DataDirect TECHNOLOGIES

SequeLink®

Administrator's Guide

Release 6.0
April 2008

© 2008 Progress Software Corporation. All rights reserved. Printed in the U.S.A.

DataDirect, DataDirect Connect, DataDirect Connect64, DataDirect Spy, DataDirect Test, DataDirect XML Converters, DataDirect XQuery, OpenAccess, SequeLink, Stylus Studio, and SupportLink are trademarks or registered trademarks of Progress Software Corporation or one of its subsidiaries or affiliates in the United States and other countries. Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. MySQL and MySQL Enterprise are registered trademarks of MySQL AB in the United States, the European Union and other countries.

Other company or product names mentioned herein may be trademarks or registered trademarks of their respective companies.

DataDirect products for the Microsoft SQL Server database:

These products contain a licensed implementation of the Microsoft TDS Protocol.

DataDirect Connect for ODBC, DataDirect Connect64 for ODBC, and DataDirect SequeLink include:

ICU Copyright © 1995-2003 International Business Machines Corporation and others. All rights reserved. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

Software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http:/www.openssl.org/). Copyright © 1998-2006 The OpenSSL Project. All rights reserved. And Copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

DataDirect SequeLink includes:

Portions created by Eric Young are Copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All Rights Reserved. OpenLDAP, Copyright © 1999-2003 The OpenLDAP Foundation, Redwood City, California, US. All rights reserved.

DataDirect OpenAccess SDK client for ODBC, DataDirect OpenAccess SDK client for ADO, DataDirect Open Access SDK client for JDBC and DataDirect OpenAccess SDK server include DataDirect SequeLink.

No part of this publication, with the exception of the software product user documentation contained in electronic format, may be copied, photocopied, reproduced, transmitted, transcribed, or reduced to any electronic medium or machine-readable form without prior written consent of DataDirect Technologies.

Licensees may duplicate the software product user documentation contained on a CD-ROM or DVD, but only to the extent necessary to support the users authorized access to the software under the license agreement. Any reproduction of the documentation, regardless of whether the documentation is reproduced in whole or in part, must be accompanied by this copyright statement in its entirety, without modification.

Table of Contents

Part 1: Configuring and Managing SequeLink® Services

Part 3: Configuring SequeLink® in Your Environment

List of Tables

Preface

This book is your guide to configuring and managing DataDirect SequeLink® 6.0 from DataDirect Technologies. Read on to find out more about your SequeLink environment and how to use this book.

What Is DataDirect SequeLink®?

DataDirect SequeLink is a middleware product that provides point-to-point connections from a client to a server for the latest data access standards, including Open Database Connectivity (ODBC), JDBC, ActiveX Data Objects (ADO), and ADO.NET.

In this documentation, references to SequeLink Server and SequeLink Client apply to both the 32-bit and 64-bit versions. Information that applies to a specific version of SequeLink Server or SequeLink Client is identified.

Using This Book

This book assumes you are familiar with your operating system and its commands; the concept of directories; the management of user accounts and security access; and your network protocol and its configuration.

This book contains the following information:

- Chapter 1 "Introduction" on page 29 introduces some concepts to help you understand how to configure and manage your SequeLink environment.

Part 1: Configuring and Managing SequeLink Services

- Chapter 2 "Using the SequeLink® Manager Snap-in" on page 47 describes how to use the SequeLink Manager Snap-in.
- Chapter 3 "Configuring SequeLink® Services Using the SequeLink Manager Snap-in" on page 63 describes how to create and manage server data sources with the SequeLink Manager Snap-in.
- Chapter 4 "Managing Data Access Activity Using the SequeLink® Manager Snap-in" on page 91 describes the tasks you perform to manage and monitor SequeLink service activity using the SequeLink Manager Snap-in.
- Chapter 5 "Using the SequeLink® Manager Command-Line Administrator" on page 99 describes how to use the SequeLink Manager Command-Line Tool, issue SequeLink Manager commands, and lists some commonly used SequeLink Manager commands.
- Chapter 6 "Using the SequeLink® Manager for z/OS" on page 113 describes how to use the SequeLink Manager for z/OS.
- Chapter 7 "Configuring SequeLink® Services Using the SequeLink® Manager for z/OS" on page 129 describes the tasks you may need to perform to configure and manage SequeLink Server for z/OS services and data access activities locally from a z/OS machine.

Part 2: Configuring and Managing SequeLink Clients

- Chapter 8 "Configuring the ODBC Client" on page 169 describes the tasks you may need to perform to configure and manage the SequeLink Client *for* ODBC.
- Chapter 9 "Configuring the ADO Client" on page 211 describes the tasks you may need to perform to configure and manage the SequeLink Client *for* ADO.
- Chapter 10 "Configuring the JDBC Client" on page 251 describes the tasks you may need to perform to configure and manage the SequeLink Client *for* JDBC.
- Chapter 11 "Configuring the .NET Client" on page 265 describes the tasks you may need to perform to configure and manage the SequeLink Client *for* .NET.

Part 3: Configuring SequeLink in Your Environment

- Chapter 12 "Configuring Transliteration" on page 277 describes how configure transliteration for SequeLink Server.
- Chapter 13 "Configuring SequeLink® Security" on page 291 offers an overview of SequeLink security options and describes how to configure SequeLink security for Linux, UNIX, and Windows and z/OS platforms.
- Chapter 14 "Configuring the SequeLink® Proxy Server" on page 351 describes how to configure SequeLink security for Java environments.
- Chapter 15 "Configuring SequeLink® Services for Your Database" on page 381 describes how to configure SequeLink services for specific databases.
- Chapter 16 "Using LDAP with the SequeLink® Clients" on page 427 explains how SequeLink Clients use LDAP directories to retrieve connection information and describes how to create and update LDAP entries for SequeLink services.

Part 4: Appendixes

- Appendix A "z/OS Workload Manager (WLM) Classification" on page 437 describes the information used by SequeLink Server to classify WLM enclaves.
- Appendix B "SequeLink® Manager Commands" on page 439 lists all available SequeLink Manager commands.
- Appendix C "Operator Interface Commands for z/OS" on page 479 lists all available Operator Interface commands by category.
- Appendix D "SequeLink® Service Attributes" on page 491 lists the SequeLink Manager attributes you can use to configure and manage your SequeLink environment.
- Appendix E "SequeLink® Events" on page 589 lists and defines the SequeLink events, the attributes associated with events, and explains how to write a filter for an event.
- Appendix F "Internationalization, Localization, and Unicode" on page 597 provides an overview of how internationalization, localization, and Unicode relate to each other.

NOTE: This book refers the reader to Web URLs for more information about specific topics, including Web URLs not maintained by DataDirect Technologies. Because it is the nature of Web content to change frequently, DataDirect Technologies can guarantee only that the URLs referenced in this book were correct at the time of publishing.

