



**PeerDirect™ Distributed Enterprise  
Configuration Guide  
for Progress® RDBMS**

While the PeerDirect Corporation believes the information included in this publication is correct as of the publication date, information in this document is subject to change without notice. The references in this manual to specific platforms supported are subject to change.

UNLESS EXPRESSLY SET FORTH IN A WRITTEN AGREEMENT SIGNED BY AN AUTHORIZED REPRESENTATIVE OF PEERDIRECT CORPORATION, PEERDIRECT MAKES NO WARRANTY OR REPRESENTATION OF ANY KIND WITH RESPECT TO THE INFORMATION CONTAINED HEREIN, INCLUDING WARRANTY OF MERCHANTABILITY AND FITNESS FOR A PURPOSE. PeerDirect assumes no responsibility or obligation of any kind for any errors contained herein or in connection with the furnishing, performance, or use of this document.

All contents of this manual are copyrighted by the PeerDirect Corporation, a Progress Software Corporation company. The information contained herein is the exclusive property of PeerDirect and shall not be copied, transferred, photocopied, translated on paper, film, electronic media, or computer-readable form, or otherwise reproduced in any way, without the express written permission of PeerDirect Corporation.

Software described in PeerDirect documentation (a) is the property of PeerDirect or a third party, (b) furnished only under license, and (c) may be copied or used only as expressly permitted under the terms of the applicable license.

PeerDirect, InnerEdge, OuterEdge, UltraSQL, and Voice of Experience are trademarks or registered trademarks of PeerDirect Corporation or its affiliate PeerDirect Company in the U.S. and other countries.

Allegrix, Leadership by Design, Object Design, ObjectStore, Progress, Powered by Progress, Progress Fast Track, Progress Profiles, Partners in Progress, Partners en Progress, Progress en Partners, Progress in Progress, P.I.P., Progress Results, ProVision, ProCare, ProtoSpeed, SmartBeans, SpeedScript, and WebSpeed are registered trademarks of Progress Software Corporation or one of its subsidiaries or affiliates in the U.S. and other countries. A Data Center of Your Very Own, Apptivity, AppsAlive, AppServer, ASPen, ASP-in-a-Box, BPM, Cache-Forward, Empowerment Center, eXcelon, EXLN, Fathom, Future Proof, inProgress for Partners, IntelliStream, Javlin, ObjectStore Browsers, OpenEdge, POSSE, POSSENET, Progress Dynamics, Progress Software Developers Network, RTEE, Schemadesigner, SectorAlliance, SmartBrowser, SmartComponent, SmartDataBrowser, SmartDataObjects, SmartDataView, SmartDialog, SmartFolder, SmartFrame, SmartObjects, SmartPanel, SmartQuery, SmartViewer, SmartWindow, Stylus, Stylus Studio, WebClient, Who Makes Progress, XIS, XIS Lite, and XPress are trademarks or service marks of Progress Software Corporation or one of its subsidiaries or affiliates in the U.S. and other countries.

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

Any other trademarks or service marks contained herein may be the property of their respective owners.

Perl is distributed pursuant to the applicable Artistic License, that is available as of the publication of this manual at [www.perl.com/language/misc/Artistic.html](http://www.perl.com/language/misc/Artistic.html), and all of the terms and conditions contained therein.

This PeerDirect product uses parts of the SFL package, Copyright © 1996-2000 iMatix Corp 1991-2000 iMatix Corporation <<http://www.imatix.com>>.

April 2003

# Contents

---

<b>Preface</b> .....	<b>ix</b>
Purpose .....	ix
Audience .....	ix
Organization of This Manual .....	ix
Typographical Conventions .....	x
Syntax Notation .....	xi
Progress Messages .....	xii
Progress Documentation .....	xiv
PeerDirect Distributed Enterprise Documentation .....	xv
<b>1. Introduction</b> .....	<b>1-1</b>
1.1 Introduction to PeerDirect Distributed Enterprise with Progress RDBMS .....	1-2
1.2 PeerDirect Distributed Enterprise Products .....	1-2
1.3 How the PeerDirect Products Work Together .....	1-3
1.4 Progress Products .....	1-4
1.5 Heterogeneous Database Products .....	1-5
1.6 PeerDirect Replication Network with All Products .....	1-5
1.7 Choosing a Layout .....	1-7
1.7.1 SDK Setup with Databases and PDRE on One Machine. ....	1-7
1.7.2 Deployment Database and PDRE on Separate Machines. ....	1-8
1.8 PeerDirect Distributed Enterprise Homogeneous Configuration .....	1-9
1.9 PeerDirect Distributed Enterprise Heterogeneous Configuration .....	1-11
1.10 Progress Fathom High Availability Replication with PeerDirect Distributed Enterprise .....	1-12

<b>2.</b>	<b>Setup in a Progress Environment</b>	<b>2-1</b>
2.1	Overview	2-2
2.1.1	Setting Up Replication on an InnerEdge Server	2-2
2.1.2	Setting Up Replication on an OuterEdge Server or Workstation	2-3
2.1.3	Setup Task Sequence	2-4
2.1.4	Recommendations	2-9
2.2	Progress Database Considerations	2-9
2.3	Setting Up a Progress Database for PeerDirect Replication	2-11
2.3.1	Dumping Table Definitions	2-11
2.3.2	Loading Table Definitions	2-11
2.3.3	Setting Up Database Logins and Permissions	2-12
2.3.4	Selecting Progress Databases to Replicate	2-14
2.3.5	Selecting Tables for Replication	2-15
2.3.6	Specifying Foreign Keys	2-15
2.3.7	Defining Fragments	2-16
2.3.8	Defining Work Sets	2-17
2.3.9	Defining Transaction Sets and Groups	2-17
2.3.10	Verifying and Saving Replication Rules	2-18
2.3.11	Preparing Progress Databases for Activation	2-19
2.3.12	Preparing to Activate with PRD	2-19
2.4	Activating Your Replication Sites	2-20
2.4.1	Preparing to Activate an InnerEdge Server	2-20
2.4.2	Activating the InnerEdge Server Progress Database	2-22
2.4.3	Activating OuterEdge Server Sites	2-25
2.5	Initiating a Replication Session	2-27
2.6	Administering a Replication Network	2-28
<b>3.</b>	<b>Setup in a Heterogeneous Environment</b>	<b>3-1</b>
3.1	Overview	3-2
3.1.1	Setting Up Replication on an InnerEdge Server	3-3
3.1.2	Setting Up Replication on an OuterEdge Server or Workstation	3-4
3.1.3	Setup Task Sequence	3-4
3.1.4	Recommendations	3-10
3.2	Database Considerations	3-10
3.3	Setting Up Databases for Replication	3-12
3.3.1	Setting Up Database Logins and Permissions	3-12
3.3.2	Selecting Databases to Replicate	3-15
3.3.3	Selecting Tables for Replication	3-16
3.3.4	Specifying Foreign Keys	3-17
3.3.5	Defining Fragments	3-18
3.3.6	Defining Worksets	3-18
3.3.7	Defining Transaction Sets and Groups	3-19
3.3.8	Verifying and Saving Replication Rules	3-20

3.4	Verifying Heterogeneous Database Compatibility .....	3–20
3.5	About Preparing Databases for Activation .....	3–21
	3.5.1 Preparing to Activate Your Replication Sites .....	3–22
	3.5.2 Preparing to Activate with PRD .....	3–23
	3.5.3 Activating the InnerEdge Server Database .....	3–24
	3.5.4 Preparing and Activating the OuterEdge Server Sites.....	3–25
3.6	Initiating a Replication Session .....	3–29
3.7	Administering a Replication Network .....	3–29
<b>Index</b>	.....	<b>Index–1</b>

**Figures**

Figure 1–1:	PeerDirect Replication Network with All Products . . . . .	1–6
Figure 1–2:	PeerDirect with Progress Database on Windows Only . . . . .	1–7
Figure 1–3:	Database and PDRE on Separate Machines . . . . .	1–8
Figure 1–4:	PeerDirect Distributed Enterprise Homogeneous Configuration . . . . .	1–10
Figure 1–5:	PeerDirect Distributed Enterprise Heterogeneous Configuration . . . . .	1–11
Figure 1–6:	Progress Fathom High Availability Replication . . . . .	1–12

---

**Tables**

Table 2–1:	Sequence of Replication Network Setup Tasks (Method A) . . . . .	2–5
Table 2–2:	Sequence of Replication Network Setup Tasks (Method B) . . . . .	2–7
Table 2–3:	PDR and PDRE Access Permissions . . . . .	2–12
Table 3–1:	Sequence of Replication Network Setup Tasks (Method A) . . . . .	3–5
Table 3–2:	Sequence of Replication Network Setup Tasks (Method B) . . . . .	3–7
Table 3–3:	PRD and PDRE Access Permissions . . . . .	3–12
Table 3–4:	Database Type Formats . . . . .	3–15



# Preface

---

## Purpose

This book documents how to configure PeerDirect™ Distributed Enterprise for the Progress® RDBMS in environments that have only the Progress database type (homogeneous) or in environments that have both the Progress database type and other database types (heterogeneous).

## Audience

This book is intended for database administrators (DBAs) and application developers who plan to use the PeerDirect Distributed Enterprise product suite with the Progress database or application. It assumes a knowledge of database design, deployment, and administration of a Progress database.

## Organization of This Manual

The chapters in this manual are set up as follows:

### [Chapter 1, “Introduction”](#)

Introduces the PeerDirect Distributed Enterprise product suite to Progress database users. It provides a visual and textual framework so that a DBA and application developer can understand the value, requirements, and expectations set if using the PeerDirect Distributed Enterprise product suites with the Progress database.

### [Chapter 2, “Setup in a Progress Environment”](#)

Takes the application developer through the steps to create rules that are exported to a file and applied against the Progress database. This chapter includes steps to deploy a Progress database and initiate replication.

### Chapter 3, “Setup in a Heterogeneous Environment”

Takes the application developer through the steps to create rules that are exported to a file and applied against the Progress database and other database types. This chapter also includes steps to deploy a database and initiate replication.

## Typographical Conventions

This manual uses the following typographical conventions:

- **Bold typeface** indicates:
  - Commands or characters that the user types
  - That a word carries particular weight or emphasis
- *Italic typeface* indicates:
  - Progress variable information that the user supplies
  - New terms
  - Titles of complete publications
- Monospaced typeface indicates:
  - Code examples
  - System output
  - Operating system filenames and pathnames

The following typographical conventions are used to represent keystrokes:

- Small capitals are used for Progress key functions and generic keyboard keys.

**END-ERROR, GET, GO**  
**ALT, CTRL, SPACEBAR, TAB**

- When you have to press a combination of keys, they are joined by a dash. You press and hold down the first key, then press the second key.

**CTRL-X**

- When you have to press and release one key, then press another key, the key names are separated with a space.

**ESCAPE H**

**ESCAPE CURSOR-LEFT**

## Syntax Notation

The following sections describe the conventions for the components in syntax notation.

### Uppercase Words

Uppercase words are keywords. Although they are always shown in uppercase, you can use either uppercase or lowercase when using them in a procedure.

In this example, **ACCUM** is a keyword:

### **SYNTAX**

```
ACCUM aggregate expression
```

### Italics

Italics identify options or arguments that you must supply. These options can be defined as part of the syntax or in a separate syntax identified by the name in italics.

In the **ACCUM** function above, the *aggregate* and *expression* options are defined with the syntax for the **ACCUM** function in the *Progress Language Reference*.

### Period or Colon

All statements (except for **DO**, **FOR**, **FUNCTION**, **PROCEDURE**, and **REPEAT**) end with a period. **DO**, **FOR**, **FUNCTION**, **PROCEDURE**, and **REPEAT** statements can end with either a period or a colon, as in this example:

```
FOR EACH Customer:  
  DISPLAY Name.  
END.
```

## Square Brackets

Square brackets ( [ ] ) around an item indicate that the item, or a choice of one of the enclosed items, is optional.

In this example, `STREAM stream`, `UNLESS-HIDDEN`, and `NO-ERROR` are optional:

### SYNTAX

```
DISPLAY [ STREAM stream ] [ UNLESS-HIDDEN ] [ NO-ERROR ]
```

**NOTE:** In some instances, square brackets are not a syntax notation, but part of the language.

For example, this syntax for the `INITIAL` option uses brackets to bound an initial value list for an array variable definition. In these cases, normal text brackets ( [ ] ) are used:

### SYNTAX

```
INITIAL [ constant [ , constant ] . . . ]
```

**NOTE:** The ellipsis ( . . . ) indicates repetition, as shown in a following description.

## Progress Messages

Progress displays several types of messages to inform you of routine and unusual occurrences:

- Execution messages inform you of errors encountered while Progress is running a procedure (for example, if Progress cannot find a record with a specified index field value).
- Compile messages inform you of errors found while Progress is reading and analyzing a procedure prior to running it (for example, if a procedure references a table name that is not defined in the database).
- Startup messages inform you of unusual conditions detected while Progress is getting ready to execute (for example, if you entered an invalid startup parameter).

After displaying a message, Progress proceeds in one of several ways:

- Continues execution, subject to the error-processing actions that you specify, or that are assumed, as part of the procedure. This is the most common action taken following execution messages.
- Returns to the Progress Procedure Editor so that you can correct an error in a procedure. This is the usual action taken following compiler messages.

- Halts processing of a procedure and returns immediately to the Procedure Editor. This does not happen often.
- Terminates the current session.

Progress messages end with a message number in parentheses. In this example, the message number is 200:

```
** Unknown table name table. (200)
```

On the Windows platform, use Progress online help to get more information about Progress messages. Many Progress tools include the following Help menu options to provide information about messages:

- Choose Help→Recent Messages to display detailed descriptions of the most recent Progress message and all other messages returned in the current session.
- Choose Help→Messages, then enter the message number to display a description of any Progress message. (If you encounter an error that terminates Progress, make a note of the message number before restarting.)
- In the Procedure Editor, press the **HELP** key (**F2** or **CTRL-W**).

On the UNIX platform, use the Progress **PRO** command to start a single-user mode character Progress client session and view a brief description of a message by providing its number. Follow these steps:

- 1 ♦ Start the Progress Procedure Editor:

```
install-dir/dlc/bin/pro
```

- 2 ♦ Press **F3** to access the menu bar, then choose Help→Messages.
- 3 ♦ Type the message number, and press **ENTER**. Details about that message number appear.
- 4 ♦ Press **F4** to close the message, press **F3** to access the Procedure Editor menu, and choose File→Exit.

## Progress Documentation

### Getting Started

#### *Progress Installation and Configuration Guide Version 9 for UNIX*

A manual that describes how to install and set up Progress Version 9.1 for the UNIX operating system.

#### *Progress Installation and Configuration Guide Version 9 for Windows*

A manual that describes how to install and set up Progress Version 9.1 for all supported Windows and Citrix MetaFrame operating systems.

### Database

#### *Progress Database Design Guide*

A guide that uses a sample database and the Progress Data Dictionary to illustrate the fundamental principles of relational database design. Topics include relationships, normalization, indexing, and database triggers.

#### *Progress Database Administration Guide and Reference*

This guide describes Progress database administration concepts and procedures. The procedures allow you to create and maintain your Progress databases and manage their performance.

### SQL-92

#### *Progress Embedded SQL-92 Guide and Reference*

A guide to Progress Embedded SQL-92 for C, including step-by-step instructions for building ESQL-92 applications and reference information about all Embedded SQL-92 Preprocessor statements and supporting function calls. This guide also describes the relationship between ESQL-92 and the ANSI standards upon which it is based.

#### *Progress JDBC Driver Guide*

A guide to the Java Database Connectivity (JDBC) interface and the Progress SQL-92 JDBC driver. It describes how to set up and use the driver and details the driver's support for the JDBC interface.

*Progress ODBC Driver Guide*

A guide to the ODBC interface and the Progress SQL-92 ODBC driver. It describes how to set up and use the driver and details the driver's support for the ODBC interface.

*Progress SQL-92 Guide and Reference*

A user guide and reference for programmers who use Progress SQL-92. It includes information on all supported SQL-92 statements, SQL-92 Data Manipulation Language components, SQL-92 Data Definition Language components, and Progress functions. The guide describes how to use the Progress SQL-92 Java classes and how to create and use Java stored procedures and triggers.

**Reference***Pocket Progress (Hard copy only)*

A reference that lets you quickly look up information about the Progress language or programming environment.

**PeerDirect Distributed Enterprise Documentation***PeerDirect Distributed Enterprise Developer's Guide*

A guide for replication system developers that describes how to plan and implement a PeerDirect Replication Network. Topics include how to plan database schemas, define replication rules, customize events, and deploy replication network sites. It includes several reference sections for quick lookup of APIs and other information.

*PeerDirect Distributed Enterprise SDK Installation Guide*

A guide that describes the procedure for installing the PeerDirect Distributed Enterprise SDK. It also describes how to define replication rules and set up replication sites in a PeerDirect Replication Network.

