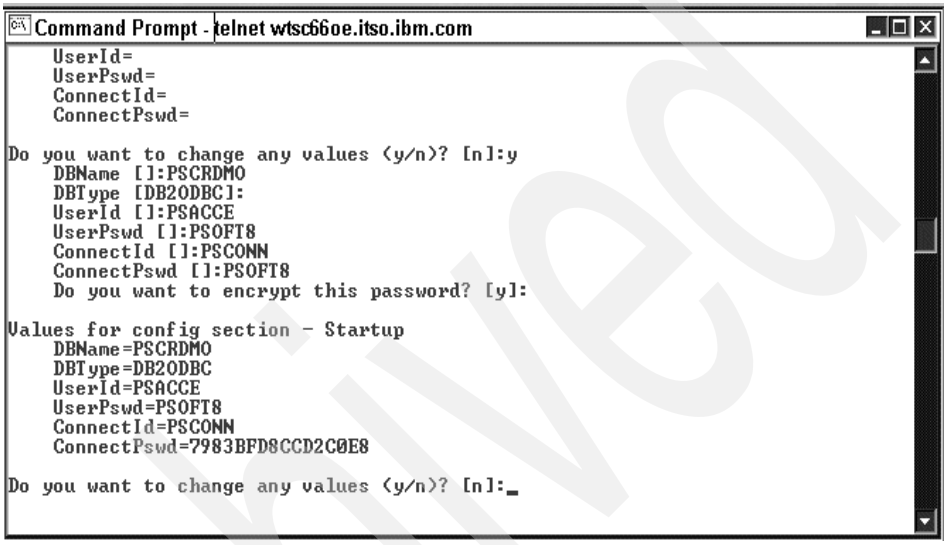

Figure 10-6 Configure process scheduler

In the z/OS block, enter the ODBC initialization file as shown in Figure 10-7.

Figure 10-7 ODBC initialization file

In the following blocks, make sure that Toolbin is configured in PTOOLS and that `DbFlags=1` (%UpdateStat disabled).

10. You can now test your process scheduler.
    a. Go to Start a Process Scheduler Server.
    b. Select your database in the Database List.

       It should indicate that your process scheduler is started.

However, our status on the last day was not good, as shown in Figure 10-8.

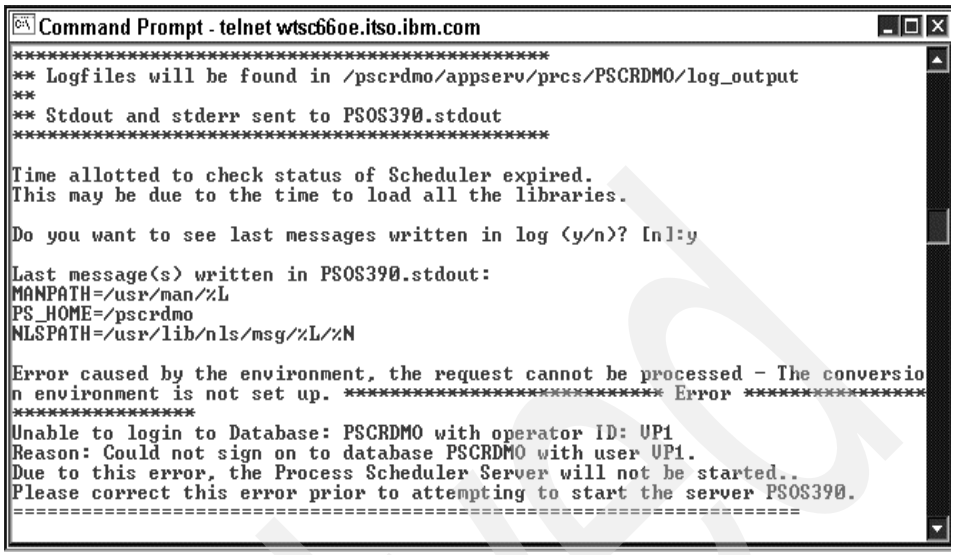

Figure 10-8 Error in starting process scheduler

11

Applying application fixes

This chapter outlines some considerations you should take into account when upgrading or applying fixes to PeopleSoft databases. It describes the method PeopleSoft uses to deliver product fixes and lists the steps you need to apply the different types of fixes.

The chapter also considers PeopleSoft upgrades.

It shows how the new Upgrade Assistant can improve upgrade quality and shorten the amount of time it takes to complete an upgrade. It also comments on other improvements in PeopleSoft V8 that assist the upgrade process.