SequeLink® Documentation

The following table provides a guide for finding information in your SequeLink documentation:

For information about...	Go to...
SequeLink concepts and planning your SequeLink environment	*Getting Started with SequeLink*
Installing the SequeLink middleware components	*SequeLink Installation Guide*
Administering your SequeLink environment	*SequeLink Administrator's Guide*
Developing ODBC, ADO, JDBC, and .NET applications for the SequeLink environment	*SequeLink Developer's Reference*
Troubleshooting and referencing error messages	*SequeLink Troubleshooting Guide and Reference*

HTML Version

All of these books can be placed on your system as HTML-based online help during a normal installation of the product. They are located in the help subdirectory of the product installation directory. To use the help, you must have one of the following browsers installed:

- Internet Explorer 5.*x* or higher
- Netscape 4.*x*, 6.1, or higher
- FireFox 1.0 or higher

If you choose to install the online books, you can access the entire help system by selecting the help icon that appears in the DataDirect program group.

UNIX

On UNIX and Linux platforms, if you want the help files, copy the /bookshtml subdirectory from the product DVD to a local directory.

To open the help system from a command-line environment, at a command prompt, enter:

```
browser_exe my_local_dir/bookshtml/help.htm
```

where *browser_exe* is the name of your browser executable and *my_local_dir* is the path to the product installation directory.

After the browser opens, the left pane displays the Table of Contents, Index, and Search tabs for the entire documentation library. When you have opened the main screen of the help system in your browser, you can bookmark it in the browser for quick access later.

NOTE: Security features set in your browser can prevent the help system from launching. A security warning message is displayed. Often, the warning message provides instructions for unblocking the help system for the current session. To allow the help system to launch without encountering a security warning message, the security settings in your browser can be modified. Check with your system administrator before disabling any security features.

Help is available from the setup dialog box for the ODBC driver and ADO data provider. When you click **Help**, your browser opens to the correct topic in the help system, without opening the help

Table of Contents. A grey toolbar appears at the top of the browser window.

This tool bar contains previous and next navigation buttons.

PDF Version

DataDirect product documentation is also provided in PDF format, which allows you to view it, perform text searches, or print it. You can view the PDF documentation using the Adobe Acrobat Reader. The PDF documentation is available on the product DVD and also on the DataDirect Technologies Web site:

http://www.datadirect.com/support/product_info/proddoc_product/index.ssp

You can download the entire library in a compressed file. When you uncompress the file, it appears in the correct directory structure.

If you want to copy the documentation library from the product DVD, you must maintain the same directory structure that is on the DVD.

- **To copy all product books**, copy the entire \bookspdf directory to your local or network drive.
- **To copy a specific book**, copy that book's directory structure (beneath the \bookspdf subdirectory) to your local or network drive. For example, to copy the *SequeLink Administrator's Guide*, you would copy the entire \admin subdirectory:

  ```
  \bookspdf\admin
  ```

Maintaining the correct directory structure allows cross-book text searches and cross-references. If you download or copy the books individually outside of their normal directory structure, their cross-book search indexes and hyperlinked cross-references to other volumes will not work. You can view a book individually, but it will not automatically open other books to which it has cross-references.

To help you navigate through the library, a file, called **books.pdf**, is provided. This file lists each online book provided for the product. We recommend that you open this file first and, from this file, open the book you want to view.

Typographical Conventions

This book uses the following typographical conventions:

Convention	Explanation
italics	Introduces new terms you may not be familiar with, and is used occasionally for emphasis.
bold	Emphasizes important information. Also indicates button, menu, and icon names on which you can act. For example, click **Next**.
UPPERCASE	Indicates keys or key combinations you can use. For example, press the ENTER key.
`monospace`	Indicates syntax examples, values that you specify, or results that you receive.
`monospaced italics`	Indicates names that are placeholders for values you specify; for example, *`filename`*.
forward slash /	Separates menus and their associated commands. For example, Select **File** / **Copy** means to select Copy from the File menu.
vertical rule |	Indicates an OR separator to delineate items.

Convention	Explanation
brackets [ ]	Indicates optional items. For example, in the following statement: SELECT [DISTINCT], DISTINCT is an optional keyword.
braces { }	Indicates that you must select one item. For example, {yes | no} means you must specify either yes or no.
ellipsis . . .	Indicates that the immediately preceding item can be repeated any number of times in succession. An ellipsis following a closing bracket indicates that all information in that unit can be repeated.

Environment-Specific Information

This book supports users of various operating environments. Where it provides information that does not apply to all supported environments, the following symbols are used to identify that information:

Symbol	Environment
	Windows. Information specific to the Microsoft Windows 2000, Windows Server 2003, Windows XP, and Windows Vista environments is identified by the Windows symbol.
UNIX	*UNIX and Linux*. Information specific to Linux and UNIX environments is identified by this symbol, which applies to all Linux and UNIX environments supported. UNIX is a registered trademark of The Open Group in the United States and other countries.
z/OS	*z/OS*. Information specific to z/OS environments is identified by the characters z/OS.

Contacting Technical Support

DataDirect Technologies offers a variety of options to meet your technical support needs. Please visit our Web site for more details and for contact information:

http://support.datadirect.com

The DataDirect Technologies Web site provides the latest support information through our global service network. The SupportLink program provides access to support contact details, tools, patches, and valuable information, including a list of FAQs for each product. In addition, you can search our Knowledgebase for technical bulletins and other information.

To obtain technical support for an evaluation copy of the product, go to:

http://www.datadirect.com/support/eval_help/index.ssp

or contact your sales representative.

When you contact us for assistance, please provide the following information:

- The serial number that corresponds to the product for which you are seeking support, or a case number if you have been provided one for your issue. If you do not have a SupportLink contract, the SupportLink representative assisting you will connect you with our Sales team.
- Your name, phone number, email address, and organization. For a first-time call, you may be asked for full customer information, including location.
- The DataDirect product and the version that you are using.
- The type and version of the operating system where you have installed your DataDirect product.

- Any database, database version, third-party software, or other environment information required to understand the problem.
- A brief description of the problem, including, but not limited to, any error messages you have received, what steps you followed prior to the initial occurrence of the problem, any trace logs capturing the issue, and so on. Depending on the complexity of the problem, you may be asked to submit an example or reproducible application so that the issue can be recreated.
- A description of what you have attempted to resolve the issue. If you have researched your issue on Web search engines, our Knowledgebase, or have tested additional configurations, applications, or other vendor products, you will want to carefully note everything you have already attempted.
- A simple assessment of how the severity of the issue is impacting your organization.

1 Introduction

This chapter introduces some concepts to help you understand how to configure and manage your SequeLink environment. For a complete discussion of planning issues, including configuration, administration, and migration issues, refer to *Getting Started with SequeLink*.

SequeLink® Server System Administration

SequeLink provides the following options for configuring and managing your SequeLink environment:

- **Local system administration** allows you to configure and manage your SequeLink environment using the SequeLink Manager installed locally on a SequeLink Server.
- **Remote system administration** allows you to configure and manage your SequeLink environment using the SequeLink Manager installed on the desktop of a networked client.

NOTE: Only SequeLink 6.0 services can be configured, managed, or monitored with the SequeLink Manager 6.0.

Local System Administration

You can use the SequeLink Manager locally from the SequeLink Server to configure and manage your SequeLink environment; however, which SequeLink Manager tool you can use locally depends on your SequeLink Server platform.