*PeerDirect InnerEdge Server Installation Guide*

A guide that describes the procedure for installing the PeerDirect InnerEdge™ Server. It also contains instructions for setting up and configuring this product as a site in a PeerDirect Replication Network.

*PeerDirect OuterEdge Server Installation Guide*

A guide that describes the procedure for installing the PeerDirect OuterEdge™ Server and Workstation. It also contains instructions for setting up and configuring this product as a site in a PeerDirect Replication Network.

*PeerDirect Replication Engine Developer's Guide and Reference*

An online guide that provides basic concepts and instructions for planning your replication project and defining replication rules. It also describes how to use the PeerDirect Replication Engine APIs, utilities, and event handling.

---

## Introduction

This chapter gives an introduction to the PeerDirect™ products and configurations for Progress® application developers and database administrators. It covers the following topics:

- [Introduction to PeerDirect Distributed Enterprise with Progress RDBMS](#)
- [PeerDirect Distributed Enterprise Products](#)
- [How the PeerDirect Products Work Together](#)
- [Progress Products](#)
- [Heterogeneous Database Products](#)
- [PeerDirect Replication Network with All Products](#)
- [Choosing a Layout](#)
- [PeerDirect Distributed Enterprise Homogeneous Configuration](#)
- [PeerDirect Distributed Enterprise Heterogeneous Configuration](#)
- [Progress Fathom High Availability Replication with PeerDirect Distributed Enterprise](#)

## 1.1 Introduction to PeerDirect Distributed Enterprise with Progress RDBMS

This guide presents the steps necessary to configure PeerDirect products in a Progress environment. PeerDirect Distributed Enterprise software is a system for replicating and synchronizing data stored across databases. You can use the PeerDirect Distributed Enterprise products to replicate data between Progress databases (known here as a homogeneous environment) or between a Progress database and a different database type, such as Oracle (known here as a heterogeneous environment).

## 1.2 PeerDirect Distributed Enterprise Products

There are several different PeerDirect Distributed Enterprise products you can install. You can install the PeerDirect Distributed Enterprise SDK, the PeerDirect InnerEdge Server, and the PeerDirect OuterEdge Server and Workstation. Each product serves a specific functional purpose as follows:

- **PeerDirect Distributed Enterprise SDK** — Provides a development environment for building and testing a PeerDirect Replication Network. The PeerDirect Distributed Enterprise SDK software installs everything you need to configure InnerEdge and OuterEdge Servers or Workstations. This is a required product for developing a PeerDirect Replication Network.

The PeerDirect Distributed Enterprise SDK is the only product that contains the PeerDirect Replication Designer (PRD). The PRD is a graphical user interface tool on Windows for defining your replication sites and rules. Replication rules define what and how the data is to be synchronized between two or more sites within a PeerDirect Replication Network. Sites contain a Progress database and other supported databases in a PeerDirect Replication Network.

- **PeerDirect InnerEdge Servers** — Support high-availability sites where 24X7 access to applications and their data is critical. The primary node in a PeerDirect Replication Network is a PeerDirect InnerEdge Server, and it contains a complete copy of replicated data. The first PeerDirect InnerEdge Server is also referred to as the *first* site in a replication network, and therefore is installed first and is generally used as the target of the first replication session for each subsequent site. The first InnerEdge Server is a required product for all PeerDirect customers. Additional InnerEdge Servers are recommended for high availability.

- **PeerDirect OuterEdge Servers** — Are remote replication sites on a PeerDirect Replication Network. They provide the ability to replicate data from a central application center to outer-edge databases, giving remote users access to their applications even when connectivity to the home office is interrupted. This is an optional product for PeerDirect customers, but it is recommended as a remote replication site.
- **PeerDirect OuterEdge Workstation** — Is a remote replication site for a single-user laptop on a PeerDirect Replication Network. Each OuterEdge Workstation provides the ability to replicate data from a central application center to outer-edge databases, giving a single user access to his application even when connectivity to the home office is interrupted. This is an optional product for PeerDirect customers, but is recommended for remote users.

All of the PeerDirect Distributed Enterprise products include the PeerDirect Replication Engine (PDRE). The PDRE replicates all data stored in databases that have been set up to replicate among activated sites in a PeerDirect Replication network. At predefined intervals, PDRE gathers the changes made to the Progress or other supported database since the last replication session, and propagates these changes to other PDRE sites in the background over TCP/IP.

All of the PeerDirect Distributed Enterprise products include the PeerDirect Administrator (PDA). PDA is a graphical user interface on Windows for administering a replication network. Administrative tasks include subscribing and unsubscribing to and from work set slices and setting up replication schedules.

A *work set* is a connected set of related tables rooted at a base table. Each base record in the base table identifies a slice of the Progress or other supported database that sites can request to subscribe to. This allows clients to see and replicate only that information in the database that they are subscribed to. Given a specific record from the work set base table, the *work set slice* comprises that base table record along with the set of related records within the tables defined to belong to that work set. Thus, each base table record identifies a specific slice of data that can be subscribed to or unsubscribed from.

The PeerDirect Administrator provides for setting up replication schedules to happen automatically.

### 1.3 How the PeerDirect Products Work Together

The InnerEdge Server and OuterEdge Server or Workstation products install one instance of the PDRE per license. Every database that PeerDirect Distributed Enterprise supports has its own make of an InnerEdge Server and OuterEdge Server or Workstation. Therefore, in an environment where there are different database types (heterogeneous), the customer might install the PeerDirect InnerEdge Server for the Oracle database, and on their remote sites install an OuterEdge Server for the Progress database.

PeerDirect products provide for the creation of a *PeerDirect Replication Network* for distributing application data. The PeerDirect Replication Network determines which site databases will be replicated, in other words, which distributed database data will be routinely updated and synchronized with database data in central data centers. All of these databases share identical data schemas and are configured to recognize each other. Each replication network is defined by the data that is replicated, not by sites themselves. So, one replication network can be set up to replicate tables A-D and another to replicate table Z, where both table sets are stored in the same database.

To set up a PeerDirect Replication Network, the customer must first install the PeerDirect Distributed Enterprise SDK in a development environment. Critical to the activation of each database site is the existence of a `.rul` file, which is created using the PeerDirect Replication Designer (PRD). The developer exports the `.rul` file and tests it before activating each database. Once the databases are activated, the PRD is not used again until there are changes in the rules or database schema changes that impact the replication network.

The system administrator can install the InnerEdge Servers and OuterEdge Servers or Workstations before or after the `.rul` file is created. The requirement is that the `.rul` file must be ready prior to database preparation for activation using the DIMPRULE utility. Once the database is prepared, the database can be activated and replication initiated from the remote site as each site activation completes. Now the user can use the Administrator to schedule replication and subscribe work set slices. The Administrator only runs on Windows. For this reason, if Linux users want to use the Administrator, they must install it on a Windows machine that can remotely communicate with the Linux machine that is running the PRDE.

For specific OS requirements for the Administrator and other PeerDirect Distributed Enterprise components, see the PeerDirect installation guides.

## 1.4 Progress Products

There are several different Progress products you can install, including the Progress Enterprise, Workgroup, or Personal databases. In addition, you must have SQL-92 installed as follows:

- **Progress RDBMS Version 9.1D with Service Pack 9.1D05 or later** — With the PeerDirect product suite, you can configure your PeerDirect Replication Network to include the Progress Enterprise database, Workgroup database, or Personal database. When you install the Progress database, you must include the Progress SQL-92 Explorer for communications between the Progress database and the PDRE. Because Progress SQL-92 is not supported on SCO or Solaris Intel, the Progress database cannot reside on either of these platforms. The Progress database is required for at least one of your sites in a PeerDirect Replication Network.

- **Progress SQL-92 Client Access** — This product resides on the same machine as the PDRE to provide communications between the PDRE and the Progress database. This product is only required if your PDRE is on a different machine than your Progress database.
- **ODBC Drivers** — The ODBC drivers are used for SQL-92 communications between the PDRE and the Progress database. See the *Progress Release Notes* for more information. This product is required.
- **Progress® Fathom™ High Availability, Replication** — Provides a hot standby from the InnerEdge Server database, also known as a source database, to a target database should a recovery database be required in case of an outage. This product is recommended.

For more information on installing Progress products, see the *Progress Installation and Configuration Guide for Windows* and the corresponding guide for UNIX provided with your Progress documentation.

## 1.5 Heterogeneous Database Products

In addition to the Progress database, PeerDirect Distributed Enterprise products also support the configuration of more than one database type in the PeerDirect Replication Network. The following additional database types are supported for a heterogeneous environment where at least one of your databases is a Progress database:

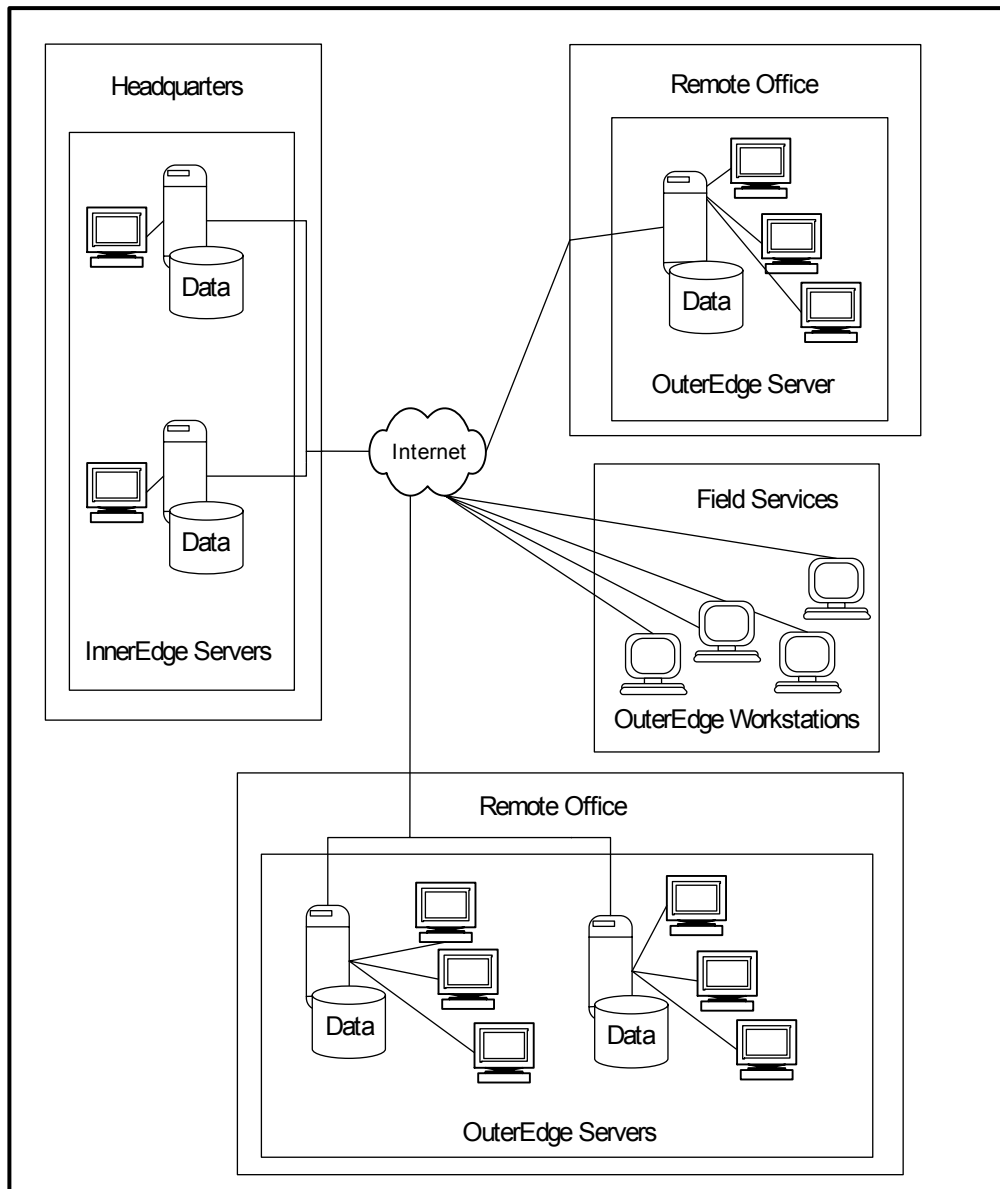
- IBM DB2 Universal Database
- Microsoft® SQL Server™
- Oracle

## 1.6 PeerDirect Replication Network with All Products

PeerDirect OuterEdge Servers and Workstations are remote replication sites on a PeerDirect Replication Network. A PeerDirect Replication Network provides the ability to replicate data from a central application to outer-edge databases, such as Progress databases, giving remote users access to their applications even when connectivity to the home office is interrupted.

The PeerDirect Replication Network is a configuration of InnerEdge and OuterEdge Servers that support distributed application management. An application runs locally using data stored in local Progress, or other supported databases, and thus is always available. In the background, the PDRE manages changes to your application by replicating and synchronizing data updates with other applications you have designated as active sites in the network.

Figure 1-1 shows a replication network that consists of three OuterEdge Servers set up in remote field or office locations.



**Figure 1-1: PeerDirect Replication Network with All Products**

You can set up a PeerDirect Replication Network in several different ways, but for greatest benefit, you should select one that matches the business requirements for distributing

application data within your system. Each PeerDirect network manages the database for one or more applications (or unrelated database tables) that you wish to connect across your organization, and there is no limit to the number of replication networks you can set up.

## 1.7 Choosing a Layout

Before you configure your PeerDirect Distributed Enterprise products, you must choose how to deploy your PeerDirect products in a PeerDirect Replication Network. The PeerDirect Replication Network can be set up in many ways, and you should choose the deployment that best suits your business needs. The PeerDirect SDK environment allows you to create your .rul file and test it before deployment. The [“SDK Setup with Databases and PDRE on One Machine”](#) section shows a PeerDirect SDK layout as well as a sample deployment environment.

### 1.7.1 SDK Setup with Databases and PDRE on One Machine

In this scenario, everything you need to define rules and test them on multiple databases is on one machine. This configuration limits the platform to Windows because the PRD and PDA require Windows. If you use a Linux machine for your PDRE, you need to install the PRD and PDA in Windows with a remote connection to Linux, which requires two machines. In a testing and application development environment, the single machine configuration is useful since it is also the easiest one to set up. [Figure 1–2](#) shows a Progress Enterprise Database connected through PDRE to a second database where another Progress database resides. The application developer can use the PRD to design rules and test them in this configuration.

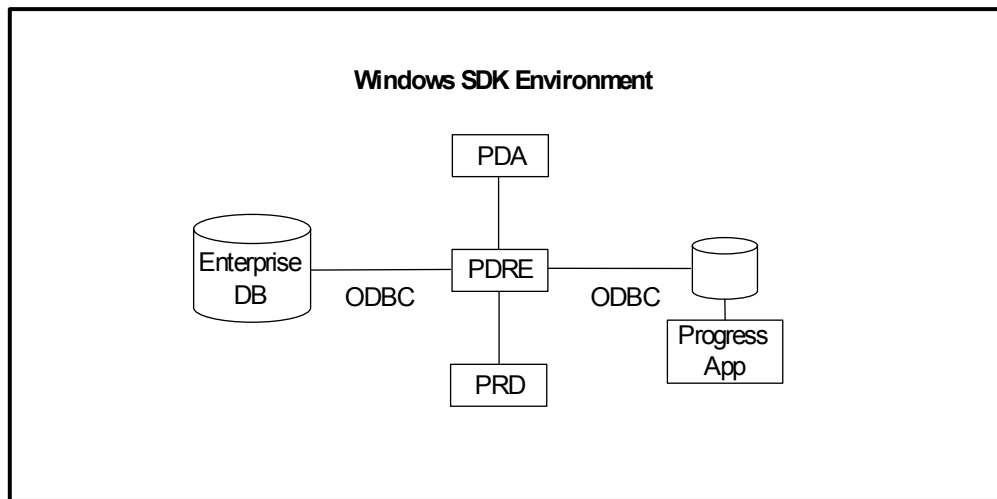


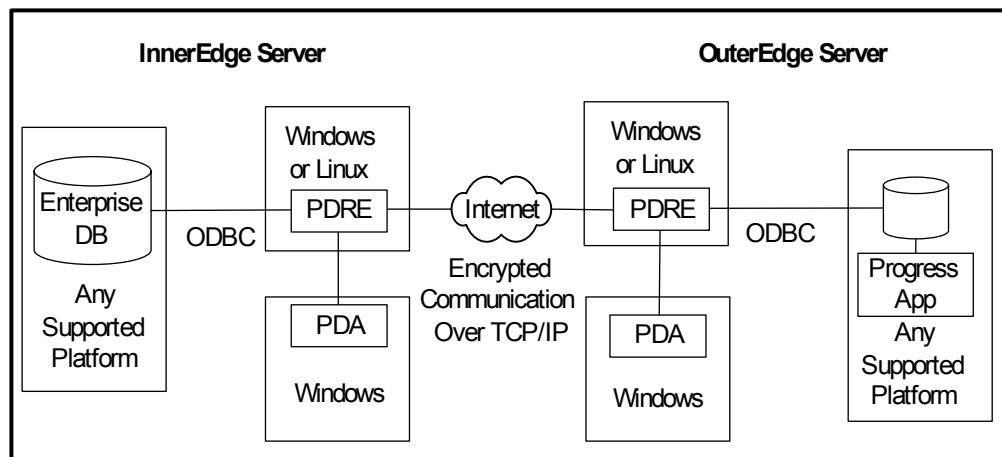
Figure 1–2: PeerDirect with Progress Database on Windows Only

In [Figure 1–2](#), a Progress database and the PDRE use ODBC to communicate, which is why Progress SQL-92 Access is required. Since this entire configuration is on a Windows operating system, the application developer has all the necessary tools on one machine to design a PeerDirect Replication network.