11.1 Applying PeopleSoft patches

PeopleSoft delivers updates and fixes on a regular basis. Review the Customer Connection at:

http://www.peoplesoft.com

to check if there are any new updates for the product.

These instructions describe the general steps for different types of patches.For details about applying a specific patch, refer to the instructions for that patch.

Applying a patch requires these steps:

1. Verify that all prerequisite patches have been applied.
2. Download the patch files from Customer Connection.
3. Apply the patch to the system.

The software change made in an application patch may depend on other changes made in prior patches.Therefore, it is important to apply patches to the PeopleSoft system in the correct order. Each Updates & Fixes Report page in Customer Connection includes a Prerequisites section.This section lists any prior patches on which the current patch relies.You must make sure you have applied all the prerequisite patches before you apply the current patch.

For patches with multiple prerequisites, the Prerequisites section also specifies the correct order for applying the prior patches. If no specific order is given, apply the patches in the order of their posted (or reposted) date, starting with the earliest date.

Skipping prerequisites or applying them in the incorrect order may result in incompatible software with unpredictable problems.

The software changes associated with a patch are usually attached to the Updates & Fixes Report page as a self-extracting ZIP archive (.EXE) file. This archive file contains the various files you need to apply the patch. The .EXE file may also include supporting documentation, or the Updates & Fixes Report page may include a separate .EXE for the documentation. If the Updates & Fixes Report page includes more than one.EXE file, read the implementation instructions carefully to make sure to download the correct files.

Note: For some cases, the Prerequisites section lists only the immediate prerequisites for the current patch.Those prior patches may themselves have additional prerequisites. Before you apply any prerequisite patch, you must make sure that you have applied all of its prerequisites.

To install the patch, follow these steps:

1. Create an empty temporary directory on a local drive on the workstation or file server.
2. Download the.EXE file from the Updates & Fixes Report page to the temporary directory.
3. Run the executable file to extract the files that make up the patch.
4. Copy the extracted files to the appropriate locations.
5. Complete any additional steps specified on the Updates &Fixes Report page.
6. Delete the temporary directory.

Applying each type of file to the PeopleSoft system

The components delivered in a patch replace prior versions of the same components on the system. If any of these components were specially configured or customized, the changes must be re-applied to the new version of the component.

COBOL File (.CBL)*

The process for handling COBOL files differs depending on the database platform. In all cases, write access to the COBOL source and executable directories is required to compile COBOL modules.

For batch servers that reside on Windows NT and 2000,the.cbl file(s) should be copied to the <PS_HOME>\src \cbl \base directory and recompiled. For batch servers that reside on UNIX,the.cbl file(s) should be copied using FTP (ASCII mode) to the <PS_HOME>/src/cbl directory in the case-sensitive format of PSXXXXXX.cbl and recompiled.

Data Mover file (.DMS,*.DAT)*

There are three types of DMS files:

- Those that contain COBOL stored statements
- Those that contain database update scripts
- Those that contain Application Upgrade DAT file import scripts

The Updates & Fixes Report page should specify which type of DMS file each one is. To run the Data Mover scripts, follow these steps:

1. Copy any DMS files that contain COBOL stored statements into the directory <PS_HOME>\src \cbl \base, unless the Updates & Fixes Report page specifies a different directory.
2. Copy any DMS files that contain database or Application Upgrade DAT file import scripts into the *scripts* directory.

3. Copy all *.DAT files into the *data* directory.
4. If the set of newly delivered files includes Storexxx.DMS (where xxx refers to the relevant PeopleSoft application) then run that script using Data Mover. Store xxx. DMS reloads all stored statements, including those in the newly delivered files. Skip the next two steps.
5. If the set of newly delivered files does not include Storexxx.DMS, use Data Mover to run each newly delivered DMS script in the <PS_HOME>\src \cbl \base directory. Either run the script files individually, or add them to the existing Storexxx.DMS file and run Storexxx.DMS.
6. If the set of newly-delivered files does not include Storexxx.DMS, consult the documentation provided with the patch for information about when to run the DMS scripts that are copied into the *scripts* directory.

Remember to set the Data Mover Input and Output directory settings in the Configuration Manager before running the scripts.

Report File (.SQR,*.SQC,*.RPT,*.XNV)*

PeopleSoft delivered report files should be placed in the <PS_HOME>\sqr directory while any user-modified SQR should be placed in a separate directory.