Remote System Administration

Remote system administration allows you to configure and manage your data access environment from the convenience of your desktop regardless of your SequeLink Server platform. For example, suppose you are responsible for administering an environment with distributed data access involving a variety of data stores across your enterprise, such as Oracle on UNIX and Windows, and Microsoft SQL Server on Windows as shown in Figure 1-1. You can install the SequeLink Manager Snap-in on a Windows XP or Windows Vista networked client and perform administration tasks, such as configuring SequeLink service settings, from the convenience of your desktop.

Figure 1-1. Remote System Administration for Data Access Environments

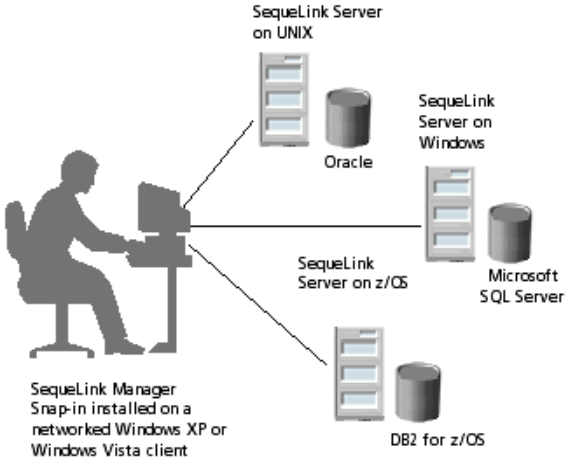

Using the SequeLink® Manager

The SequeLink Manager tool can be used to perform administrative and monitoring requests.

Administrative Requests

The type of administrative requests you can issue to a SequeLink Agent and the SequeLink Manager tool you can use to issue the requests depends on the platform you are administering. The following list describes the types of administrative requests you can issue:

- **Configuration**
  - Creating and managing SequeLink services
  - Creating and managing SequeLink server data sources
  - Configuring monitoring profiles, which determine the data access events that can be monitored (viewed) using the SequeLink Manager
  - Configuring event-tracing profiles, which determine the data access events that are written to an event trace file
- **Management**
  - Starting and stopping SequeLink services
  - Stopping active data access user sessions
  - Reviewing traced events to analyze a problem during an earlier data access activity

Monitoring Requests

The SequeLink Manager allows you to perform the following monitoring tasks:

- Viewing details about active services
- Viewing active user sessions and information about live data access activities

For example, you can easily view the number of transactions that have been processed or the number of rows that have been fetched by all user sessions. Also, if a user session is not performing correctly (such as the session repeatedly fetches thousands of rows), you can use the SequeLink Manager to identify and end that specific user session.

Additionally, the SequeLink Manager allows you to troubleshoot previous events. For example, if an error occurs during a nightly data processing job, you can view an event trace to troubleshoot the problem.

SequeLink® Manager Implementations

SequeLink provides the following implementations of the SequeLink Manager:

- **SequeLink Manager Snap-in** is a GUI designed as a snap-in to the Microsoft Management Console (MMC). It can be used to configure and manage SequeLink services, and to monitor data access activity.
- **SequeLink Manager Command-Line Tool** is a command-line interface that can be used to configure and manage SequeLink services. Similarly, it can be used to monitor data access activity.
- **SequeLink Manager for z/OS** is an ISPF dialog tool that can be used to create data access services on the z/OS platform, as well as configure and manage SequeLink services and monitor

data access activity on z/OS. It can be installed only on z/OS platforms.

Table 1-1 shows the platforms on which you can install and run the different implementations of the SequeLink Manager.

Table 1-1. Installation Platforms for the SequeLink Manager Tools

SequeLink Manager	Windows	Linux/ UNIX	z/OS
SequeLink Manager Snap-in	X		
SequeLink Manager Command-Line Tool	X	X	
SequeLink Manager for z/OS			X

About SequeLink® Services and Data Sources

SequeLink Server installs the following server software service components to provide data connectivity, performance, and administration for two-tier client/server and *n*-tier Web/application server environments:

- **SequeLink data access services** handle data access requests from any SequeLink Client. Multiple SequeLink data access services can run on the same SequeLink Server. For example, SequeLink Server for Oracle and SequeLink Server for Microsoft SQL Server can run side-by-side on the same machine.
- **SequeLink Agent services** carry out configuration, management, and monitoring requests from any SequeLink Manager. The SequeLink Agent can service multiple SequeLink services on the same SequeLink Server.

When you complete the installation of the SequeLink Server software as documented in the *SequeLink Installation Guide*, a SequeLink data access service is configured for the type of SequeLink Server you installed (for example, SequeLink Server for Oracle). In addition, a SequeLink Agent is configured to handle configuration, management, and monitoring requests from any SequeLink Manager.

Data Sources

SequeLink uses two types of data sources—server data sources and client data sources.

Server Data Sources

Server data sources are data sources configured on the SequeLink Server. These data sources contain settings that affect how the SequeLink service operates and settings that affect how data is accessed by SequeLink Clients. Centralizing this information on the server, instead of distributing it among hundreds of SequeLink Clients, provides easier management of your entire data access infrastructure. When you install SequeLink Server, a default server data source, named *Default*, is automatically created on the server, using the values you specified during installation. If necessary, you can modify the definition of the default server data source. Server data sources apply to a SequeLink service. Each SequeLink service can have multiple server data sources.

Unlike server data sources, client data sources are minimal data sources configured on the SequeLink Client that contain connection instructions to a SequeLink data access service. The data access functionality of a session is governed by a set of data source and service attributes for the SequeLink data access service. These attributes are configured on the server.

If you do not specify a server data source for your connection, the attributes of the Default server data source govern the data access functionality of the connection. For example, if you configured two server data sources named DS1 and DS2 as shown in Figure 1-2, and configured a SequeLink Client that did not specify a particular server data source, the data access functionality of the connection between the SequeLink Client and the SequeLink Server would be governed by the Default server data source.

Figure 1-2. SequeLink Clients Specifying Server Data Sources

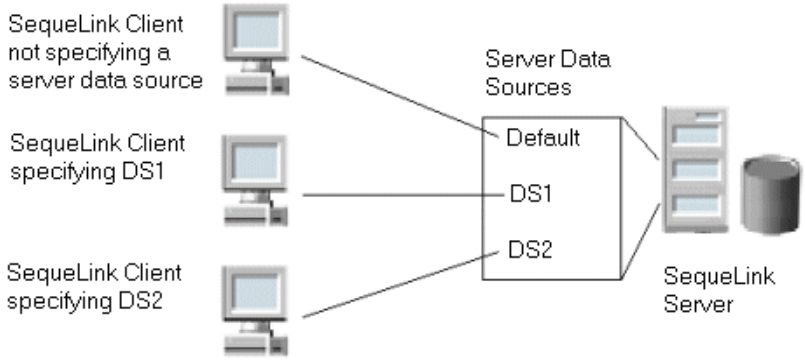

Server data sources use SequeLink service attributes to define functionality. For example, if you set DataSourceReadOnly= Select, the client application will only be able to perform Select statements when using that service.

NOTES:

- SequeLink service attributes beginning with "DataSource", such as DataSourceReadOnly, are server data source attributes. Server data source attributes are always dynamic.
- SequeLink Agent services do not have server data source attributes.