### 1.7.2 Deployment Database and PDRE on Separate Machines

In this deployment scenario, there is an InnerEdge Server and an OuterEdge Server separated by the Internet. One machine has a database and a second machine has a copy of the PDRE installed at a site where there is either an InnerEdge or an OuterEdge Server. There are two sites and at each site there is one PDRE for each database. Each copy of the PDRE interacts with its database as well as the PDRE at the other site. This scenario mimics a typical WAN deployment in which two regional servers exist in separate LANs protected by firewalls. [Figure 1–3](#) shows a total of six machines, two on each LAN, which have a database on one machine and an instance of PDRE on the other.

The scenario in [Figure 1–3](#) allows for the databases to reside on any platform in which they are supported.



**Figure 1–3: Database and PDRE on Separate Machines**

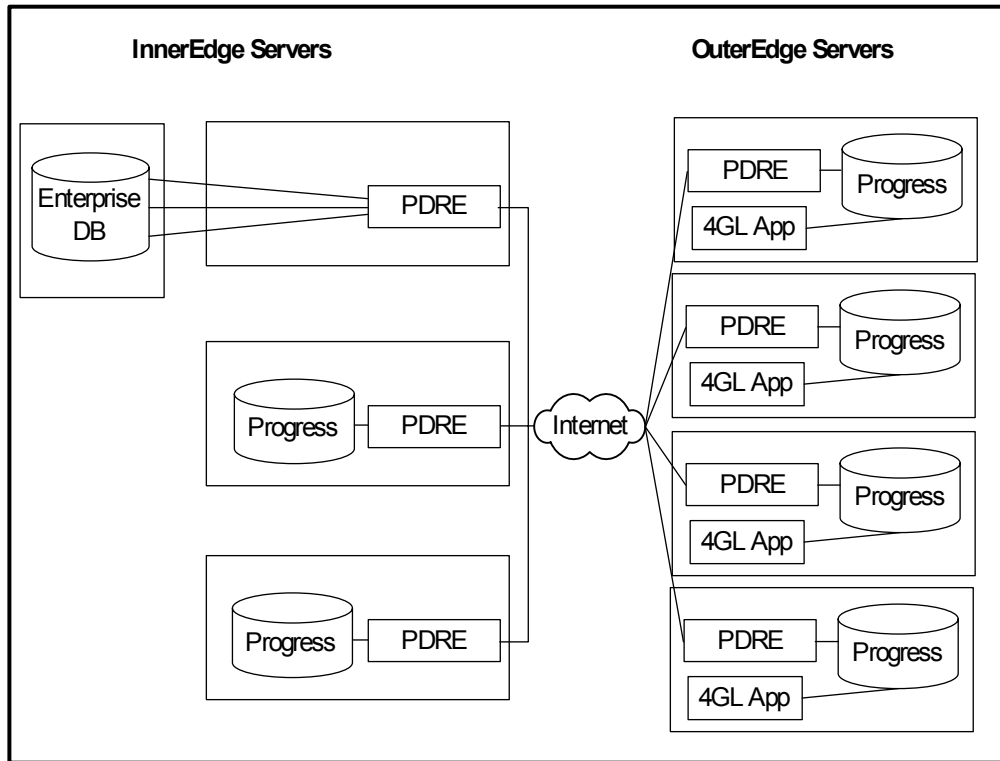
[Figure 1–3](#) is a homogeneous environment because there are only Progress databases in the configuration. In a PeerDirect Replication Network, one of these databases could just as easily be a different database type as long as it is one of the supported databases. The only Progress requirement in a heterogeneous environment is that one of the databases must be a Progress database. This Progress database can even reside on the OuterEdge Server machine.

## 1.8 PeerDirect Distributed Enterprise Homogeneous Configuration

In a homogeneous environment, replication occurs between two or more Progress databases. [Figure 1-4](#) shows one example of how an application developer might deploy a PeerDirect Replication Network using only Progress databases. The Enterprise database with a complete copy of the database resides on its own machine to which the InnerEdge Server PDRE communicates. Having the InnerEdge Server PDRE on a separate machine provides more resources for the database and the InnerEdge Server PDRE, as they are not competing on the same machine for resources. Also, since PDRE must reside on either Windows or Red Hat Linux, having the Enterprise database on a separate machine increases the choices of operating systems that you can use for your database.

However, there are also benefits to having the database as close as possible to the PDRE. The setup is more secure and the communications between the database and PDRE might be faster.

Also note that each PDRE connection uses two connections of a Progress database license. It is important for the application developer to take into account the two connections used for each PDRE connection when determining the number of Progress users licences needed in a PeerDirect Replication Network that uses a Progress database.



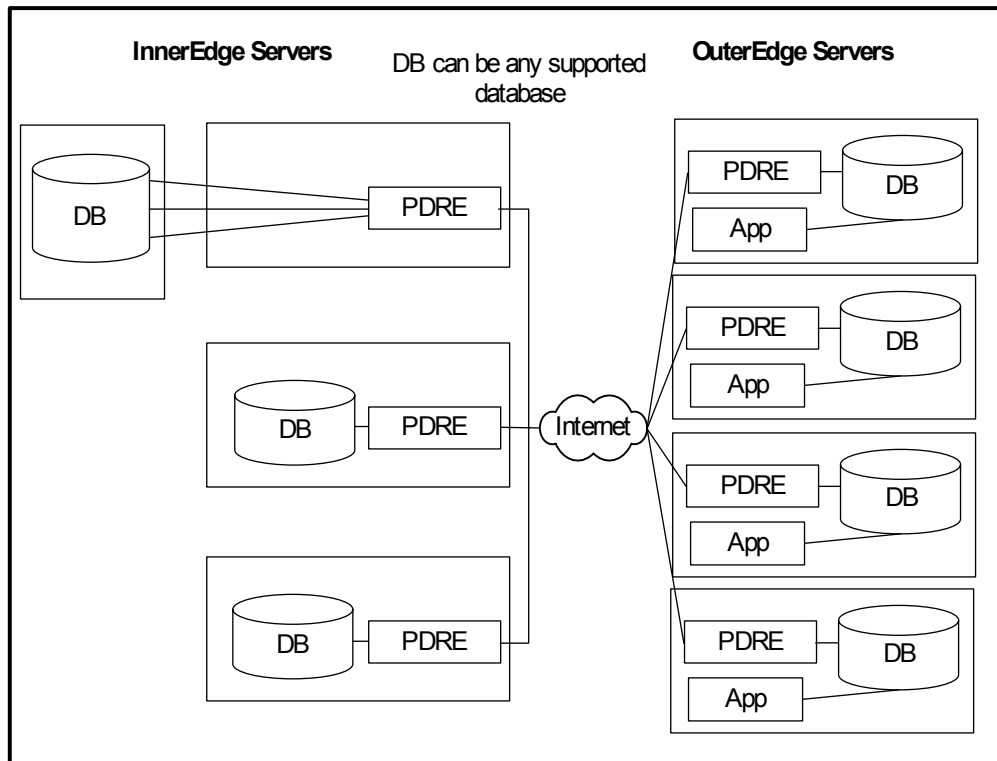
**Figure 1–4: PeerDirect Distributed Enterprise Homogeneous Configuration**

As shown in [Figure 1–4](#), there are three InnerEdge servers with an instance of PDRE on each machine. By having multiple InnerEdge Servers, you increase the number of sites an OuterEdge Servers can replicate to, as the PeerDirect Replication Network has multiple sites with a complete copy of the database. This can improve overall performance and provide a warm backup to the first InnerEdge server should this be required. With multiple InnerEdge Servers, if one InnerEdge Server becomes unavailable, other InnerEdge Servers can offload the replication during the interruption of service to the first InnerEdge Server.

[Figure 1–4](#) also shows multiple OuterEdge Servers, giving remote users access to their applications even when connectivity to the home office is interrupted.

## 1.9 PeerDirect Distributed Enterprise Heterogeneous Configuration

In a heterogeneous environment, replication occurs between at least one Progress database and another supported database. For a list of supported databases, see the “[Heterogeneous Database Products](#)” section in this chapter. [Figure 1–5](#) shows one example of a heterogeneous configuration.

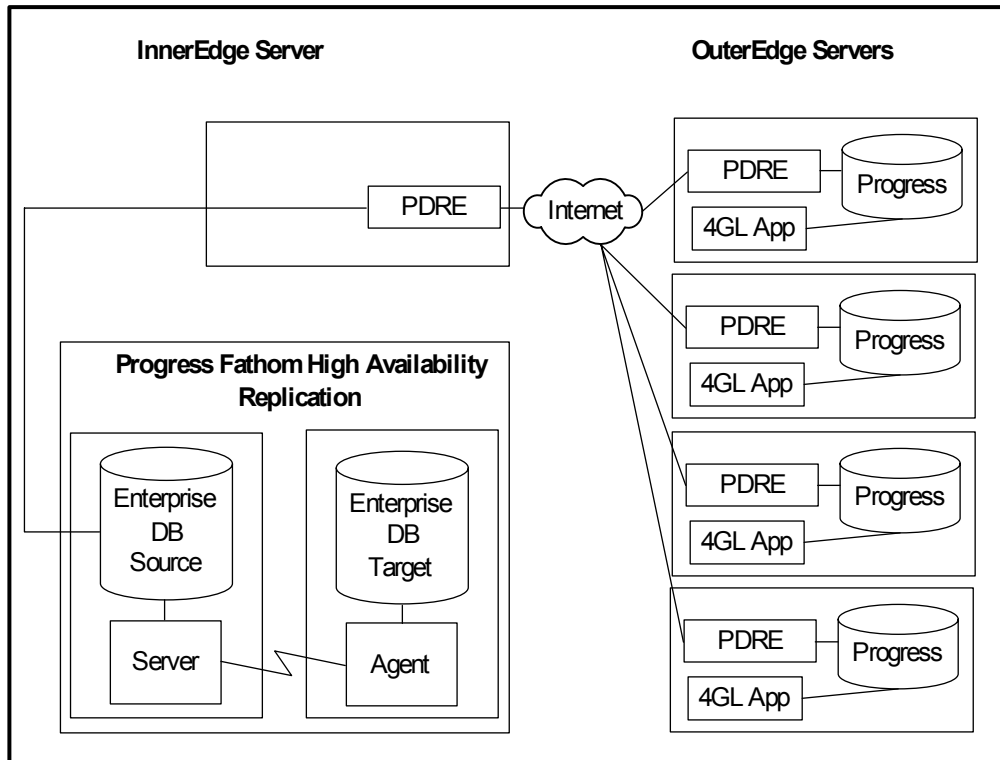


**Figure 1–5: PeerDirect Distributed Enterprise Heterogeneous Configuration**

The only Progress product requirement in a heterogeneous configuration is that one of the databases must be a Progress database. Therefore, in [Figure 1–5](#), one of the DB types is a Progress database.

## 1.10 Progress Fathom High Availability Replication with PeerDirect Distributed Enterprise

The first database in your PeerDirect Replication Network contains a complete database to which other PDREs replicate. If this database is a Progress database, you can provide for a warm-standby should a recovery database be required by using the Progress Fathom High Availability Replication product. Figure 1–6 shows a PeerDirect Replication Network with Progress Fathom High Availability Replication providing for a warm-standby of the InnerEdge complete database should one be needed.



**Figure 1–6: Progress Fathom High Availability Replication**

For more information on Progress Fathom High Availability Replication, you can access the *Progress Fathom High Availability User's Guide* at the Progress Products Documentation Web Site, <http://www.progress.com/v9/documentation/index.htm>.

---

## Setup in a Progress Environment

This chapter describes how to set up PeerDirect replication in a Progress environment using your PeerDirect SDK software. This chapter takes you through each phase of the PeerDirect Replication Engine setup and includes these sections:

- [Overview](#)
- [Progress Database Considerations](#)
- [Setting Up a Progress Database for PeerDirect Replication](#)
- [Activating Your Replication Sites](#)
- [Initiating a Replication Session](#)
- [Administering a Replication Network](#)

## 2.1 Overview

All successful replication networks begin with a deployment strategy and plan. Begin by considering the different layout options that PeerDirect offers and choose the one that offers you the benefits and features you need.

The prerequisite setup for Progress requires that you have installed your Progress database and SQL-92 Server Engine on the database machines. If you install PeerDirect on a different machine than where your Progress database is located, you must install SQL-92 Client Access on your PeerDirect machine to get the ODBC driver for Progress. For detailed requirements for the prerequisite Progress products, see the “[Progress Products](#)” section in [Chapter 1, “Introduction,”](#) of this guide. For installation instructions on Progress products, see the *Progress Installation and Configuration Guide for Windows* and the corresponding guide for UNIX provided with your Progress documentation.

The initial setup of a replication network requires that you install the PRD available in the Windows PeerDirect SDK product. You need the PRD to define the replication rules for the database. These rules allow the PDRE to know what and how the data is to be synchronized between two or more sites within a PeerDirect Replication Network. You can then do the following:

- Export the replication rules you define to a file and apply it against an InnerEdge Server using the DIMPRULE utility.
- Back up and restore this database to other sites.
- Activate the first database and subsequent databases. To activate a database, you use the DINST utility. The DIMPRULE and DINST utilities are present in both the Windows and Linux versions of PDRE.

### 2.1.1 Setting Up Replication on an InnerEdge Server

You must define replication rules *before* you can complete setup of the first InnerEdge Server site. To set up replication on the first Inner Edge Server, follow these steps:

- 1 ♦ Install the PeerDirect Distributed Enterprise SDK software and use the PDR to create replication rules. See the *PeerDirect Distributed Enterprise SDK Installation Guide* and the *PeerDirect Distributed Enterprise Developer’s Guide* for information about how to create rules.
- 2 ♦ Install an InnerEdge Server running PDRE on a primary datacenter server that hosts the database with which you plan to replicate OuterEdge Servers. See the *PeerDirect InnerEdge Server Installation Guide* for installation instructions.

- 3 ♦ Create a development database using the Progress dump and load utilities to set up a database with only data definition files.
- 4 ♦ Add a PDUSER account through Progress SQL-92 Explorer to the database with the data definitions file loaded.
- 5 ♦ Connect to the PRD using the database with only the data definitions file loaded.
- 6 ♦ Choose the tables to replicate.
- 7 ♦ Verify and save your replication rules.
- 8 ♦ Prepare to activate your database with the data definitions files and export your rules to a .rul file.
- 9 ♦ Prepare to deploy the first site by enabling PeerDirect replication using the PROUTIL utility, adding the PDUSER account, running DIMPRULE utility to load rules, and making a backup copy of this database to deploy to other sites.
- 10 ♦ Using the DINST utility, activate the InnerEdge Server site at your first site.
- 11 ♦ Restore your first site database to remaining Inner Edge Server site databases and use the DINST utility to activate them.

**NOTE:** You must activate the site on the InnerEdge Server first and then activate all remaining InnerEdge and OuterEdge sites.

### **2.1.2 Setting Up Replication on an OuterEdge Server or Workstation**

You should already have installed one InnerEdge Server on a primary server that hosts the database you plan to use to replicate data. If you chose Method A (see [Table 2-1](#)), you should also have completed the preparation of your first InnerEdge Server database and activated the site.

If you chose Method B (see [Table 2-2](#)), after you install both the InnerEdge and OuterEdge Servers, the next step is to complete the setup of the InnerEdge Server and then set up the OuterEdge Server.

To set up replication on an OuterEdge Server or Workstation, follow these steps:

- 1 ♦ Install an OuterEdge Server or Workstation that runs PDRE. See the *PeerDirect OuterEdge Server Installation Guide* for installation instructions.
- 2 ♦ Deploy the OuterEdge Server or Workstation sites by restoring the backup copy of the first InnerEdge Server site database and running the DINST utility to activate them. See the “[Activating OuterEdge Server Sites](#)” section in this chapter for more information.

After the first replication, you can use the PDA console to set up the PeerDirect Replication Engine to automatically replicate data on a schedule.