Do the following:

1. Copy any *.SQR and *.SQC files to the <PS_HOME>\sqr directory on the file server.
2. Copy any *.RPT files to the <PS_HOME>\crw directory on the file server.
3. Copy any *.XNV files to the <PS_HOME>\nVision \Layouts directory on the file server, unless the Updates & Fixes Report page specifies a different location.
4. Copy any *.SQR and *.SQC files to the <PS_HOME>\sqr directory on the Windows NT/2000 batch server, or FTP any *.SQR and *.SQC files to the <PS_HOME>/sqr directory on the UNIX batch server.

Application upgrade project

The process for loading Application Upgrade projects into the PeopleSoft database differs between PeopleSoft V8 and previous releases. PeopleSoft V8 Application Upgrade projects are provided as a maintenance project. This is a set of files containing the definitions for a PeopleSoft project and its related database objects.

Note: For DMS files that contain Application Upgrade DAT file import scripts, refer to “Application upgrade project” on page 130.

PeopleSoft recommends that patches be initially applied to an application demo database. Once the fix is applied to the demo database, run an upgrade to compare with the development, test, or production databases, and identify whether the delivered patch has an impact on any modifications that have made. For information on running an upgrade compare, refer to the PeopleTools documentation.

The project file allows customers to quickly apply patches to the database directly through the Application Designer. Maintenance projects are standard PeopleTools upgrade projects with one additional feature. When a maintenance project is copied into a database, it inserts data about the applied patch into a database maintenance log for reference purposes.

To apply a maintenance project, complete these steps:

1. Log into the database where the project is to be applied (the application demo database in most cases).
2. Start Application Designer.
3. Select **File->Copy Project from File**.
4. In the dialog box that appears, navigate to the directory that contains the maintenance project.
5. Select the project file and modify the copy options if necessary.
6. Click the **Copy** button to load the project and its related objects into the database.
7. Complete any other steps specified on the Updates & Fixes Report page.

Recreating views and altering tables

Occasionally, the instructions for applying an application patch tell customers to recreate database views or alter a database table. This section provides an overview of how to perform these tasks. Read the documentation for the patch carefully before you make these database changes. Perform these steps only for objects explicitly identified in the instructions.

To recreate a view, perform these steps:

1. Start Application Designer.
2. Select **File ->Open**.
3. Enter or select **Record** as the Object Type to open.

4. Open the record representing the view to be recreated.
5. Select **Build->Current Object**.The Build dialog box appears.
6. Select the **Create Views** checkbox.

Note: To apply a maintenance project to a PeopleSoft V8 database, the change control must be turned off.

7. Select one of the Build Execute Options: **Build Script**, **Execute SQL now** or **Execute and Build Script**.
8. Click the **Build** button to perform the change or generate the script.
9. If **Build Script** is selected, review the resulting script file, and then run the file using an SQL editor.

To alter a table, follow these steps:

1. Start Application Designer.
2. Select **File ->Open**.
3. Enter or select **Record** as the Object Type to open.
4. Open the record representing the table to alter.
5. Select **Build->Current Object.** The Build dialog box appears.
6. Select the **Alter Tables** checkbox. The **Create Indexes** checkbox is automatically selected, as is the **Build Script** checkbox in the Build Execute Options box.
7. Click the **Build** button to generate the script.
8. Review the resulting script file, then run the file according to the platform and administrative requirements.

Related publications

How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

ibm.com/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the **CD-ROMs** button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

J

M

N

O

P

R

S

T

U

W

Z

Implementing PeopleSoft 8.4 on zSeries

(0.2"spine)
0.17"<->0.473"
90<->249 pages

Implementing PeopleSoft 8.4 on zSeries

Redbooks

Overview of PeopleSoft on zSeries

Installation of PeopleSoft 8.4

Configuration of PeopleSoft 8.4

This IBM Redbook will help you install PeopleSoft 8.4 on zSeries with DB2. This information is based on experience gained while installing PeopleSoft 8.4 at the IBM Silicon Valley Laboratory in San Jose, California and on customer support experiences.

This book will be especially useful for those who are installing and implementing PeopleSoft 8.4 DB2 on zSeries for the first time. A basic knowledge of DB2, Windows NT/2000, SQL is presumed.

Database administrators, system administrators and system designers who need to install PeopleSoft V8.4 using DB2 UDB V7.1 on UNIX and Windows NT/2000 platforms will find useful information pertaining to those tasks.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6549-00 ISBN 0738424587