Only commonly used service attributes are included in the default configuration of a SequeLink service. To configure other attributes, you must add that attribute explicitly to your SequeLink configuration.

For instructions on configuring server data sources using the:

- SequeLink Manager MMC Snap-in, see Chapter 3 "Configuring SequeLink® Services Using the SequeLink Manager Snap-in" on page 63.
- SequeLink Manager Command-Line Tool, see Chapter 5 "Using the SequeLink® Manager Command-Line Administrator" on page 99.
- SequeLink Manager for z/OS, see Chapter 7 "Configuring SequeLink® Services Using the SequeLink® Manager for z/OS" on page 129.

Client Data Sources

Client data sources are minimal data source containing the host, port, and server data source, that are configured on the SequeLink Client and provide connection instructions to a SequeLink data access service. When a client application connects to a SequeLink data access service using a SequeLink Client, the data access functionality of the session is governed by a set of data source and service attributes for the SequeLink data access service. These attributes are configured on the server. These attributes are configured on the server.

Client data sources are required when configuring the SequeLink Client *for* ODBC or the SequeLink Client *for* ADO. For SequeLink Client *for* JDBC, you can configure either a client data source or a connection URL. SequeLink Client for .NET clients do not require a client data source; instead, you configure a connection string.

See the following chapters for instructions on configuring client data sources for the SequeLink Client *for* ODBC and SequeLink Client *for* ADO Client:

- Chapter 8 "Configuring the ODBC Client" on page 169
- Chapter 9 "Configuring the ADO Client" on page 211

See "Specifying JDBC Driver Connection URLs" on page 252 for instructions on specifying connection URLs for SequeLink Client *for* JDBC. Configuring JDBC client data sources is considered an advanced topic.

See "Specifying Connection Properties" on page 265 for instructions on specifying connection options for the SequeLink Client *for* .NET.

Service Attributes

When a client application connects to a SequeLink data access service using a SequeLink Client, the data access functionality of the session is governed by a set of service and data source attributes for the SequeLink data access service.

When you create a SequeLink service, only commonly used service attributes are included in the default configuration of a SequeLink service. The newly created SequeLink service contains a default server data source called *Default*, which is configured with the default values entered during installation.

To configure other service or data source attributes, you must add each attribute explicitly to your SequeLink configuration. You configure SequeLink services and their attributes using the SequeLink Manager.

SequeLink service attributes beginning with *Service*, such as ServiceCodePageMap, are SequeLink data access service attributes that are associated with a server data source. SequeLink service attributes beginning with "DataSource," such

as DataSourceReadOnly, are SequeLink data access service attributes that are associated with a server data source.

SequeLink service attributes are static or dynamic:

- **Static attributes** require you to restart a SequeLink service when you add or change the attribute before the change becomes effective.
- **Dynamic attributes** become effective after the attribute is added or changed and the configuration is saved. Most dynamic attributes affect the behavior of a database connection; therefore, when you add or change an attribute, the new values are used for the next connection. Active connections do not use the new values.

See Appendix D "SequeLink® Service Attributes" on page 491 for a complete list and description of SequeLink service attributes.

SequeLink® Service Templates

When you install SequeLink Server, at least one SequeLink data access service is installed using default attributes for that service. Default service attributes are defined in the SequeLink service templates. Using the SequeLink Manager, you can create additional services based on the SequeLink service templates. Examples of SequeLink service templates include:

- [SequeLink 6.0] Agent Service
- [SequeLink 6.0] DB2 for LUW service
- [SequeLink 6.0] DB2 for z/OS service
- [SequeLink 6.0] Informix service
- [SequeLink 6.0] JDBC Socket service
- [SequeLink 6.0] ODBC Socket service
- [SequeLink 6.0] Oracle 10 service
- [SequeLink 6.0] SQL Server service
- [SequeLink 6.0] Sybase service

See "Creating a SequeLink® Service" on page 66 for information on creating a SequeLink service and a complete list of SequeLink service templates.

Monitoring SequeLink® Service Activity

SequeLink can monitor services, sessions, statements, and data access events. To configure what you want SequeLink to monitor, you can use the SequeLink Manager. For information about setting monitoring profiles using the:

- SequeLink Manager Snap-in, see Chapter 3 "Configuring SequeLink® Services Using the SequeLink Manager Snap-in" on page 63.
- SequeLink Manager Command-Line Tool, see Chapter 5 "Using the SequeLink® Manager Command-Line Administrator" on page 99.
- SequeLink Manager for z/OS, see Chapter 7 "Configuring SequeLink® Services Using the SequeLink® Manager for z/OS" on page 129.

You can integrate SequeLink monitoring with the Windows Performance Monitor tool, which allows you to access monitoring information from this Windows tool. See "Integrating SequeLink® Monitoring with the Windows Performance Tool" on page 85 for instructions on integrating SequeLink monitoring with the Windows Performance Monitor tool.

Event Handling

All important server actions, such as data access activity and stopping and starting the server, cause an event to be generated. Depending on which SequeLink profiles are active, the information generated by the event is displayed as it occurs

on the runtime monitor and stored in a file named the *event trace file*. By default, the event trace file is located in the *installdir*/tracing directory, where *installdir* is your SequeLink Server installation directory.

By setting profiles in your SequeLink service configuration to control which events are traced, you can inspect information generated by these events. Event tracing allows you to monitor ongoing activity, troubleshoot problems, and fine-tune your data access infrastructure. For example, if you want to monitor the number of transactions a SequeLink Server processes for capacity planning purposes, you could set a profile in the SequeLink service configuration to return only that information. In addition, the information stored in the event trace file is persisted, meaning that you can inspect the information at a later time.

Events are identified by:

- An event ID
- The service in which the event occurred
- The time the event occurred
- List of attributes and their values

Examples of information that can be monitored and traced are SQL statements, number of transactions, failures, and authentication information.

See "Configuring Event Tracing" on page 86 for instructions on configuring the events to be monitored and traced by the SequeLink Manager.

About SequeLink® Threading Models

The type of connection model you choose for your SequeLink configuration partly depends on your SequeLink Server platform, the scalability requirements, and whether you are using the distributed transaction functionality of a DBMS. SequeLink provides the following types of connection models:

- The **ThreadPool** connection model starts SequeLink with a preallocated minimum number of threads that can be increased when needed to a specified maximum number of threads. These threads can be shared by multiple SequeLink Clients connected to the SequeLink Server. This connection model provides optimum scalability—many client connections can be serviced with the same system resources on the server. It is the default connection model for all platforms. (ServiceConnectionModel=ThreadPool)

  NOTE: If you are using distributed transactions with DB2 Universal Database (UDB) on Linux, UNIX, or Windows platforms, do **not** use the ThreadPool connection model; use the Process/Connection model.

  See "Allocating the Number of Threads to the Thread Pool" on page 42 and "Returning Threads to the Thread Pool" on page 42 for more information about configuring how the thread-pool engine operates.

- The **Process/Connection** connection model creates a separate operating system process for each SequeLink Client connection request. This connection model is **not** valid on z/OS. (ServiceConnectionModel=Process/Connection)

  NOTE: If you are using distributed transactions with DB2 UDB on Linux, UNIX, or Windows platforms, use this model.