### 2.1.3 Setup Task Sequence

Setting up a PeerDirect Replication Network requires that you perform a sequence of tasks. There are two ways to perform this sequence. You can choose to set up your replication sites one by one, beginning with your first InnerEdge Server (see [Table 2-1](#)) or you can perform all your installation tasks first and then prepare and activate each site separately (see [Table 2-2](#)). You should choose the method that you think works best for your particular deployment purposes. [Table 2-1](#) and [Table 2-2](#) describe the two methods. For detailed instructions, see the referenced resources.

Table 2–1 gives a sequence of steps to follow to set up your PeerDirect Replication Network.

**Table 2–1: Sequence of Replication Network Setup Tasks (Method A) (1 of 2)**

Step	Task	Description
1	Plan Replication Network.	<p>Plan and design your distributed application system and replication network as an up-front activity. Review your database configuration and schema design to make sure it meets the PeerDirect requirements for replication of data.</p> <p>You will find some typical replication network layouts described in the “Choosing a Layout” section in <a href="#">Chapter 1, “Introduction,”</a> of this guide.</p>
2	Install the Distributed Enterprise SDK (Windows version).	<p>When you install the SDK, you must also install the PRD to create the database rules for your PeerDirect Replication Network.</p> <p>The PRD is only available as part of the PeerDirect Distributed Enterprise SDK software product.</p> <p>See the <i>PeerDirect Distributed Enterprise SDK Installation Guide</i> for detailed instructions on installing the SDK.</p>
3	Define Replication Rules.	<p>Use the PRD to define the database replication rules for your PeerDirect Replication Network. This is where the work you did in Step 1 pays off.</p> <p>See the <i>PeerDirect Distributed Enterprise Developer’s Guide</i>.</p>
4	Export Replication Rules.	<p>After you have created replication rules and have verified them, select <b>Tools &gt; Prepare to Activate</b> and then <b>File &gt; Export Rules</b>. This exports the rules to a .ru1 file that you use to set up your replication databases at each site.</p> <p>See the “Preparing to Activate with PRD” section in this chapter for more information.</p>
5	Install the InnerEdge Server as the first site.	<p>Install the InnerEdge Server on a central data server that is highly available to the other sites in your PeerDirect Replication Network.</p> <p>See the <i>PeerDirect InnerEdge Server Installation Guide</i>.</p>

**Table 2–1: Sequence of Replication Network Setup Tasks (Method A) (2 of 2)**

Step	Task	Description
6	Prepare the InnerEdge Server database.	<p>Prepare the first site database using the DIMPRULE utility. Be sure to use the -FIRSTSITE parameter.</p> <p>See the <a href="#">“Activating Your Replication Sites”</a> section in this chapter for more information.</p>
7	Install an InnerEdge Server, OuterEdge Server, or Workstation as a second site.	<p>After you install and prepare to activate your first site, install your second replication site.</p> <p>Later, you will repeat this step for the other sites in your replication network.</p> <p>See the <i>PeerDirect InnerEdge Server Installation Guide</i> or <i>PeerDirect OuterEdge Server and Workstation Installation Guide</i>.</p>
8	Back up the first site database and restore it to the second site.	<p>Back up and restore of the first database to other sites. You do not need to rerun the DIMPRULE utility.</p> <p>Later, you will repeat the restore step for all remaining sites you have installed as part of your replication network.</p>
9	Activate the InnerEdge Server site.	<p>Activate the InnerEdge Server site using the DINST utility. Be sure to use the -FIRSTSITE parameter on the first InnerEdge Server site.</p> <p>See the <a href="#">“Activating the InnerEdge Server Progress Database”</a> section in this chapter for more information.</p>
10	Activate the InnerEdge Server, OuterEdge Server, or Workstation site as a second site.	<p>Activate the second site using the DINST utility. As part of this process, the initial replication session is started with the first site.</p> <p>Later, you will repeat this step for all remaining sites.</p> <p>See the <a href="#">“Activating OuterEdge Server Sites”</a> section in this chapter for more information.</p>
11	Repeat Steps 7–10.	<p>Activate the remaining sites in your PeerDirect Replication Network.</p>

Table 2–2 gives an alternate sequence of steps to follow to set up your PeerDirect Replication Network.

**Table 2–2: Sequence of Replication Network Setup Tasks (Method B) (1 of 2)**

Step	Task	Description
1	Plan Replication Network.	<p>Plan and design your distributed application system and replication network as an up-front activity. Review your database configuration and schema design to make sure it meets the PeerDirect requirements for replication of data.</p> <p>You will find typical replication network layouts described in <a href="#">“Choosing a Layout”</a> section in <a href="#">Chapter 1, “Introduction,”</a> of this guide.</p>
2	Install the Distributed Enterprise SDK (Windows version).	<p>You must install the PRD to create the database rules for your PeerDirect Replication Network.</p> <p>The PRD is only available as part of the PeerDirect Distributed Enterprise SDK software product.</p> <p>See the <i>PeerDirect Distributed Enterprise SDK Installation Guide</i> for detailed instructions on installing the SDK.</p>
3	Define Replication Rules.	<p>Use the PRD to define the database replication rules for your PeerDirect Replication Network. This is where the work you did in Step 1 pays off.</p> <p>See the <i>PeerDirect Distributed Enterprise Developer’s Guide</i>.</p>
4	Export Replication Rules to a file.	<p>After you have created replication rules and have verified them, select <b>Tools &gt; Prepare to Activate</b> and then <b>File &gt; Export Rules</b>. This exports the rules to a .rul file that you use to set up your replication databases at each site.</p> <p>See the <a href="#">“Preparing to Activate with PRD”</a> section in this chapter for more information.</p>
5	Install the InnerEdge Server as the first site.	<p>Install the InnerEdge Server on a central data server that is highly available to the other sites in your PeerDirect Replication Network.</p> <p>See the <i>PeerDirect InnerEdge Server Installation Guide</i> for more information.</p>

**Table 2–2: Sequence of Replication Network Setup Tasks (Method B) (2 of 2)**

Step	Task	Description
6	Install an InnerEdge Server, OuterEdge Server, or Workstation on the second and remaining sites in the network.	<p>After you install and prepare to activate your first site, install your second replication site.</p> <p>Repeat this step for the other sites in your replication network.</p> <p>See the <i>PeerDirect InnerEdge Server Installation Guide</i> or <i>PeerDirect OuterEdge Server and Workstation Installation Guide</i> for more information.</p>
7	Prepare the InnerEdge Server database as the first site.	<p>Prepare the first site database using the DIMPRULE utility. Be sure to use the <code>-FIRSTSITE</code> parameter.</p> <p>See the “<a href="#">Activating Your Replication Sites</a>” section in this chapter for more information.</p>
8	Back up the first site database and restore it to the second site.	<p>Back up and restore of the first database to other sites. You do not need to rerun the DIMPRULE utility.</p> <p>Later, you will repeat the restore step for all remaining sites you have installed as part of your replication network.</p>
9	Activate the InnerEdge Server site.	<p>Activate the InnerEdge Server site using the DINST utility. Be sure to use the <code>-FIRSTSITE</code> parameter on the first InnerEdge Server site.</p> <p>See the “<a href="#">Activating the InnerEdge Server Progress Database</a>” section in this chapter for more information.</p>
10	Activate the InnerEdge Server, OuterEdge Server, or Workstation site as a second site.	<p>Activate the second site using the DINST utility. As part of this process, the initial replication session is started with the first site.</p> <p>Later, you will repeat this step for all remaining sites.</p> <p>See the “<a href="#">Activating OuterEdge Server Sites</a>” section in this chapter for more information.</p>
11	Repeat Steps 7–10.	<p>Activate the remaining sites in your PeerDirect Replication Network.</p>

Each InnerEdge Server, OuterEdge Server, or Workstation runs a PDRE component that replicates data among the databases in your PeerDirect Replication Network. You can select to install PDRE on your server or workstation when you run the PeerDirect install program (on **Windows**, select the **Runtime** edition).

For instructions on how to install an OuterEdge Server, see the *PeerDirect OuterEdge Server Installation Guide*. For instructions on how to install an InnerEdge Server, see the *PeerDirect InnerEdge Server Installation Guide*.

## 2.1.4 Recommendations

You will find instructions in this chapter for setting up replication on both the PeerDirect InnerEdge and OuterEdge Servers. As you work, it is important that you follow all the instructions in the sequence in which they are presented. Do not skip any steps.

## 2.2 Progress Database Considerations

When replicating databases with the PeerDirect Replication Engine, you should pay attention to the recommendations described in the following sections.

### Foreign Keys

Foreign keys are important if you want to impact the order in which tables are replicated. If you plan to use work sets or transaction sets, you **must** define foreign keys.

It is important that PDRE know the foreign keys for each table involved in replication. Database table relationships are defined from the bottom up, that is from the first child to the parent to which it is related by a foreign key. PDRE uses this information to determine the order in which tables are replicated, as parent tables are always replicated before child tables. Foreign key relationships are also used in defining work sets and transaction sets, and to ensure that deletions are cascaded to dependent tables.

For instructions about how to define foreign keys, see the [“Specifying Foreign Keys”](#) section.

### Cascading Table Deletions

If your database includes local (nonreplicated) tables that are related to replicated tables, cascading deletions will not extend to these local tables. This means that you must either use triggers to ensure that deletions are cascaded to local tables or enable cascading deletions in your database software or in your business application.

Note that if you implement referential integrity constraints between these local and replicated tables without ensuring that deletions are properly cascaded to the local tables, PDRE deletions will cause these constraints to “break.”

## Consistent Naming

Ensure that all tables you plan to replicate have the same names across replicated databases. This is not a problem if you enable a first database and then copy it to other locations for replication.

Columns in replicated tables must also have identical names, and when you are replicating between two different vendor databases, the columns must have compatible data types.

There are no requirements for tables or columns that you do not plan to replicate.

## Unsupported Data Types

Tables you plan to replicate should not include unsupported data types or associated settings or properties. By default, PDRE replicates all columns in a table, but you can modify this behavior using a feature called fragments. For more information on fragments, see the “[Defining Fragments](#)” section in this chapter.

For a list of supported data types and database-specific types, see the *PeerDirect Distributed Enterprise Developer’s Guide*.

## Unique Primary Keys

All tables you plan to replicate must have a unique primary key. Primary keys can be *simple*—comprised of a single column—or *composite*—comprised of two or more columns. However, no component of a primary key can contain a NULL value.

Primary keys for the Progress database must be set up in the Progress database. The PRD will import them. They cannot be identified in the PRD. However, other database types can use the PRD to identify primary keys.

## Database Versions

PDRE supports Progress Version 9.1D with Service Pack V9.1D05 or later. Any changes to this requirement can be found in the Release Notes.

## Clock Synchronization

The date and time must be set correctly on all machines in a replication network, including all database servers, PDRE servers, application servers (machines running database clients), and remotely connected machines. You should install and use one of the following programs to synchronize the clocks on your machines:

- Tardis (<http://www.kaska.demon.co.uk>)
- Dimension 4 (<http://www.thinkman.com/dimension4/>)

## 2.3 Setting Up a Progress Database for PeerDirect Replication

To set up a PeerDirect Replication Network for Progress, you must perform a series of tasks. The sequence of these tasks is outlined in the “[Overview](#)” section of this chapter. The sections below describe each task in detail.

### 2.3.1 Dumping Table Definitions

To facilitate the development necessary to prepare a Progress database for deployment in a PeerDirect Replication Network, the application developer should dump the data definitions from an existing database and load those definitions into a new database. Doing this saves time because all that is needed to define rules are the data definitions. To create a development database for PeerDirect configuration, use the Progress dump and load utilities. Follow these steps to dump database or table definitions:

- 1 ♦ Access the appropriate data tool (the Data Administration tool if you are using a graphical interface or the Data Dictionary if you are using a character interface).
- 2 ♦ Make sure you are connected to the database (the source) from which you want to dump the table definitions.
- 3 ♦ Choose **Admin**→ **Dump Data and Definitions**→ **Data Definitions (.df file)**. The Data Administration tool lists all the tables defined for the database alphabetically.
- 4 ♦ Choose **All**.
- 5 ♦ Accept the default filename or enter a different name and choose **OK**. The Data Administration tool displays each object name as it writes its definition to the data definitions file.
- 6 ♦ Choose **OK** to return to the Data Administration main window.

### 2.3.2 Loading Table Definitions

After creating an empty database, follow these steps to load table definitions into your empty database:

- 1 ♦ Access the appropriate data tool (the Data Administration tool if you are using a graphical interface or the Data Dictionary if you are using a character interface).
- 2 ♦ Make sure you are connected to the database (the target) in which you want to load the table definitions.

- 3 ♦ Choose **Admin**→ **Load Data and Definitions**→ **Data Definitions (.df files)**. The Data tool prompts you for the name of the file that contains the data definitions you want to load into the current database. The default filename is the logical name of the current working database, with a .df extension.
- 4 ♦ Specify a file or accept the default.
- 5 ♦ Specify whether you want to stop the load as soon as the Data tool encounters a bad definition statement. The Data tool displays each item as it loads the definitions for that object. The Data tool displays a status message and prompts you to continue. If you choose not to stop on errors, the load will continue with the next entry.

**NOTE:** Whatever you choose, if the Data tool encounters any errors, it backs out any loaded definitions.

- 6 ♦ Choose **OK** to return to the Data Administration or Data Dictionary main window.

The database now contains the table, field, index, or sequence, but none of the data. You can use this database to design your PeerDirect replication rules.

### 2.3.3 Setting Up Database Logins and Permissions

Both PRD and PDRE require access permissions that allow them to manipulate tables in the Progress database. [Table 2-3](#) shows the permissions required for the PRD and PDRE.

**Table 2-3: PDR and PDRE Access Permissions**

Each installed...	Requires...
PRD	Access to all application tables to be replicated. Permissions to create and update new tables in the database.
PDRE	Insert, update, delete, and select access to all replicated tables, all PeerDirect system tables, and all PeerDirect control tables in the database. Exclusive use of a specific login ID if you plan to implement triggers to maintain control fields.

If your database requires users to log in, you will need to provide login IDs that have appropriate permissions to access and update your database with PRD and PDRE.

## Creating a Login ID with Adequate Permissions on Your Progress Database

To create a login ID with adequate permissions, follow these steps:

- 1 ♦ Start a server on the development database:

```
PROSERVE dbname -S 3333
```

where *dbname* is the name of your development database with only the definitions files loaded.

- 2 ♦ Create the PDUSER account.

The account does not have to be named PDUSER. Any account name can be used as long as it is granted the appropriate permissions.

Only create a PDUSER account if there **are already active accounts** in the database. To determine if there are active accounts, connect to the database. If you are prompted for a user name and password, this indicates there are already active accounts in the database.

Create the PDUSER account by running the SQL Explorer and connecting to the database as the user that owns the database. You might need to create a SYSPROGRESS account through the 4GL dictionary to be able to connect to the database in some cases. Once connected, enter the following commands:

```
create user 'PDUSER', 'PASSWORD';  
grant resource, dba to 'PDUSER';  
commit work;
```

If there **are no active accounts** for the database, then enter the following commands:

```
grant resource, dba to 'PDUSER';  
commit work;
```

### 2.3.4 Selecting Progress Databases to Replicate

Use the PRD to select and configure the Progress database you want to replicate as part of a replication network. In this case, the Progress database will only contain data definitions without the data, since you are using the database where you loaded the data definitions file (.df).

When you start PRD for the first time, PRD asks you detailed questions about the Progress database, connection, and replication information.

#### Selecting a Progress Database to Replicate

Follow these steps to select a database to replicate:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→ **Programs**→ **PeerDirect**→ **PDRE**→ **Designer**.
- 2 ♦ The Database Information screen appears. Complete the following fields:
  - a) Database Type — Select the Progress database that contains the schema you have loaded.
  - b) Database Locator — Enter the location of the database you want to replicate using the format *host:port:dbname*, where *host* is the machine where your database is located, *port* is the port number to connect to, and *dbname* is the name of your Progress database with the data definition files in it.
  - c) Database Connection Information — (optional) If you need a user/login ID and password to access your database through ODBC, enter the following information:

**NOTE:** The user name PDUSER must have already been granted the necessary permissions to the database.

- UserName — Enter the user name or login ID required to access the database.
  - Password — Enter the password required to access the database.
  - Optional Parameters — Enter any additional connection parameters. For example, if your database uses a system database to validate users, you must enter its name here. The syntax is `systemdb=myfile.mdw`.
- 3 ♦ Click **Next**.
  - 4 ♦ The Database Information screen appears. Complete the following fields:
    - a) Project — Enter the name of your project.

- b) **Network** — Enter the name of the replication network. The network will be comprised of the databases you identify for replication.

**NOTE:** Only databases that are part of the same project and network can replicate together. Databases in two different networks will not interact.

- c) **Release** — Enter the name of this release. Each project/network can have multiple releases as new replication features are added or database schema changes are implemented.

5 ♦ Click **Next**.

### 2.3.5 Selecting Tables for Replication

Use the main PRD window to specify which tables to replicate (tables are initially marked as “Not Replicated”) and to configure the advanced features of PDRE.