- The **Thread/Connection** connection model provides a dedicated thread for each SequeLink Client connection to a SequeLink Server. Use the Thread/Connection connection

model for client applications that are database-intensive, such as bulk load or bulk transfer applications. (ServiceConnectionModel=Thread/Connection)

See Appendix D "SequeLink® Service Attributes" on page 491 for more information about service attributes.

Allocating the Number of Threads to the Thread Pool

SequeLink can accommodate both low and high user activity by using a minimum number of pre-started threads in the thread pool that can be dynamically increased to accommodate peak user activity. When the SequeLink Server is started, the number of threads specified by the ServiceMinThreads service attribute will populate the thread pool to wait for data access requests from SequeLink Clients. If, during the working day, frequent user activity causes the number of threads specified by ServiceMinThreads to be active concurrently, SequeLink Server will dynamically create additional threads up to the number specified by the ServiceMaxThreads service attribute.

z/OS NOTE: On z/OS, a thread is equivalent to an attached TCB.

Returning Threads to the Thread Pool

SequeLink allows you to accommodate idle periods and heavy workload traffic by setting service attributes that let you govern when threads are returned to the thread pool, and consequently, when the threads become available to service other client connections.

The DataSourceThreadMaxRpc attribute specifies the maximum number of data access requests to be accepted from the same client before the thread allocated to that connection is returned to the thread pool. For example, if DataSourceThreadMaxRpc=10,

the thread will not be returned to the thread pool until after 10 requests have been made. When the time specified by the DataSourceThreadRpcTimeOut attribute has been exceeded, the thread is returned to the thread pool to serve another connection. This attribute avoids a client connection monopolizing a thread by not responding to it in a timely manner.

The ServiceThreadLockThreshold attribute specifies a percentage of the value specified by the ServiceMaxThreads attribute. When the number of active threads is less than this percentage, a connection that has executed more RPCs than the value specified by the DataSourceThreadMaxRpc attribute on the current thread is allowed to lock this thread for the time specified by the DataSourceThreadRpcTimeOut attribute. For example, if ServiceMaxThreads=10 and ServiceThreadLockThreshold=50, and only 4 threads are active, a connection can lock the thread it is using for another time period specified by the DataSourceThreadRpcTimeOut attribute.

These data source attributes ensure that the SequeLink Server can continue to process additional client data access requests, even under heavy workload conditions that can cause all threads to be active concurrently.

Part 1: Configuring and Managing SequeLink® Services

This part contains the following chapters:

- Chapter 2 "Using the SequeLink® Manager Snap-in" on page 47 describes how to use the SequeLink MMC Snap-In Administrator.
- Chapter 3 "Configuring SequeLink® Services Using the SequeLink Manager Snap-in" on page 63 describes how to create and manage server data sources with the SequeLink MMC Snap-In Administrator.
- Chapter 5 "Using the SequeLink® Manager Command-Line Administrator" on page 99 describes how to use the SequeLink Command-Line Administrator Tool, issue SequeLink Manager commands, and lists some commonly used SequeLink Manager commands.
- Chapter 4 "Managing Data Access Activity Using the SequeLink® Manager Snap-in" on page 91 describes the tasks you perform to manage and monitor SequeLink service activity using the SequeLink MMC Snap-In Administrator.
- Chapter 6 "Using the SequeLink® Manager for z/OS" on page 113 describes how to use the SequeLink Manager for z/OS.
- Chapter 7 "Configuring SequeLink® Services Using the SequeLink® Manager for z/OS" on page 129 describes the tasks you may need to perform to configure and manage SequeLink Server for z/OS services locally from a z/OS machine.

2 Using the SequeLink® Manager Snap-in

On the Windows platforms on which SequeLink Server runs, you can use the SequeLink Manager Snap-in to configure, manage, and monitor your SequeLink environment on the same machine or on a remote networked machine. The SequeLink Manager MMC Snap-in is designed as a snap-in tool to the Microsoft Management Console (MMC). This chapter describes how to use the SequeLink Manager MMC Snap-in. See "SequeLink® Server System Administration" on page 29 for more information about the SequeLink Manager.

z/OS NOTE: To configure and manage SequeLink services on z/OS or to create z/OS-specific core entities such as UID maps, DB2 interfaces, use the SequeLink Manager for z/OS. See Chapter 7 "Configuring SequeLink® Services Using the SequeLink® Manager for z/OS" on page 129 for more information. Monitoring can be performed using any SequeLink Manager.

Adding the SequeLink® Manager Snap-in to the MMC

Before you can use the SequeLink Manager MMC Snap-in to administer a remote SequeLink Server, you must add it to the MMC. You do not need to add the SequeLink Manager Snap-in to the MMC if you are administering a local SequeLink Server because a default .MSC file, which defines the local SequeLink Manager configuration, is installed when you install SequeLink Server. For local administration, you can simply open the .MSC

file in the MMC. The default .MSC file is named sladmin60.msc and is installed in the *installdir*\admin directory, where *installdir* is your SequeLink Server installation directory (for example, C:\Program Files\DataDirect\slserver60\admin).

When you add the SequeLink Manager MMC Snap-in to the MMC, you must choose a configuration option to connect to a SequeLink Agent on the same machine or connect to a SequeLink Agent on another machine. Once you have added the configuration to the MMC, you can save the configuration in an .MSC file.

To add the SequeLink Manager MMC Snap-in to the MMC:

1 Start the MMC. Select **Start** / **Run**, and type `mmc` in the Open field; then, click **OK**. An MMC console window appears.

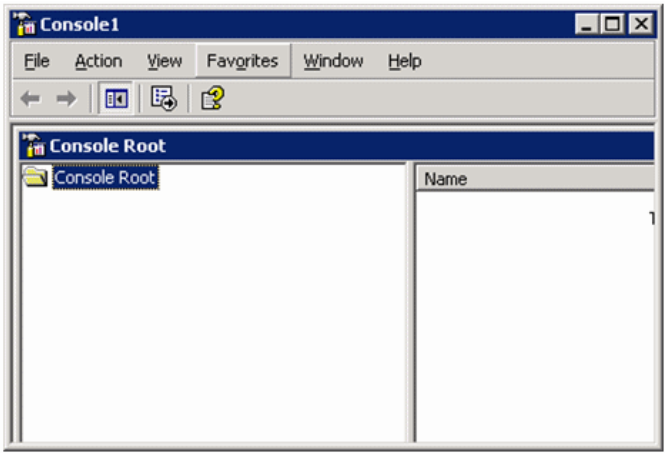

2 From the MMC toolbar, select **File / Add/Remove Snap-in**. The Add/Remove Snap-in window appears.

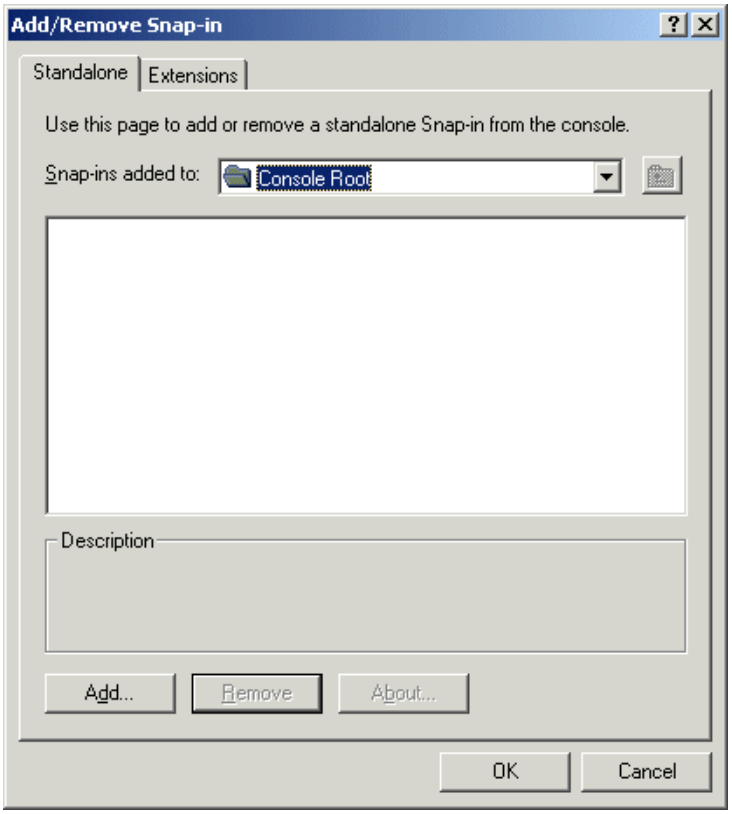

3 Click **Add**. The Add Standalone Snap-in window appears.

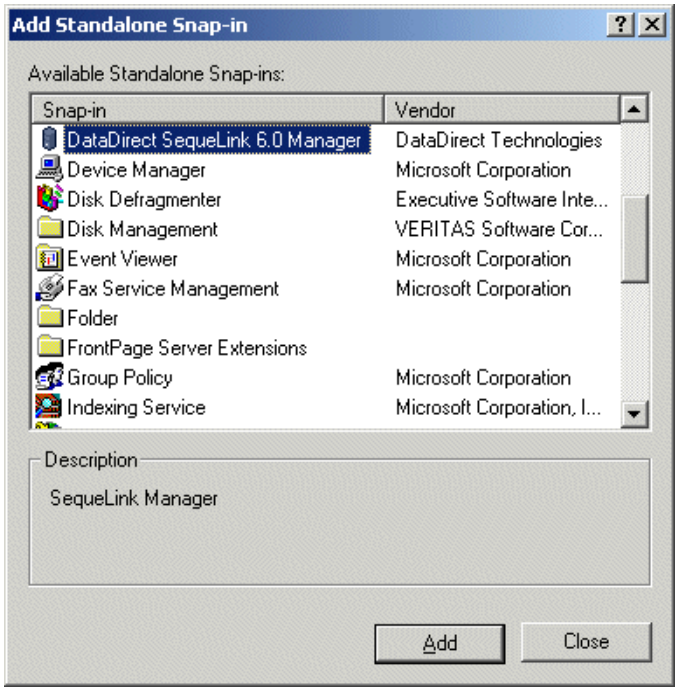

4 From the Available Standalone Snap-ins list, select the **DataDirect SequeLink 6.0 Manager**, and click **Add**. The Add SequeLink Manager window appears.

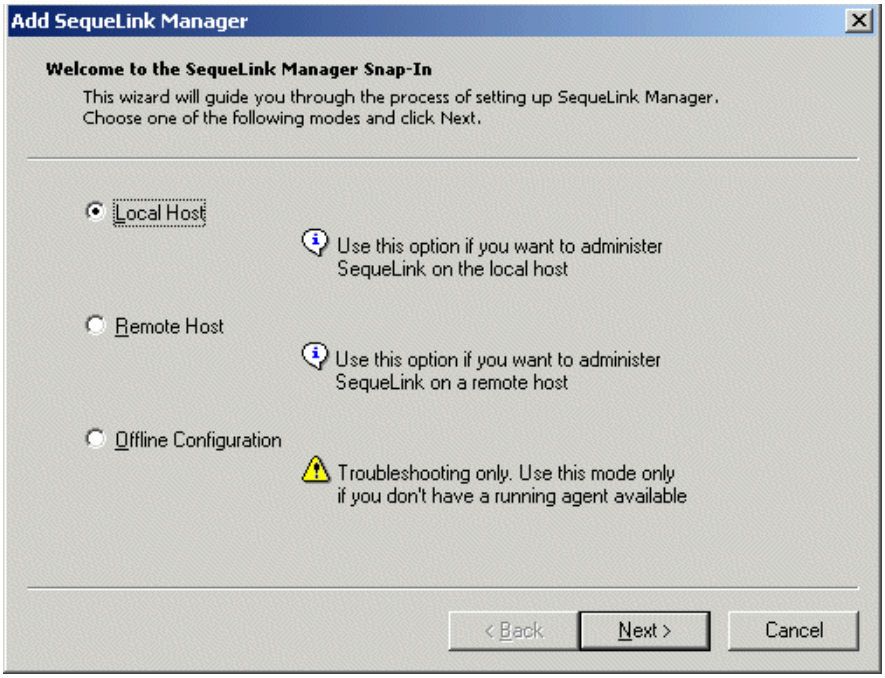

5 Choose one of the following SequeLink Manager configuration options:

- Choose the **Local Host** option to configure and manage SequeLink services on the same machine. Continue with "Local Host Configuration" on page 52.
- Choose the **Remote Host** option to configure and manage SequeLink services on another machine. Continue with "Remote Host Configuration" on page 54.

- Choose the **Offline Configuration** option to open the local configuration file in offline mode. You must specify the local configuration file in the Configuration File field. The local configuration file is *installdir*\cfg\swandm.ini where *installdir* is the SequeLink Server installation directory.

  IMPORTANT: Only use this option when instructed to do so by DataDirect Technologies technical support.

Local Host Configuration

1 On the Add SequeLink Manager window, select the **Local Host** option; then, click **Next**.

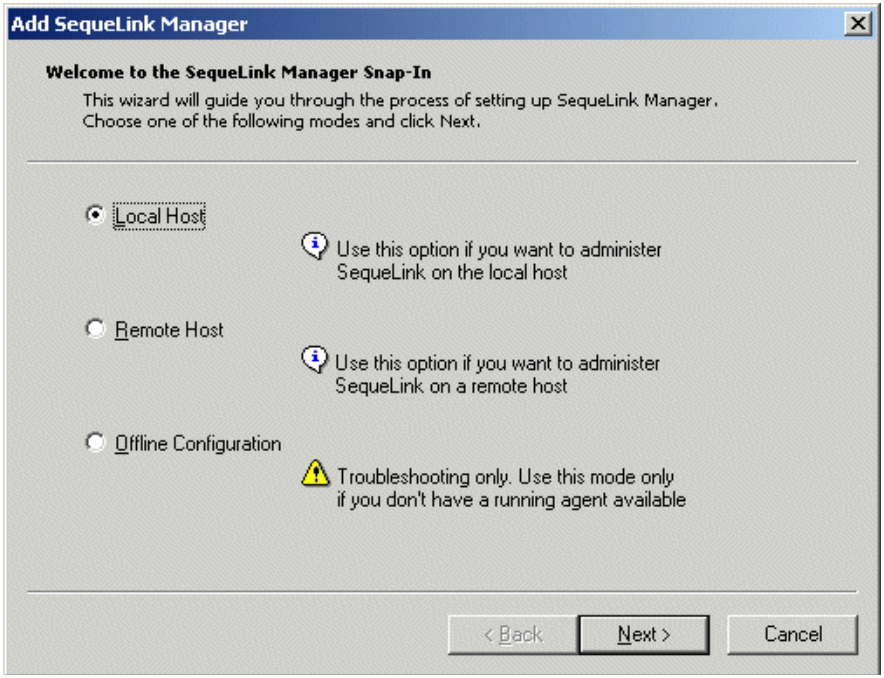

2 The Local window appears showing the location and name of the local configuration file in the Configuration File field. The local configuration file defines SequeLink Server configuration information such as SequeLink services, server data sources, and profiles.