Follow these steps to select a table for replication:

- 1 ♦ In the main PRD window, select the tables in your Progress database that you want to replicate by clicking on the table listed in the Not Replicated tab. Select multiple tables by holding down the **CTRL** or **SHIFT** keys.
- 2 ♦ Choose **Actions** → **Replicate Table**.
- 3 ♦ Click the **Replicated** tab at the bottom of the left panel. The tables marked for replication are shown in the replicated list.

### 2.3.6 Specifying Foreign Keys

With Progress, you must manually define foreign keys in the PRD. The PRD will import foreign key references defined in a Progress database.

It is important that PDRE know what the foreign keys are for each table it replicates. Foreign keys determine the order in which PDRE replicates tables, replicating parent tables before child tables. Foreign keys are also required to define work sets and transaction sets. See the “[Foreign Keys](#)” section in this chapter for more information.

Follow these steps to specify a foreign key relationship:

- 1 ♦ In the main PRD window, click the **Foreign Key** tab.
- 2 ♦ In the Create new relationship from <TableName> to parent table field, select a parent table for this child table from the drop-down menu. The primary key fields for the parent table and candidate foreign key fields appear.

- 3 ♦ Select the appropriate foreign key fields that correspond to the primary key fields.
- 4 ♦ Click **Add Relationship**. The relationship is added to the list of **Existing Relationships**.

Repeat these steps for each foreign key in your database.

Follow these steps to modify a foreign key relationship:

- 1 ♦ Select a relationship from the Existing Relationships list.
- 2 ♦ Select the primary key and foreign key fields from the drop-down menu.
- 3 ♦ Click **Apply Changes**.

To delete a relationship, select the relationship from the Existing Relationships list, and click **Delete Relationship**.

### 2.3.7 Defining Fragments

You can define fragments that allow changes to flow and merge when they do not occur in the same group of data. *Fragments* are the smallest unit of replication. By grouping fields that should change together into fragments, the application developer can prevent false collisions and detect real ones. You can create or add to a fragment by following these steps in the PRD:

- 1 ♦ Select a table in the table selection window.
- 2 ♦ Select the **Structure** tab.
- 3 ♦ Click the left mouse button in a field that you want in a new fragment.
- 4 ♦ From the Actions menu, select **Create Fragment**.
- 5 ♦ Enter the name of the fragment.
- 6 ♦ Choose **OK**. The table name appears in the Fragment column of the Structure tab.

### 2.3.8 Defining Work Sets

PeerDirect Distributed Enterprise allows you to define the business rules for subsetting base tables. This subsetting is known as work sets. Follow these steps to define a work set:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→ **Programs**→ **PeerDirect**→ **PDRE**→ **Designer**.
- 2 ♦ Select a base table in the Table Selection window, then right-click and select **Create Work Set**.
- 3 ♦ Select the **Work Set** tab.
- 4 ♦ To add a table to a work set, select the parent table in the Work Set window.
- 5 ♦ Select the table you wish to add to the work set from the Related Tables list.
- 6 ♦ Select the **Add to Work Set**.

### 2.3.9 Defining Transaction Sets and Groups

PeerDirect Distributed Enterprise allows you to define the business rules for data integrity of transactions. A *transaction* is a set of insert, update, or delete operations bound together and treated as a single unit of work. Defining rules for transactions is known as transaction set and transaction group definition. A *transaction set* is a group of records across related tables. A *transaction group* is one or more tables in their entirety. Follow these steps to define a transaction set:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→ **Programs**→ **PeerDirect**→ **PDRE**→ **Designer**.
- 2 ♦ Make sure the Transaction tab is selected.
- 3 ♦ Select a base table in the **Table Selection** window, then right-click and select **Create Transaction Set**.

Follow these steps to define a transaction group:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→ **Programs**→ **PeerDirect**→ **PDRE**→ **Designer**.
- 2 ♦ Make sure the Transaction tab is selected.
- 3 ♦ At the bottom of the Transaction tab, select the **Group** tab.

- 4 ♦ Select a table from the table drop down list.
- 5 ♦ Click on the **Create New Group** button.
- 6 ♦ The New Transaction Group dialog appears. Enter a name for the group.
- 7 ♦ Click **OK**.
- 8 ♦ In the Existing Tables Not In a Transaction, you can click on a table and select the **Add to Group** button to add the table to the Transaction Group.

### 2.3.10 Verifying and Saving Replication Rules

The next step is to verify and save your Progress database's replication rules. During verification the Progress database performs checks for cyclic dependencies, such as circular foreign key references, and whether each replicated table has a primary key specified. During the verification, the error messages are generated one at a time.

Follow these steps to verify and save Progress Database replication rules:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→**Programs**→**PeerDirect**→**PDRE**→**Designer**.
- 2 ♦ In the main PRD window, choose **Tools > Verify Rules**.
- 3 ♦ If PRD finds errors in your Progress database, they are displayed. Correct them and rerun the verify tool.
- 4 ♦ Repeat Step 2 until all error messages are resolved.
- 5 ♦ Choose **File**→**Save**.

**NOTE:** PRD lets you save your changes even if there are errors, but you cannot continue with the next steps unless you correct all errors. To view the error messages, select **View**→**Verification Messages**.

It is good practice at this point to export the script file of your rules for future development.

### 2.3.11 Preparing Progress Databases for Activation

PeerDirect recommends that you run PRD in your SDK development environment to create replication rules, and then import these rules to the first site of your production environment (which should be an InnerEdge Server) using the DIMPRULE utility (see the “[Activating Your Replication Sites](#)” section in this chapter for more information). Be sure to specify the DIMPRULE -FIRSTSITE parameter at the first InnerEdge Server site.

If you use PRD to activate your first site, you must be sure to check the **Mark starter data for replication** check box for InnerEdge Servers. All data in the OuterEdge databases will be initialized during a special *starter data reconciliation phase* that occurs during the initial replication with the InnerEdge Server database.

### 2.3.12 Preparing to Activate with PRD

To prepare to activate, you export a .ru1 file. The PRD uses this file to set up and activate other servers. It is important that you create the .ru1 file for the first site you install so that when you do a backup and restore of this site, you propagate the rules to each deployment site.

Follow these steps to prepare to activate your InnerEdge Server Progress database:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start→ Programs→ PeerDirect→ PDRE→ Designer**
- 2 ♦ Choose **Tools→ Prepare to Activate....**
- 3 ♦ If you know the network address of the InnerEdge Server site in your replication network, enter it. If you are using single engine replication, you can type **localhost**.
- 4 ♦ If you have initial data in the Progress database (starter data), leave the **Mark starter data for replication** check box checked.
- 5 ♦ Click **OK**.
- 6 ♦ A dialog box appears telling you that the database has been prepared for activation.
- 7 ♦ Click **OK**.

After you have successfully completed the Prepare to Activate operation, you must export the replication rules to a file. The default file extension is `.rul`. Follow these steps to export rules:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start** → **Programs** → **PeerDirect** → **PDRE** → **Designer**
- 2 ♦ Choose **File** → **Export Rules** and save the replication rules to a file, such as `xPD060102.rul`. You must use this rule file to set up your replication sites.
- 3 ♦ Exit the PeerDirect Replication Designer.

## 2.4 Activating Your Replication Sites

Begin by preparing your first InnerEdge Server for activation. Then activate your InnerEdge Server, OuterEdge Server, and Workstation sites as described in the following sections.

### 2.4.1 Preparing to Activate an InnerEdge Server

To prepare to activate your InnerEdge Server you must enable your database and run `DIMPRULE` on the first InnerEdge Server site by following these steps:

- 1 ♦ Open a command-line prompt window. Progress suggests using `PROENV` for this.
- 2 ♦ In a command-line shell or window, enable the Progress change capture mechanism for replication on the **deployment** database by entering the following command:

```
PROUTIL dbname -C enablepdr
```

where *dbname* is the name of your deployment Progress database. This database contains the data definitions and the data.

**NOTE:** You must enable the change capture mechanism for replication with the `proutil` command before you create your `PDUSER` account.

- 3 ♦ Start the server on your actual deployment database. When you start the server, turn on networking and set a large enough lock table size to accommodate the largest table in the database.
- 4 ♦ Create the `PDUSER` account.

The account does not have to be named `PDUSER`. Any account name can be used as long as it is granted the appropriate permissions.

Only create a PDUSER account if there **are already active accounts** in the database. To determine if there are active accounts, connect to the database. If you are prompted for a user name and password, this indicates there are already active accounts in the database.

Create the PDUSER account by running the SQL Explorer and connecting to the database as the user that owns the database. You might need to create a SYSPROGRESS account through the 4GL dictionary to be able to connect to the database in some cases. Once connected, enter the following commands:

```
create user 'PDUSER', 'PASSWORD';
grant resource, dba to 'PDUSER';
commit work;
```

If there **are no active accounts** for the database, then enter the following commands:

```
grant resource, dba to 'PDUSER';
commit work;
```

- 5 ♦ Type the following command to import your rules into your actual deployment database at the first site:

```
DIMPRULE -FIRSTSITE -DA=db_type -DB=db_locator -ODBCN=username
-ODBCP=password rule_file_name.ru1 [-logwide] [-v]
```

where

*db\_type* is the type of database to be activated and will always be the word Progress in a Progress environment.

*db\_locator* is the location of the deployment database. For Progress, the format for *db\_locator* is *host:port:db\_name*, where *host* is the machine name, *port* is the port for connection, and *db\_name* is the name of the database.

*username* is the user name you used in the PRD. You should use *PDUSER*. See the [“Setting Up Database Logins and Permissions”](#) section.

*password* is the password that corresponds to the user name above. The parameter is optional and defaults to "password". Be sure to enclose the password in double-quotes if it contains spaces.

*rule\_file\_name.ru1* is the name of the file you created in the PRD when you set up the rules for your first InnerEdge Server. You can get the name of this file from that site.

-logwide and -V (verbose) are optional parameters that enable detailed logging activity.

-FIRSTSITE designates this site to be the first site in the replication network and is used at only one site per replication network.

**NOTE:** The -ODBCN and -ODBCP parameters represent the user ID and password that PDRE will use to log in to and access the database. You must therefore have granted sufficient permissions to this user ID to allow it to read and update all replicated tables in the database. This is the PDUSER account. See the [“Setting Up Database Logins and Permissions”](#) section.

An example of using DIMPRULE with a Progress database follows:

```
DIMPRULE -FIRSTSITE -DA=Progress -DB=outeredge1:5432:db1 RULE.ru1
-LOGWIDE -VERBOSE
```

- 6 ♦ Back up the database so that you can restore it to InnerEdge Server, OuterEdge Server, and Workstation database sites. Make sure you have stopped the server on the database. A sample backup command follows:

```
PROBKUP dbname dbname.bak
```

where *dbname* is the name of your first deployment database.

## 2.4.2 Activating the InnerEdge Server Progress Database

After you have prepared the InnerEdge Server, you must activate it. You must have a license file before you run the DINST utility on your database. Ensure that you have exclusive access to the database during this process or DINST will fail.

### Activating the Database at the First InnerEdge Server Site

Follow these steps to activate the database at the first InnerEdge Server site:

- 1 ♦ Start your database server.
- 2 ♦ Open a command-line prompt window. Progress recommends using PROENV.

- 3 ♦ Enter the following command:

```
DINST -DA=db_type -DB=db_locator -ODBCN=username -ODBCP=password -FI  
-PNL=license.pn1
```

where

*db\_type* is the type of database to be activated. For Progress databases, *db\_type* will always be Progress.

*db\_locator* is the location of the deployment database.

*username* is the user name you used in the PRD. PeerDirect strongly recommends that you use *PDUSER* as the user name. See the [“Setting Up Database Logins and Permissions”](#) section.

*password* is the password that corresponds to the user name above. The parameter is optional and defaults to (""), an empty string. Be sure to enclose the password in double-quotes if it contains spaces.

-FIRSTSITE designates this site to be the first site in the replication network and is used at only one site per replication network. When you activate subsequent InnerEdge sites after your first site, do not use the -FIRSTSITE parameter.

-PNL is the replication network license file and specifies the network license to use with your replication network.

*license.pn1* is the full path and filename of the license. If you do not supply this information, a dialog box appears asking you to supply the license file location.

**NOTE:** The -ODBCN and -ODBCP parameters represent the user ID and password that PDRE will use to log in to and access the Progress database. You must therefore have granted sufficient permissions to this user ID to allow it to read and update all replicated tables in the database. See the [“Setting Up Database Logins and Permissions”](#) section.

For a complete list of DINST arguments and other PDRE utility commands, use the "-?" argument; for example, "DINST -?".

- 4 ♦ A dialog box appears informing you that the activation of the database was successful.
- 5 ♦ Click **OK**.

## Activating the Database at a Second InnerEdge Server Site

If you have prepared a second InnerEdge Server site, you must activate it. You must have a license file before you run the DINST utility on your database. Ensure that you have exclusive access to the database during this process or DINST will fail. Follow these steps:

- 1 ♦ Restore the deployment database to the location you desire an InnerEdge server database site. An example of the restore command follows:

```
PROREST dbname dbname.bak
```

Where *dbname* is the name of the deployment database you backed up in the [“Activating Your Replication Sites”](#) section of this chapter.

- 2 ♦ Start your database server.
- 3 ♦ Open a command-line prompt window. Progress recommends using PROENV.
- 4 ♦ Enter the following command:

```
DINST -DA=db_type -DB=db_locator -ODBCN=username -ODBCP=password  
-SITETYPE=high -PA=partner_site
```

where

*db\_type* is the type of database to be activated. For Progress databases, *db\_type* will always be Progress.

*db\_locator* is the location of the deployment database.

*username* is the user name you used in the PRD. You should use *PDUSER* as the user name. See the [“Setting Up Database Logins and Permissions”](#) section.

*password* is the password that corresponds to the user name above. The parameter is optional and defaults to ("") an empty string. Be sure to enclose the password in double-quotes if it contains spaces.

*high* specifies that the site is available for replication at all times.

*partner\_site* specifies the first PDRE site you installed, which should also be the first InnerEdge Server site installed. It is the site of the PDRE, which might not be the same site as the first database site.

**NOTE:** The `-ODBCN` and `-ODBCP` parameters represent the user ID and password that PDRE will use to log in to and access the Progress database. You must therefore have granted sufficient permissions to this user ID to allow it to read and update all replicated tables in the database. See the [“Setting Up Database Logins and Permissions”](#) section.

For a complete list of DINST arguments and other PDRE utility commands, use the `"-?"` argument; for example, `"DINST -?"`.

- 5 ♦ A dialog box appears informing you that activation of the database was successful.
- 6 ♦ Click **OK**.

### 2.4.3 Activating OuterEdge Server Sites

Once you have prepared your OuterEdge Server site, you must activate it. You must have a license file before you run the DINST utility on your database. Ensure that you have exclusive access to the database during this process or DINST will fail. Follow these steps to activate you OuterEdge Server site:

- 1 ♦ Restore the deployment database to the location you desire for an OuterEdge server database site. An example of the restore command follows:

```
PROREST dbname dbname.bak
```

Where *dbname* is the name of the database you backed up in the [“Activating Your Replication Sites”](#) section of this chapter.

- 2 ♦ Start your database server.
- 3 ♦ Open a command-line prompt window. Progress recommends using PROENV for this.

- 4 ♦ Enter the following command:

```
DINST -DA=Progress -DB=db_locator -ODBCN=username -ODBCP=password  
-SITETYPE=site_type -PA=partner_site
```

where

*db\_type* is the type of database to be activated. For Progress databases, *db\_type* will always be Progress.

*db\_locator* is the location of the deployment database.

*username* is the user name you used in the PRD. You should use *PDUSER* as the user name. See the [“Setting Up Database Logins and Permissions”](#) section.

*password* is the password that corresponds to the user name above. The parameter is optional and defaults to ("") an empty string. Be sure to enclose the password in double-quotes if it contains spaces.

*site\_type* specifies one of the following values. For an OuterEdge Server, select either low or high . Select complete only if you are configuring your OuterEdge Server to be a failover server:

- a) low — (default) Site is available only for replication from time to time.
- b) high — Site is available for replication at all times.
- c) complete — Site is available for replication at all times and the database stores all of the data in the replication network.

*partner\_site* specifies a PDRE InnerEdge Server site installed with complete data. It is the site of the PDRE, which might not be the same site as the first database site.

**NOTE:** The -ODBCN and -ODBCP parameters represent the user ID and password that PDRE will use to log in to and access the database. You must therefore have granted sufficient permissions to this user ID to allow it to read and update all replicated tables in the database. See the [“Setting Up Database Logins and Permissions”](#) section.

For a complete list of DINST arguments and other PDRE utility commands, use the "-?" argument; for example, "DINST -?".

- 5 ♦ A dialog box appears informing you that the activation of the database was successful.
- 6 ♦ Click **OK**. The first replication between the InnerEdge Server and OuterEdge site occurs.