NOTE: The default local configuration file is *installdir*\cfg\swandm.ini, where *installdir* is the SequeLink Server installation directory.

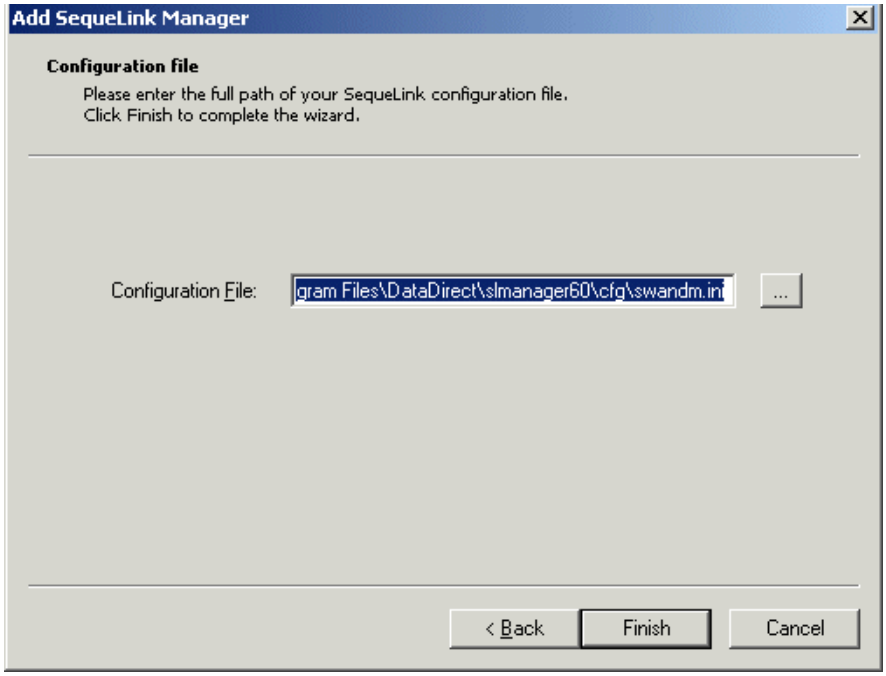

Click **Finish**. The SequeLink Manager MMC Snap-in is added to the MMC.

You can now use the SequeLink Manager MMC Snap-in to configure and manage SequeLink services on the same machine.

3 To save the SequeLink Manager MMC Snap-in to an MMC file (.MSC), select **Console** / **Save** from the MMC console window.

Remote Host Configuration

IMPORTANT: Before you can add the SequeLink Manager MMC Snap-in to the MMC as a remote host configuration, make sure that the SequeLink Agent is active on the remote host.

1 On the Add SequeLink Manager window, select the **Remote Host** option. Then, click **Next**.

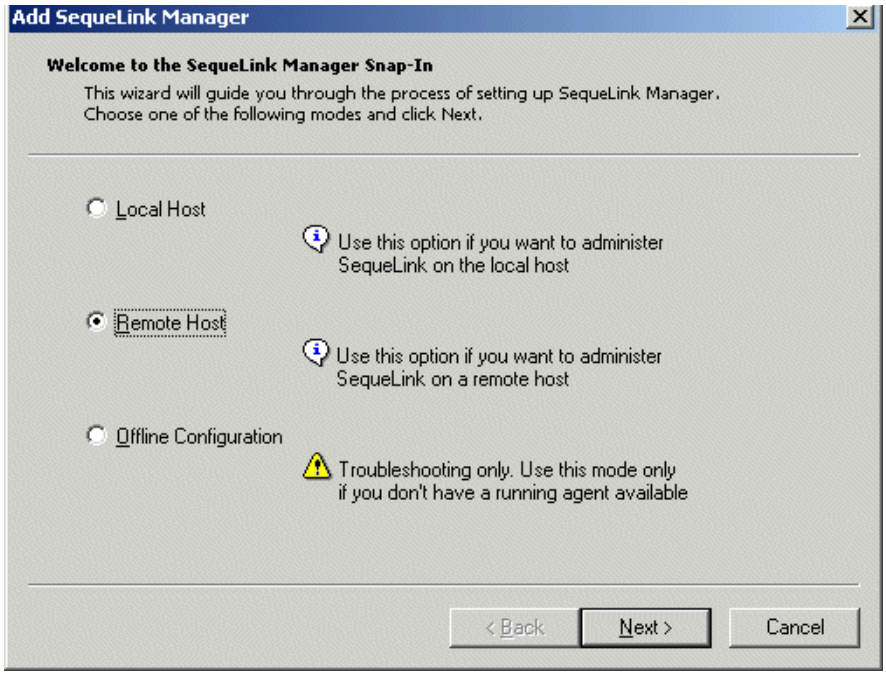

2 The Add SequeLink Manager window appears.

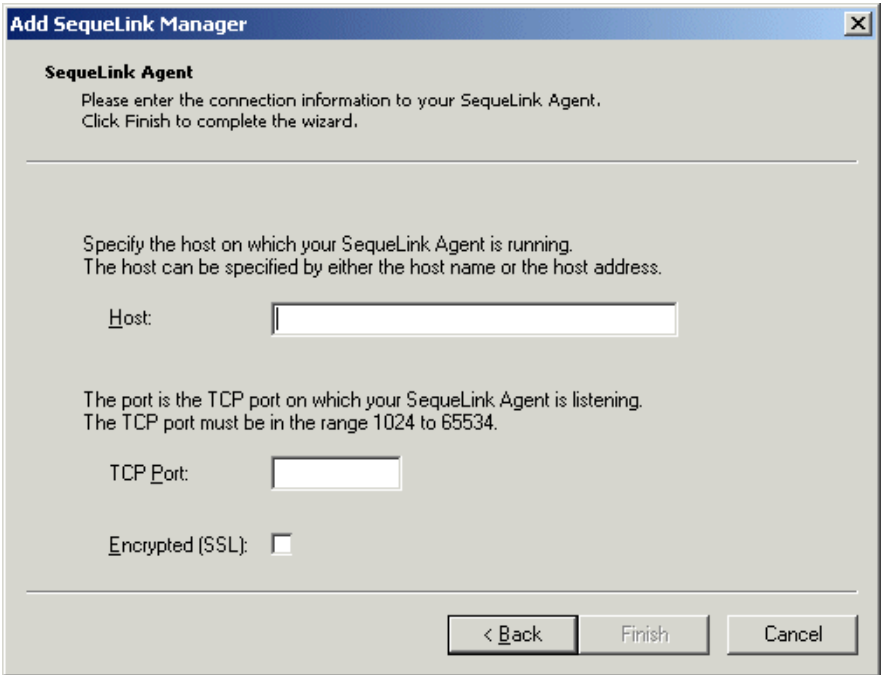

Perform the following actions:

a In the Host field, type the host name of the remote SequeLink server.