**NOTE:** If you are activating the OuterEdge Server database at another workstation, ensure PDRE is running on an InnerEdge Server in order for the first replication to complete. The first replication at a new site normally takes longer than subsequent replications, because PDRE guarantees that all data in common at both sites are exactly synchronized. Subsequent replications normally send only what has changed.

## 2.5 Initiating a Replication Session

PDRE initiates replication bidirectionally from either the InnerEdge Server or the OuterEdge Server. If your PeerDirect-prepared databases have activated successfully, follow these steps to manually initiate replication sessions.

- 1 ♦ Open the PDRE user interface. On Windows, you start the PDRE user interface with **Start → Programs → PeerDirect → PDRE → Replication Engine**. On Linux, you start the PDRE with the command **pdreui**.
- 2 ♦ Choose **Action → Initiate Replication**.
- 3 ♦ The Initiate Replication dialog box appears. Select the database you want to initiate, typically the second database site.
- 4 ♦ Click **Replicate Now**.

If one of the replication sites has a site type of **LOW**, PDRE might tell you that it is unable to find a suitable replication partner site. This happens if you select a **HIGH** site and click **Replicate Now**, because **LOW** sites are never chosen by default as partner sites. If you want to initiate replication from a **HIGH** site to a **LOW** site, you must use the **Advanced** button and explicitly select the two sites you want to replicate between. The easiest way to avoid this message is to always select the **LOW** site and initiate the replication from the **LOW** site to the **HIGH** site.

## 2.6 Administering a Replication Network

To administer a replication network, the Windows PeerDirect Administrator is required. This application is used to schedule or subscribe to work set slices.

If your replication network is running on Linux, you must install the Windows PeerDirect InnerEdge Server or OuterEdge Server and remotely connect to Linux. When you install, follow these steps:

- 1 ♦ Choose the **Custom Installation** option and the **Client Access Edition**.
- 2 ♦ When you have completed the installation, then you can activate a client database site by using the DINST utility.

This allows the PeerDirect Administrator to connect to a PeerDirect Replication Engine on Linux and remotely administer the specified database in the network.

---

## Setup in a Heterogeneous Environment

This chapter describes how to set up PeerDirect replication in a heterogeneous environment (where Progress is just one of the several database types) using your PeerDirect SDK software. This chapter takes you through each phase of the PeerDirect Replication Engine setup and includes these sections:

- [Overview](#)
- [Database Considerations](#)
- [Setting Up Databases for Replication](#)
- [Verifying Heterogeneous Database Compatibility](#)
- [About Preparing Databases for Activation](#)
- [Initiating a Replication Session](#)
- [Administering a Replication Network](#)

## 3.1 Overview

All successful replication networks begin with a deployment strategy and plan. Begin by considering the different layout options that PeerDirect offers and choose the one that offers you the benefits and features you need.

The prerequisite setup requires that you have installed your databases on the appropriate machines. For Progress, you also need to ensure you have the necessary Progress SQL-92 access and supporting ODBC drivers. In order for heterogeneous databases to replicate with each other, they must have matching schema.

The initial setup of a replication network requires that you install the PRD available in the Windows PeerDirect SDK product. You need the PRD to define the replication rules for the database. These rules allow the PDRE to know which data and how the data is to be synchronized between two or more sites within a PeerDirect Replication Network. Once you have created your rules, you then verify and export them.

In a heterogeneous environment, both the schema of your databases and the base SQL data types must match. For more information on compatible data types for each database type, see the *PeerDirect Distributed Enterprise Developer's Guide*. Each database type has a hashed fingerprint. This fingerprint represents the structure of the database. If it matches the other database's fingerprint, your databases are compatible and you can enable them for replication.

The fingerprint can be found in two files produced by the PRD. The filenames correspond to the ODBC data source name prefixed by x, for example xpd3get4m. The extensions of the files are .ref, a summary of the schema and PDRE system information, and .out, a detailed description of the database including internal data types and SQL data types. These files are generated when valid rules are saved using the PRD.

Comparing the compatibility fingerprint of your two replication databases will quickly tell you whether they are truly compatible for the purposes of replication. If there are discrepancies, you must investigate and fix them before proceeding or your databases will not replicate.

Once you have verified that your databases are compatible, you can apply your .rul file against an InnerEdge Server, OuterEdge Server, or Workstation database using the DIMPRULE utility. This step enables the database for replication and prepares it for activation into a PeerDirect Replication Network. To activate the database, you use the DINST utility. The DIMPRULE and DINST utilities are present in both the Windows and Linux versions of PDRE.

### 3.1.1 Setting Up Replication on an InnerEdge Server

You must define replication rules **before** you can complete the setup of the first InnerEdge Server site. To set up replication on the first Inner Edge Server, follow these steps:

- 1 ♦ Install the PeerDirect Distributed Enterprise SDK software and use the PDR to create replication rules. See the *PeerDirect Distributed Enterprise SDK Installation Guide* and the *PeerDirect Distributed Enterprise Developer's Guide* for information about how to create rules.
- 2 ♦ Install an InnerEdge Server running PDRE on a primary server machine that hosts the database with which you plan to replicate OuterEdge Servers. See the *PeerDirect InnerEdge Server Installation Guide* for installation instructions.
- 3 ♦ Set up database logins and permissions so that PDRE can access the required data for replication. See the [“Setting Up Database Logins and Permissions”](#) section.
- 4 ♦ Connect to the PRD using one of your databases. See the [“Selecting Databases to Replicate”](#) section in this chapter for more information.
- 5 ♦ Choose the tables to replicate. See the [“Selecting Tables for Replication”](#) section in this chapter for more information.
- 6 ♦ Verify and save your replication rules. See the [“Verifying and Saving Replication Rules”](#) section in this chapter for more information.
- 7 ♦ Prepare to activate your database with the data definitions files and export your rules to a .rul file. See the [“Preparing to Activate with PRD”](#) section in this chapter for more information.
- 8 ♦ Verify the compatibility of your heterogeneous databases by comparing each databases .ref and .out file. See the [“Verifying Heterogeneous Database Compatibility”](#) section in this chapter for more information.
- 9 ♦ Using the DIMPRULE utility, prepare to activate the database. See the [“Preparing to Activate Your Replication Sites”](#) section.
- 10 ♦ Using the DINST utility, activate the InnerEdge Server site. See the [“Activating the InnerEdge Server Database”](#) section.

**NOTE:** You must activate the site on the InnerEdge Server first and then activate all remaining sites, be they InnerEdge or OuterEdge.

### 3.1.2 Setting Up Replication on an OuterEdge Server or Workstation

You should already have installed one InnerEdge Server on a primary server that hosts the database you plan to use to replicate data. If you choose Method A (see [Table 3–1](#)), you should also have completed the preparation of your first InnerEdge Server database and activated the site.

If you chose Method B (see [Table 3–2](#)), after you install both the InnerEdge and OuterEdge Servers, the next step is to complete the setup of the InnerEdge Server and then set up the OuterEdge Server.

To set up replication on an OuterEdge Server or Workstation, follow these steps:

- 1 ♦ Install an OuterEdge Server or Workstation that runs PDRE. See the *PeerDirect OuterEdge Server Installation Guide* for installation instructions.
- 2 ♦ Set up database logins and permissions so that PDRE can access the required data for replication. See the “[Setting Up Database Logins and Permissions](#)” section.
- 3 ♦ Prepare the OuterEdge Server or Workstation database. See the “[Preparing and Activating the OuterEdge Server Sites](#)” section.
- 4 ♦ Activate the OuterEdge Server or Workstation site. As part of this process, the first replication session will be initiated automatically. See the “[Preparing and Activating the OuterEdge Server Sites](#)” section.

After the first replication, you can use the PDA console to set up the PeerDirect Replication Engine to automatically replicate data on a schedule.

### 3.1.3 Setup Task Sequence

Setting up a PeerDirect Replication Network requires that you perform a sequence of tasks. There are two ways to perform this sequence. You can choose to set up your replication sites one by one, beginning with your first InnerEdge Server (see [Table 3–1](#)), or you can perform all your installation tasks first and then prepare and activate each site separately (see [Table 3–2](#)). You should choose the method that you think works best for your particular deployment purposes. Tables [Table 3–1](#) and [Table 3–2](#) describe the two methods. For detailed instructions, see the referenced resources.

Table 3–1 gives a sequence of steps to follow to set up your PeerDirect Replication Network.

**Table 3–1: Sequence of Replication Network Setup Tasks (Method A) (1 of 3)**

Step	Task	Description
1	Plan your PeerDirect Replication Network.	<p>Plan and design your distributed application system and replication network as an up-front activity. Review your database configuration and schema design to make sure it meets the PeerDirect requirements for replication of data.</p> <p>You will find some typical replication network layouts described in the “<a href="#">Choosing a Layout</a>” section in <a href="#">Chapter 1, “Introduction,”</a> of this guide.</p>
2	Install the Distributed Enterprise SDK (Windows version).	<p>When you install the SDK, you must also install the PRD to create the database rules for your PeerDirect Replication Network.</p> <p>The PRD is only available as part of the PeerDirect Distributed Enterprise SDK software product.</p> <p>See the <i>PeerDirect Distributed Enterprise SDK Installation Guide</i> for detailed instructions on installing the SDK.</p>
3	Define replication rules.	<p>Use the PRD to define the database replication rules for your PeerDirect Replication Network. See the <i>PeerDirect Distributed Enterprise Developer’s Guide</i>.</p>
4	Export replication rules.	<p>After you have created replication rules and have verified the rules, select <b>Tools</b> → <b>Prepare to Activate</b> and then <b>File</b> → <b>Export Rules</b>. This exports the rules to a .rul file that you use to set up your replication databases at each site.</p> <p>See the “<a href="#">Exporting Rules</a>” section in this chapter for more information.</p>
5	Install the InnerEdge Server as the first site.	<p>Install the InnerEdge Server on a central data server that is highly available to the other sites in your PeerDirect Replication Network. See the <i>PeerDirect InnerEdge Server Installation Guide</i>.</p>

**Table 3–1: Sequence of Replication Network Setup Tasks (Method A) (2 of 3)**

Step	Task	Description
6	Prepare the InnerEdge Server database.	Prepare the first site database using the DIMPRULE utility. Be sure to use the -FIRSTSITE parameter. See the <a href="#">“Preparing to Activate Your Replication Sites”</a> section in this chapter for more information.
7	Activate the InnerEdge Server site.	Activate the InnerEdge Server site using the DINST utility. Be sure to use the -FIRSTSITE parameter. See the <a href="#">“Activating the InnerEdge Server Database”</a> section in this chapter for more information.
8	Install an InnerEdge Server, OuterEdge Server, or Workstation as a second site.	After you install and activate your first site, install your second replication site.  Later, you will repeat this step for the other sites in your replication network.  <i>See the <a href="#">PeerDirect InnerEdge Server Installation Guide</a> or the <a href="#">OuterEdge Server and Workstation Installation Guide</a>.</i>
9	Verify Rules for an InnerEdge Server, OuterEdge Server, or Workstation database as a second site.	From your first InnerEdge Server site, you will save the rules to a second replication site and check the .ref and .out files for compatibility.  Later, you will repeat this step for all remaining sites you have installed as part of your replication network.  <a href="#">See the “Verifying Heterogeneous Database Compatibility”</a> section in this chapter for more information.
10	Prepare the InnerEdge Server or OuterEdge Server or Workstation database as a second site.	Prepare the second site’s database by using the DIMPRULE utility.  Later, you will repeat this step for all remaining sites you have installed as part of your replication network.  <a href="#">See the “Preparing and Activating the OuterEdge Server Sites”</a> section in this chapter for more information.

**Table 3–1: Sequence of Replication Network Setup Tasks (Method A)** (3 of 3)

Step	Task	Description
11	Activate the InnerEdge Server, OuterEdge Server, or Workstation site as a second site.	<p>Activate the second site using the DINST utility. As part of this process, the initial replication session is started with the first site.</p> <p>Later, you will repeat this step for all remaining sites.</p> <p>See the <a href="#">“Preparing and Activating the OuterEdge Server Sites”</a> section in this chapter for more information.</p>
12	Repeat Steps 8–11	Install and prepare to activate the remaining sites in your PeerDirect Replication Network.

[Table 3–2](#) gives an alternate sequence of steps to follow to set up your PeerDirect Replication Network.

**Table 3–2: Sequence of Replication Network Setup Tasks (Method B)** (1 of 3)

Step	Task	Description
1	Plan your PeerDirect Replication Network.	<p>Plan and design your distributed application system and replication network as an up-front activity. Review your database configuration and schema design to make sure it meets the PeerDirect requirements for replication of data.</p> <p>You will find typical replication network layouts described in the <a href="#">“Choosing a Layout”</a> section in <a href="#">Chapter 1, “Introduction,”</a> of this guide.</p>
2	Install the Distributed Enterprise SDK (Windows version).	<p>You must install the PRD to create the database rules for your PeerDirect Replication Network.</p> <p>The PRD is only available as part of the PeerDirect Distributed Enterprise SDK software product.</p> <p>See the <i>PeerDirect Distributed Enterprise SDK Installation Guide</i> for detailed instructions on installing the SDK.</p>

**Table 3–2: Sequence of Replication Network Setup Tasks (Method B) (2 of 3)**

Step	Task	Description
3	Define replication rules.	<p>Use the PRD to define the database replication rules for your PeerDirect Replication Network. This is where the work you did in Step 1 pays off.</p> <p>See the <i>PeerDirect Distributed Enterprise Developer's Guide</i>.</p>
4	Export replication rules to a file	<p>After you have created replication rules and are satisfied, select <b>Tools</b>→<b>Prepare to Activate</b> and then <b>File</b> → <b>Export Rules</b>. This exports the rules to a rules file that you use to set up your replication databases at each site.</p> <p>See the “<a href="#">Exporting Rules</a>” section in this chapter for more information.</p>
5	Install the InnerEdge Server as the first site.	<p>Install the InnerEdge Server on a central data server that is highly available to the other sites in your PeerDirect Replication Network.</p> <p>See the <i>PeerDirect InnerEdge Server Installation Guide</i> for more information.</p>
6	Install an InnerEdge Server, OuterEdge Server, or Workstation on the second and remaining sites in the network.	<p>After you install and activate your first site, install your second replication site.</p> <p>Repeat this step for the other sites in your replication network.</p> <p>See the <i>PeerDirect InnerEdge Server Installation Guide</i> or the <i>PeerDirect OuterEdge Server and Workstation Installation Guide</i> for more information.</p>
7	Verify rules for an InnerEdge Server, OuterEdge Server, or Workstation database as a second site.	<p>From your first InnerEdge Server site, you will save a different rules file for each database type. Be sure to check the .ref and .out files for compatibility.</p> <p>Later, you will repeat this step for each database type you have installed as part of your replication network.</p> <p>See the “<a href="#">Verifying Heterogeneous Database Compatibility</a>” section in this chapter for more information.</p>

**Table 3–2: Sequence of Replication Network Setup Tasks (Method B)** (3 of 3)

Step	Task	Description
8	Prepare the InnerEdge Server database as the first site.	Prepare the first site database using the DIMPRULE utility. Be sure to use the -FIRSTSITE parameter.  See the “ <a href="#">Preparing to Activate Your Replication Sites</a> ” section in this chapter for more information.
9	Activate the InnerEdge Server site as the first site.	Activate the InnerEdge Server site using the DINST utility. Be sure to use the -FIRSTSITE parameter.  See the “ <a href="#">Activating the InnerEdge Server Database</a> ” section in this chapter for more information.
10	Prepare the InnerEdge Server, OuterEdge Server, or Workstation database on the second site and remaining sites in the network.	Prepare the second site’s database by using the DIMPRULE utility.  Repeat this step for all remaining sites you have installed as part of your replication network.  See the “ <a href="#">Preparing to Activate Your Replication Sites</a> ” and the “ <a href="#">Preparing and Activating the OuterEdge Server Sites</a> ” sections in this chapter for more information.
11	Activate the InnerEdge Server, OuterEdge Server, or Workstation site on the second site and remaining sites in the network.	Activate the second site using the DINST utility. As part of this process, the initial replication session is started with the first site.  Repeat this step for all remaining sites.  See the “ <a href="#">Preparing to Activate Your Replication Sites</a> ” and the “ <a href="#">Preparing and Activating the OuterEdge Server Sites</a> ” sections in this chapter for further instruction.

Each InnerEdge Server, OuterEdge Server, or Workstation runs a PDRE component that replicates data among the databases in your PeerDirect Replication Network. You can select to install PDRE on your server or workstation when you run the PeerDirect install program (on **Windows**, select the **Runtime** edition).

For instructions on how to install an OuterEdge Server, see the *PeerDirect OuterEdge Server Installation Guide*. For instructions on how to install an InnerEdge Server, see the *PeerDirect InnerEdge Server Installation Guide*.

### 3.1.4 Recommendations

You will find instructions in this chapter for setting up replication on both the PeerDirect InnerEdge and OuterEdge Servers. As you work, it is important that you follow all the instructions in the sequence in which they are presented. Do not skip any steps.

## 3.2 Database Considerations

When replicating databases with the PeerDirect Replication Engine, you should pay attention to the recommendations described in the following sections.

### Heterogeneous Replication Networks

When you configure a database to work with PDRE, you have the ability to replicate with databases that are different (heterogeneous), such as Progress and Oracle databases. The only requirement is that there be at least one instance of the Progress database in your configuration. For a list of supported databases, see the “[Heterogeneous Database Products](#)” section in [Chapter 1, “Introduction.”](#)

When you replicate between different vendor-specific databases, you must ensure that the data types being used are compatible. Replication cannot occur unless the documented “basic data type” match for a given field across the two different database platforms. The length and precision of each column must also be identical across the two database environments.

For detailed information about replicating with databases of different types, see the *PeerDirect Distributed Enterprise Developer’s Guide*.

### Foreign Keys

Foreign keys are important if you want to impact the order in which tables are replicated. If you plan to use work sets or transaction sets, you **must** use foreign keys.

It is important that PDRE know the foreign keys for each table involved in replication. Database table relationships are defined from the bottom up, that is from the first child to the parent to which it is related by a foreign key. PDRE uses this information to determine the order in which tables are replicated, because parent tables are always replicated before child tables. Foreign key relationships are also used to define work sets and transaction sets, and to ensure that deletions are cascaded to dependent tables.

For instructions about how to define foreign keys, see the “[Specifying Foreign Keys](#)” section.

## Cascading Table Deletions

If your database includes local (nonreplicated) tables that are related to replicated tables, cascading deletions will not extend to these local tables. This means that you must either use triggers to ensure that deletions are cascaded to local tables or enable cascading deletions in your database software or in your business application.

Note that if you implement referential integrity constraints between these local and replicated tables without ensuring that deletions are properly cascaded to the local tables, PDRE deletions will cause these constraints to “break.”

## Consistent Naming

Ensure that all tables you plan to replicate have the same names across replicated databases. This is not a problem if you enable a first database and then copy it to other locations for replication.

Columns in replicated tables must also have identical names, and when replicating between two different vendor databases, they must have compatible data types.

There are no requirements for tables or columns that you do not plan to replicate.

## Unsupported Data Types

Tables you plan to replicate should not include unsupported data types or associated settings or properties. By default, PDRE replicates all columns in a table, but you can modify this behavior using a feature called fragments. For more information on fragments, see the “[Defining Fragments](#)” section in this chapter.

For a list of supported data types and database-specific types, please refer to the *PeerDirect Distributed Enterprise Developer’s Guide*.

## Unique Primary Keys

All tables you plan to replicate must have a unique primary key. With the exception of Microsoft SQL Server and Progress RDBMS, it is not necessary to have primary key constraints physically defined in the database.

To enhance performance, PeerDirect recommends that you set a unique index on each table’s primary key. Primary keys can be *simple*—comprised of a single column—or *composite*—comprised of two or more columns. However, no component of a primary key can contain a NULL value.

Primary keys for the Progress database must be set up in the Progress database. The PRD will import them. They cannot be identified in the PRD. However, other non-Progress database types can use the PRD to identify primary keys.

## Database Versions

PDRE supports specific versions of most database platforms. Be sure to consult the *PeerDirect InnerEdge Server Installation Guide*, the *PeerDirect OuterEdge Server Installation Guide*, and the PeerDirect Release Notes for the exact versions of your database software that can be used with PDRE. If you have an earlier or a later version of the database software than the one specified, you might encounter problems if you attempt to use it with PDRE.

## Clock Synchronization

The date and time must be set correctly on all machines in a replication network, including all database servers, PDRE servers, application servers (machines running database clients), and remotely connected machines. You should install and use one of the following programs to synchronize the clocks on your machines:

- Tardis (<http://www.kaska.demon.co.uk>)
- Dimension 4 (<http://www.thinkman.com/dimension4/>)

## 3.3 Setting Up Databases for Replication

To set up a PeerDirect Replication Network, you must perform a series of tasks. The sequence of these tasks is outlined in the “[Overview](#)” section of this chapter. The sections below describe each task in detail.

### 3.3.1 Setting Up Database Logins and Permissions

Both PRD and PDRE require access permissions that allow them to manipulate tables in the database. [Table 3–3](#) shows the required permissions of the PRD and PDRE.

**Table 3–3: PRD and PDRE Access Permissions**

Each installed...	Requires...
PRD	Access to all application tables to be replicated. Permissions to create and update new tables in the database.
PDRE	Exclusive use of a specific login ID if you plan to implement triggers to maintain control fields.

If your database requires users to log in, you will need to provide login IDs that have appropriate permissions to access and update your database.

## Creating a Login ID with Adequate Permissions on Your Progress Database

Before you begin to use PRD or PDRE with a Progress database, you must enable Progress for PeerDirect replication and create a PDUSER account. To create the PDUSER account, you must have installed Progress SQL-92. If PeerDirect is located on a machine other than the Progress database, you must have installed Progress SQL-92 Client Access on the PeerDirect machine. For information regarding required software for using PeerDirect with Progress, see the “Progress Products” section in Chapter 1, “Introduction.” For complete instructions on installing Progress products, see the *Progress Installation and Configuration Guide for Windows* and the corresponding guide for UNIX provided with your Progress documentation.

To enable and create a PDUSER account so that you can begin using PeerDirect with Progress, follow these steps:

- 1 ♦ Open a command-line prompt window.
- 2 ♦ In a command-line shell or window, enable the Progress change capture mechanism for replication on your database by entering the following command:

```
PROUTIL dbname -C enablepdr
```

where *dbname* is the name of your Progress database that you want to replicate.

**NOTE:** You must enable the change capture mechanism for replication with the `proutil` command before you create your PDUSER account.

- 3 ♦ Start a server on the Progress database that you want to replicate:

```
PROSERVE dbname -S 3333
```

where *dbname* is the name of your Progress database.

- 4 ♦ Create the PDUSER account.

The account does not have to be named PDUSER. Any account name can be used as long as it is granted the appropriate permissions.

Only create a PDUSER account if there **are already active accounts** in the database. To determine if there are active accounts, connect to the database. If you are prompted for a user name and password, this indicates there are already active accounts in the database.

Create the PDUSER account by running the SQL Explorer and connecting to the database as the user that owns the database. You might need to create a SYSPROGRESS account through the 4GL dictionary to be able to connect to the database in some cases. Once connected, enter the following commands:

```
create user 'PDUSER', 'PASSWORD';  
grant resource, dba to 'PDUSER';  
commit work;
```

If there **are no active accounts** for the database, then enter the following commands:

```
grant resource, dba to 'PDUSER';  
commit work;
```

### Creating a Login ID with Adequate Permissions in a Non-Progress Database

Before you begin to use PRD or PDRE, you should do the following to create a login ID with adequate permissions:

- 1 ♦ Create a database user called “PDUSER.”
- 2 ♦ Provide PDUSER with insert, update, delete, and select access to all replicated tables, all PeerDirect system tables, and all PeerDirect control tables in the database.
- 3 ♦ Provide PDUSER with all permissions equivalent to those of the database owner.
- 4 ♦ Set up PRD and PDRE to use the PDUSER login ID.

Alternatively, you can have PDRE use the login ID of the user who owns the database tables you want to replicate. Whatever you use, make sure it is a user ID that is only used by PDRE.

**NOTE:** Do not use an alias or the Oracle “system” user. Do not alias the PDUSER login to any other user.

### 3.3.2 Selecting Databases to Replicate

Use the PRD to select and configure the databases you want to replicate as part of a replication network.

**NOTE:** For Microsoft SQL Server and Progress RDBMS, be sure to define primary keys in the database before you run PRD. Otherwise you will encounter replication errors.

When you start PRD for the first time, PRD asks you detailed questions about the database, connection, and replication information.

#### Selecting a Database to Replicate

Follow these steps to select a database to replicate:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→ **Programs**→ **PeerDirect**→ **PDRE**→ **Designer**.
- 2 ♦ The Database Information screen appears. Complete the following fields:
  - a) Database Type — Select the type of database you want to replicate. Database
  - b) Locator — Enter the location of the database you want to replicate. Depending on the database type you select, you will see one of these formats as shown in [Table 3–4](#).

**Table 3–4: Database Type Formats**

Database Type	Format
Progress RDBMS	<host>:<port>:<db_name>
IBM DB2 UDB	<db_name>
Microsoft SQL Server	<server_name>:<db_name>[:<extras>]
Oracle 8	<server_name>
Oracle 8.1.6 and above	<host>:<SID>:<port>

- c) Database Connection Information — (optional) If you need a user/login ID and password to access your database through ODBC, enter the following information:

**NOTE:** The user name PDUSER must have already been granted the necessary permissions to the database.

- UserName — Enter the user name or login ID required to access the database.
- Password — Enter the password required to access the database.
- Optional Parameters — Enter any additional connection parameters. For example, if your database uses a system database to validate users, you must enter its name here. The syntax is `systemdb=myfile.mdw`.

3 ♦ Click **Next**.

4 ♦ The Database Information screen appears. Complete the following:

- a) Project — Enter the name of your project.
- b) Network — Enter the name of the replication network. The network will be comprised of the databases you identify for replication.

**NOTE:** Only databases that are part of the same project and network can replicate together. Databases in two different networks will not interact.

- c) Release — Enter the name of this release. Each project/network can have multiple releases as new replication features are added or database schema changes are implemented.

5 ♦ Click **Next**.

### 3.3.3 Selecting Tables for Replication

Use the main PRD window to specify which tables to replicate (tables are initially marked as “Not Replicated”) and to configure the advanced features of PDRE.

Follow these steps to select a table for replication:

- 1 ♦ In the main PRD window, select the tables in your database that you want to replicate.
- 2 ♦ Choose **Actions** → **Replicate Table**.
- 3 ♦ Click the **Replicated** tab at the bottom of the left panel. The tables marked for replication are shown in the replicated list.

### 3.3.4 Specifying Foreign Keys

With Progress databases, you must manually define foreign keys in the PRD. The PRD will not import foreign key references defined in a Progress database.

With databases other than Progress, if you have foreign key constraints defined in your database, PRD will automatically detect them. These foreign keys display when you select the **Foreign Keys** tab. With these databases, you can add more foreign key relationships.

It is important that PDRE know what the foreign keys are for each table it replicates. Foreign keys determine the order in which PDRE replicates tables, replicating parent tables before child tables. Foreign keys are also required to define work sets and transaction sets. See the “[Foreign Keys](#)” section in this chapter.

Follow these steps to specify a foreign key relationship:

- 1 ♦ In the main PRD window, click the **Foreign Key** tab.
- 2 ♦ In the Create new relationship from <TableName> to parent table field, select a parent table for this child table from the drop-down menu. The primary key fields for the parent table and candidate foreign key fields appear.
- 3 ♦ Select the appropriate foreign key field(s) that correspond to the primary key field(s).
- 4 ♦ Click **Add Relationship**. The relationship is added to the list of Existing Relationships.

Repeat these steps for each foreign key in your database.

Follow these steps to modify a foreign key relationship:

- 1 ♦ Select a relationship from the Existing Relationships list.
- 2 ♦ Select the primary key and foreign key fields from the drop-down menu.
- 3 ♦ Click **Apply Changes**.

To delete a relationship, select the relationship from the Existing Relationships list, and click **Delete Relationship**.

### 3.3.5 Defining Fragments

You can define fragments that allow changes to flow and merge when they do not occur in the same group of data. *Fragments* are the smallest unit of replication. By grouping fields that should change together into fragments, the application developer can prevent false collisions and detect real ones. You can create or add to a fragment by following these steps in the PRD:

- 1 ♦ Select a table in the table selection window.
- 2 ♦ Select the **Structure** tab.
- 3 ♦ Click the left-mouse button in a field that you want in a new fragment.
- 4 ♦ From the Actions menu, select **Create Fragment**.
- 5 ♦ Enter the name of the fragment.
- 6 ♦ Choose **OK**. The table name appears in the Fragment column of the Structure tab.

### 3.3.6 Defining Worksets

PeerDirect Distributed Enterprise allows you to define the business rules for subsetting around base tables. This subsetting is known as work sets. Follow these steps to define work sets:

- 1 ♦ Select a base table in the Table Selection window, then right-click and select **Create Work Set**.
- 2 ♦ Select the **Work Set** tab.
- 3 ♦ To add a table to a work set, select the parent table in the Work Set window.
- 4 ♦ Select the table you want to add to the work set from the Related Tables list.
- 5 ♦ Select **Add to Work Set**.

### 3.3.7 Defining Transaction Sets and Groups

PeerDirect Distributed Enterprise allows you to define the business rules for data integrity of transactions. A *transaction* is a set of insert, update, or delete operations bound together and treated as a single unit of work. Defining rules for transactions is known as transaction set and transaction group definition. A *transaction set* is a group of records across related tables. A *transaction group* is one or more tables in their entirety. Follow these steps to define a transaction set:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→ **Programs**→ **PeerDirect**→ **PDRE**→ **Designer**.
- 2 ♦ Make sure the **Transaction** tab is selected.
- 3 ♦ Select a base table in the Table Selection window, then right click and select **Create Transaction Set**.

Follow these steps to define a transaction group:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→ **Programs**→ **PeerDirect**→ **PDRE**→ **Designer**.
- 2 ♦ Make sure the **Transaction** tab is selected.
- 3 ♦ At the bottom of the Transaction tab, select the **Group** tab.
- 4 ♦ Select a table from the table drop down list.
- 5 ♦ Select **Create New Group**.
- 6 ♦ The New Transaction Group dialog box appears. Enter a name for the group.
- 7 ♦ Click **OK**.
- 8 ♦ In the Existing Tables Not In a Transaction, you can click on a table and select the **Add to Group** button to add the table to the Transaction Group.

### 3.3.8 Verifying and Saving Replication Rules

The next step is to verify and save your database's replication rules. During verification the database performs checks for cyclic dependencies, such as circular foreign key references, and whether each replicated table has a primary key specified. During verification, the error messages are generated one at a time.

Follow these steps to verify and save database replication rules:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→**Programs**→**PeerDirect**→**PDRE**→**Designer**.
- 2 ♦ In the main PRD window, choose **Tools** → **Verify Rules**.
- 3 ♦ If PRD finds errors in your database, they are displayed. Correct them and rerun the verify tool.
- 4 ♦ Repeat Step 2 until all error messages are resolved.
- 5 ♦ Choose **File** → **Save**.

**NOTE:** PRD lets you save your changes even if there are errors, but you cannot continue with the next steps unless you correct all errors. To view the error messages, select **View**→**Verification Messages**.

It is good practice at this point to export the script file of your rules for future development.

## 3.4 Verifying Heterogeneous Database Compatibility

In the PRD that is connected to your development database, follow these steps to ensure that your .rul file is compatible across all database types in your replication network:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→**Programs**→**PeerDirect**→**PDRE**→**Designer**.
- 2 ♦ Choose **File**→**Save As**.

You must be sure that the machine you are working on has access (local or via a network) to the second database.

- 3 ♦ Save your .rul file to the second database. Include the login ID and password if necessary.

Leave the **Display New Database After Saving** check box selected. This activates your second database in the PRD.

- 4 ♦ In the main PRD window, choose **Tools** → **Verify Rules**.

If your two databases are compatible and you received no errors when verifying rules on the first database, you should not receive any errors.

- 5 ♦ Choose **File**→ **Save**.

- 6 ♦ Compare the .ref and .out files from both databases to make sure they are compatible. You must verify that the compatibility fingerprints match and that the table definitions match.

**NOTE:** It is very important that you choose **File**→ **Save As** for all your .rul files for each database type before you activate your databases. You could have to redesign your PeerDirect Replication Network if you want to load a new database type.

### 3.5 About Preparing Databases for Activation

You must prepare each database in your replication network for activation using the DIMPRULE utility. If the first site was prepared for activation by PRD (see the [“Preparing to Activate with PRD”](#) section), you do not need to run the DIMPRULE utility.

The DIMPRULE utility:

- Writes information to PeerDirect system tables.
- Creates a control table for each table you are replicating.
- Preprocesses any preexisting user data—also referred to as *starter data*—found in replicated tables by creating control table records for each user table record. (This step is performed at the first InnerEdge Server site but only if you check the **Mark starter data for replication** check box.)

**NOTE:** It is very important that you mark starter data for replication only at your **first site**, which should be an InnerEdge Server site.

PeerDirect recommends that you run PRD in your SDK development environment to create replication rules, and then import these rules to the first site of your production environment (which should be an InnerEdge Server) using the DIMPRULE utility (see the “[Preparing to Activate Your Replication Sites](#)” section). Be sure to specify the DIMPRULE FIRSTSITE parameter at the first InnerEdge Server site.

If you use PRD to activate your first site, you must be sure to check the **Mark starter data for replication** check box for InnerEdge Servers. All data in the OuterEdge databases will be initialized during a special *starter data reconciliation phase* that occurs during the initial replication with the InnerEdge Server database.

### 3.5.1 Preparing to Activate Your Replication Sites

Begin by preparing your first InnerEdge Server for activation. PeerDirect recommends that you create your .rul file using PRD in your SDK development environment and then run the DIMPRULE utility to prepare your first InnerEdge Server site.

Follow these steps to prepare to activate an InnerEdge Server using the DIMPRULE Utility:

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

```
DIMPRULE -FIRSTSITE -DA=db_type -DB=db_locator -ODBCN=username  
-ODBCP=password rule_file_name.rul [-LOGWIDE] [-V]
```

where

*db\_type* is the type of database to be activated.

*db\_locator* is the location of the database. See [Table 3-4](#).

*username* is the user name you used in the PRD. PDRE will use this name to connect to the SQL database. You should use "PDUSER." See the “[Setting Up Database Logins and Permissions](#)” section.

*password* is the password that corresponds to the user name above. The parameter is optional and defaults to "password". Be sure to enclose the password in double-quotes if it contains spaces.

*rule\_file\_name.rul* is the name of the file you created when you set up the first InnerEdge Server. You can get the name of this file from that site.

-logwide and -V (verbose) are optional parameters that enable detailed logging activity.

-FIRSTSITE designates this site to be the first site in the replication network and is used at only one site per replication network.

**NOTE:** The -ODBCN and -ODBCP parameters represent the user ID and password that PDRE will use to log in to and access the database. You must therefore have granted sufficient permissions to this user ID to allow it to read and update all replicated tables in the database. See the “[Setting Up Database Logins and Permissions](#)” section.

An example of the DIMPRULE utility on a Progress database as a first site follows:

```
DIMPRULE -FIRSTSITE -DA=Progress -DB=outeredge1:5432:db1 -ODBCN=username
-ODBCP=password rule_file.rul -LOGWIDE -V
```

### 3.5.2 Preparing to Activate with PRD

This preparation lets you export a .rul file which PeerDirect can then use to set up and activate other servers. It is important that you create the .rul file for the first site you install and then apply the same .rul file to all other databases of the same type you plan to replicate. You will need a separate .rul file for each database type you plan to activate in the PeerDirect Replication Network. Be sure to write down the name of the .rul file you create for future PDRE installations.

Follow these steps to prepare to activate your InnerEdge Server database:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→ **Programs**→ **PeerDirect**→ **PDRE**→ **Designer**.
- 2 ♦ Choose **Tools** → **Prepare to Activate...**
- 3 ♦ If you know the network address of the InnerEdge Server site in your replication network, enter it. If you are using single engine replication, you can type **localhost**.
- 4 ♦ If you have initial data in the database (starter data), leave the **Mark starter data for replication** check box checked.
- 5 ♦ Click **OK**.
- 6 ♦ A dialog appears telling you that the database has been prepared for activation.
- 7 ♦ Click **OK**.

## Exporting Rules

After you have successfully completed the Prepare to Activate operation, you must export the replication rules to a file. The default file extension is `.rul`. Follow these steps:

- 1 ♦ Start the PeerDirect Replication Designer (PRD) by selecting **Start**→**Programs**→**PeerDirect**→**PDRE**→**Designer**.
- 2 ♦ Choose **File**→**Export Rules** and save the replication rules to a file, such as `xPD060102.rul`. You must use this `.rul` file to set up other replication sites.
- 3 ♦ Exit the PeerDirect Replication Designer.

### 3.5.3 Activating the InnerEdge Server Database

After you have prepared the InnerEdge Server, you must prepare the OuterEdge Server and associate it with the PeerDirect Replication Engine on the InnerEdge Server. This will activate your database, making it available for replication and creating the default PeerDirect Administrator account for the site.

Follow these steps to activate the database at the InnerEdge Server site:

- 1 ♦ In a command-line shell or window, type the following command:

```
DINST -DA=db_type -DB=db_locator -ODBCN=username -ODBCP=password -FI  
-PNL=license.pn1
```

where

*db\_type* is the type of database to be activated and *db\_locator* is the location of the database. See [Table 3–4](#).

*username* is the user name you used in the PRD. You should use "*PDUSER*" as the user name. See the [“Setting Up Database Logins and Permissions”](#) section.

*password* is the password that corresponds to the user name above. The parameter is optional and defaults to ("") the empty string. Be sure to enclose the password in double-quotes if it contains spaces.

`-FI` designates this site to be the first site in the replication network and is used at only one site per replication network.

If activating an additional InnerEdge Server site, you would use the `-PA` parameter instead of `-FI`. The `-PA` parameter would point to a partner InnerEdge Server site.

-PNL is the replication network file and specifies the network license to use with your replication network.

*license.pnl* is the full path and filename of the license. If you do not supply this information, a dialog appears asking you to supply the license file location.

**NOTE:** The -ODBCN and -ODBCP parameters represent the user ID and password that PDRE will use to log in to and access the database. You must therefore have granted sufficient permissions to this user ID to allow it to read and update all replicated tables in the database. See [“Setting Up Database Logins and Permissions.”](#)

For a complete list of DINST arguments and other PDRE utility commands, use the "-?" argument; for example, "DINST -?".

- 2 ♦ A dialog box appears informing you that activation of the database was successful.
- 3 ♦ Click **OK**.
- 4 ♦ If this is **not** a Progress database, initiate the replication triggers by typing the following in a command-line shell:

```
DTRIGGER odbc_datasource -ODBCN=username -ODBCP=password
```

where

*odbc\_datasource* is the name of the data source.

*username* is the user name that PDRE uses to connect to the database.

*password* is the password that corresponds to the user name above.

You do not need to follow this step for a Progress database.

### 3.5.4 Preparing and Activating the OuterEdge Server Sites

This task must be performed from the command-line console of the machine that is running PDRE for the OuterEdge Server. Depending on your configurations, this might not be the same machine on which the OuterEdge database resides.

Once activation of the OuterEdge database is completed, the first replication between the InnerEdge and OuterEdge Server sites occurs.

If you are activating the second database at another workstation, you will encounter errors if PDRE is not already running at the first site. The first replication at a new site normally takes longer than subsequent replications because PDRE guarantees that all data in common at both sites are exactly synchronized. Subsequent replications normally send only what changed.

Follow these steps to prepare to activate the database on the OuterEdge Server:

- 1 ♦ Open a command prompt window.
- 2 ♦ You must run `DIMPRULE` to activate the `.rul` file you created when you set up the first InnerEdge Server. To do this, type the following command:

```
DIMPRULE -DA=db_type -DB=db_locator -ODBCN=username -ODBCP=password  
rule_file_name.rul [-logwide] [-v]
```

where

*db\_type* is the type of database to be activated.

*db\_locator* is the location of the database. See [Table 3–4](#).

*username* is the user name you used in the PRD. PDRE will use this name to connect to the SQL database. You should use "PDUSER." See the [“Setting Up Database Logins and Permissions”](#) section.

*password* is the password that corresponds to the user name above. The parameter is optional and defaults to "password." Be sure to enclose the password in double-quotes if it contains spaces.

*rule\_file\_name.rul* is the name of the file you created when you set up the first InnerEdge server.

`-logwide` and `-v` (verbose) are optional parameters that enable detailed logging activity.

**NOTE:** The `-ODBCN` and `-ODBCP` parameters represent the user ID and password that PDRE will use to log in to and access the database. You must therefore have granted sufficient permissions to this user ID to allow it to read and update all replicated tables in the database. See the [“Setting Up Database Logins and Permissions”](#) section.

An example of the DIMPRULE utility on an Oracle database follows:

```
DIMPRULE -DA=Oracle -DB=outeredge:ora:5432 rulefile.rul
```

To activate an OuterEdge Server site, use the DINST Utility as follows:

- 1 ♦ In a command-line shell or window, type the following DINST command:

```
DINST -DA=db_type -DB=db_locator -ODBCN=username -ODBCP=password  
-SITEYPE=site_type -PA=partner_site
```

*db\_type* is the type of database to be activated.

*db\_locator* is the location of the deployment database.

*username* is the user name you used in the PRD. You should use "PDUSER" as the user name. See the [“Setting Up Database Logins and Permissions”](#) section.

*password* is the password that corresponds to the user name above. The parameter is optional and defaults to ("") the empty string. Be sure to enclose the password in double-quotes if it contains spaces.

*site\_type* specifies one of the following values. For an OuterEdge Server, select either low or high. Select complete only if you are configuring your OuterEdge Server to be a failover server.

- a) Low — (default) Site is available only for replication from time to time.
- b) High — Site is available for replication at all times.
- c) Complete — Site is available for replication at all times and the database stores all of the data in the replication network, that is, all work set slices are auto-subscribed to.

*partner\_site* specifies the first PDRE site you installed, which should also be the first InnerEdge Server site installed.

**NOTE:** The `-ODBCN` and `-ODBCP` parameters represent the user ID and password that PDRE will use to log in to and access the database. You must therefore have granted sufficient permissions to this user ID to allow it to read and update all replicated tables in the database. See the [“Setting Up Database Logins and Permissions”](#) section.

For a complete list of DINST arguments and other PDRE utility commands, use the `"-?"` argument; for example, `"DINST -?"`.

- 2 ♦ A dialog box appears informing you that activation of the database was successful.
- 3 ♦ Click **OK**. The first replication between the InnerEdge Server and OuterEdge site occurs.
- 4 ♦ If you are activating the OuterEdge Server database at another workstation, you will encounter errors if PDRE is not already running on the InnerEdge Server. The first replication at a new site normally takes longer than subsequent replications because PDRE guarantees that all data in common at both sites are exactly synchronized. Subsequent replications normally send only what changed.
- 5 ♦ In a command-line shell or window, if this is **not** a Progress database, initiate replication triggers by typing:

```
DTRIGGER odbc_datasource -ODBCN=username -ODBCP=password
```

where

*odbc\_datasource* is the name of the data source.

*username* is the user name that PDRE uses to connect to the database.

*password* is the password that corresponds to the user name above.

### 3.6 Initiating a Replication Session

PDRE initiates replication bidirectionally from either the InnerEdge Server or the OuterEdge Server. If your PeerDirect-prepared databases have activated successfully, follow these steps to manually initiate replication sessions:

- 1 ♦ Open the PDRE user interface. On Windows, you start the PDRE user interface with **Start**→**Programs**→**PeerDirect**→**PDRE**→**Replication Engine**. On Linux, you start the PDRE with the **pdreui** command.
- 2 ♦ Choose **Action**→**Initiate Replication**.
- 3 ♦ The Initiate Replication dialog appears. Select the database you want to initiate, typically the second database site.
- 4 ♦ Click **Replicate Now**.

If one of the replication sites has a site type of low, PDRE might tell you that it is unable to find a suitable replication partner site. This happens if you select a high site and click **Replicate Now**, because low sites are never chosen by default as partner sites. If you want to initiate replication from a high site to a low site, you must use the **Advanced** button and explicitly select the two sites for replication. The easiest way to avoid this message is to always select the low site and initiate the replication from the low site to the high site.

### 3.7 Administering a Replication Network

To administer a replication network, the Windows PeerDirect Administrator is required. This application is used to schedule or subscribe to work set slices.

If your replication network is running on Linux, you must install the Windows PeerDirect InnerEdge Server or OuterEdge Server and remotely connect to Linux. Follow these steps:

- 1 ♦ Choose the **Custom Installation** option and the **Client Access Edition**.
- 2 ♦ When you have completed the installation, then you can activate a client database site by using the **DINST** utility.

This allows the PeerDirect Administrator to connect to a PeerDirect Replication Engine on Linux and remotely administer the specified database in the network.



# Index

---

## A

- Activate 2–6
- Activating the InnerEdge Server 3–9
- Activating the InnerEdge Server site 2–6, 2–8, 3–6
- Activating the OuterEdge Server 2–8, 3–7, Index–1
- Activating the OuterEdge Server site 3–9
- Activating the OuterEdge Workstation 3–9
- Activation, About preparing for 2–19, 3–21
- Administering the PeerDirect Replication Network 2–28, 3–29
- Agent 1–12

## C

- Cascading table deletions 2–9, 3–11
- Clock synchronization 2–10, 3–12

COMPLETE site type 1–9, 2–26, 3–27

Consistent naming 2–10, 3–11

Control tables 3–21

## D

Data types 2–10, 3–11

Database

- Logins and permissions 2–12, 3–12
- Versions 2–10, 3–12

Databases

- Activating 2–20, 3–22
- Design considerations 2–9, 3–10
- Preparing 2–20, 3–22
- Selecting for replication 2–14, 3–15
- Selecting tables 2–15, 3–16
- Specifying type 2–21, 3–22, 3–26

Defining replication rules 2–5, 2–7, 3–5, 3–8

Deployment 1–8

DIMPRULE Utility 3–21

- Commands 2–21, 3–22

## **E**

- Encrypted communication 1–7, 1–8
- Exporting replication rules 2–5, 2–7, 3–5, 3–8

## **F**

- First site 1–2
- FIRSTSITE parameter 2–19, 3–22
- Foreign keys 2–9, 2–15, 3–10, 3–17
  - Modifying key relationships 2–16, 3–17

## **H**

- Heterogeneous configuration 1–2, 1–11
- Heterogeneous replication networks 3–10
- High Availability 1–2, 1–12
- HIGH site type 2–26, 3–27
- Homogeneous configuration 1–2, 1–9
- Homogeneous replication networks 3–10

## **I**

- InnerEdge Server 1–7, 1–9
  - Activating the database 2–24, 3–24
  - Activating the site 2–6, 2–8, 3–6, 3–9
  - Preparing the database 2–6, 2–8, 3–6, 3–9
  - Preparing to activate 2–19, 3–23
  - Setting up replication 2–2, 3–3
- InnerEdge server 1–11

## **L**

- LAN 1–8
- License 1–9
- LOW site type 2–26, 3–27

## **M**

- Marking starter data 3–21

## **O**

- ODBC 1–7, 1–8
- ODBCN parameter 2–22, 2–26, 3–23, 3–26, 3–28
- ODNCP parameter 2–22, 2–26, 3–23, 3–26, 3–28
- OuterEdge Server 1–7, 1–9, 1–11
  - Activating the database 2–8, 3–7, Index–1
  - Activating the site 2–25, 3–9, 3–25
  - Preparing the database 3–6, 3–9
  - Preparing to activate the database 3–26
  - Setting up replication 2–3, 3–4
- OuterEdge Workstation
  - Activating the database 2–8, 3–7, Index–1
  - Activating the site 3–9
  - Preparing the database 3–6, 3–9

## **P**

- PDA 1–7, 1–8
- PDR 1–7, 1–8
- PDRE 1–7, 1–8, 1–9, 1–11, 2–12, 3–12
- PDUSER 2–21, 3–22, 3–26

- PeerDirect Replication Designer 2–19, 3–21, 3–22
  - PeerDirect Replication Network 1–5
    - Administering 2–28, 3–29
    - Planning 2–5, 2–7, 3–5, 3–7
  - PeerDirect site types 2–25, 2–26, 3–27
  - Preparing for activation 2–19, 3–21
    - How to steps 2–20
  - Preparing the InnerEdge Server database 2–6, 2–8, 3–6, 3–9
  - Preparing the OuterEdge Server database 3–6, 3–9
  - Preparing the OuterEdge Workstation 3–9
  - Preparing to activate
    - How to steps 3–22
  - Preparing to activate the OuterEdge Server database 3–26
  - Primary keys 2–10, 3–11
  - Progress
    - Database considerations 2–9
    - Heterogeneous 1–2
    - Homogeneous 1–2
    - PDUSER 2–12
    - Primary keys 2–10
    - Products 1–4
    - Proutil 2–20
    - Version 1–4
- R**
- Recovery database 1–12
  - Replication
    - Setting up 2–2, 3–2
  - Replication network types 3–10
  - Replication rules 2–18, 3–20
    - Defining 2–5, 2–7, 3–5, 3–8
    - Exporting 2–5, 2–7, 2–20, 3–5, 3–8, 3–24
  - Replication sessions
    - Initiating 2–27, 3–29
  - Rule (RUL) file 2–22, 3–22, 3–26
- S**
- Selecting database tables 2–15, 3–16
  - Selecting databases 2–14, 3–15
  - Server 1–12
  - Setting up a PeerDirect Replication Network
    - Task sequence 2–4, 3–4
  - Setting up replication 2–2, 3–2
    - OuterEdge Servers 2–3, 3–4
    - OuterEdge Workstations 2–3, 3–4
    - Task sequence 2–4, 3–4
  - Setup task sequence 2–4, 3–4
  - SITETYPE parameter 2–25, 2–26, 3–27
  - Source 1–12
  - SQL-92 1–7
  - Supported databases 1–11
  - System tables 3–21
- T**
- Tables
    - Selecting 2–15, 3–16

## **U**

Unsupported data types 2–10, 3–11

## **W**

WAN 1–8

Warm standby 1–12