b In the TCP port field, type the TCP/IP port the SequeLink Agent is listening on for connection requests. The port you specify must be the same as the one specified for the SequeLink Agent service when the SequeLink Server was installed; the default is 19995.

c If the remote SequeLink Agent service is configured for SSL encryption, select the Encrypted (SS) check box. This check box *must* be selected when connecting to a SequeLink Agent service enabled for SSL. See "ServiceSSLEnabled" on page 584 for detailed information.

d Click **Finish**. The SequeLink Manager MMC Snap-in is added to the MMC.

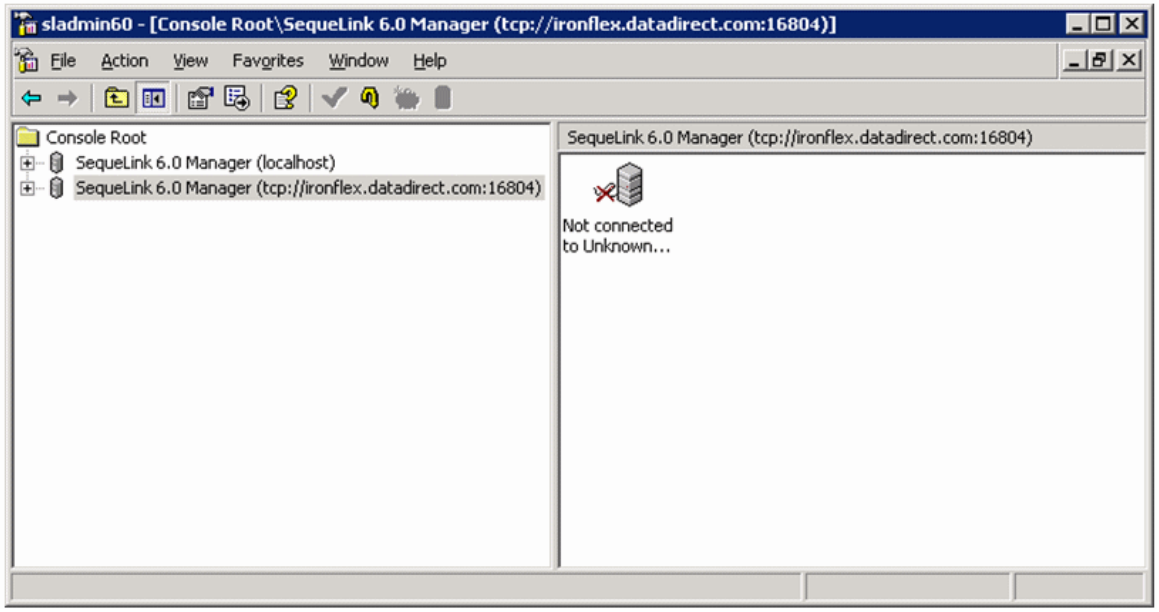

You can now use the SequeLink Manager MMC Snap-in to configure and manage SequeLink services on another machine.

3 To save the SequeLink Manager MMC Snap-in to an MMC file (.MSC), select **Console** / **Save** from the MMC console window.

Working with the SequeLink® Manager Snap-in

The MMC Console window, by default, is divided into two panes. The left pane shows the *console tree* in the Tree tab. Each node in the tree for the SequeLink Manager MMC Snap-in represents an item in your SequeLink configuration. You can expand any node in the console tree by double-clicking that node or by single-clicking the + (plus sign) for that node.

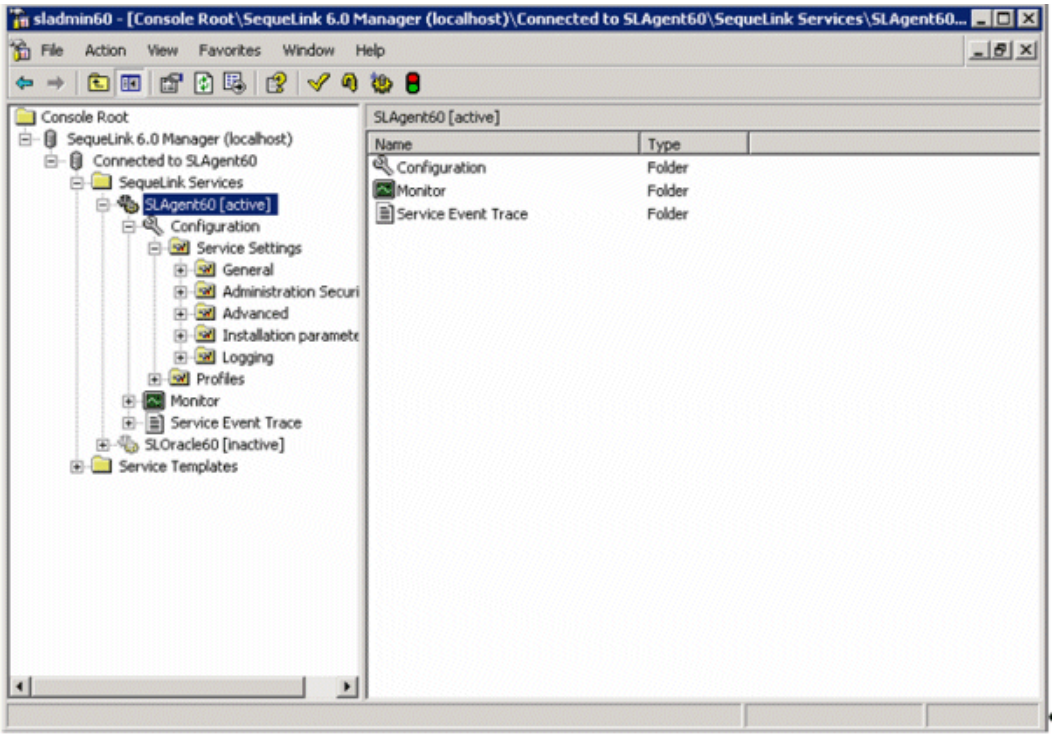

The Favorites tab of the left pane contains views you have added to your list of favorites. When you select an item in the left pane, the right pane (*details pane*) displays information about the item you selected or shows wizards you can use to perform common tasks that affect that item. To use any wizard, click the wizard icon.

Connecting to SequeLink® Agents

To view or change information about a local or remote SequeLink service, you must connect to the SequeLink Agent servicing that SequeLink service. To connect to a SequeLink Agent, double-click the SequeLink Agent in the left pane. If a user name and password are required to connect to the SequeLink Agent, the SequeLink Manager MMC Snap-in will prompt you for that information. If prompted for a user name and password, enter the appropriate user name and password in the connection dialog box.

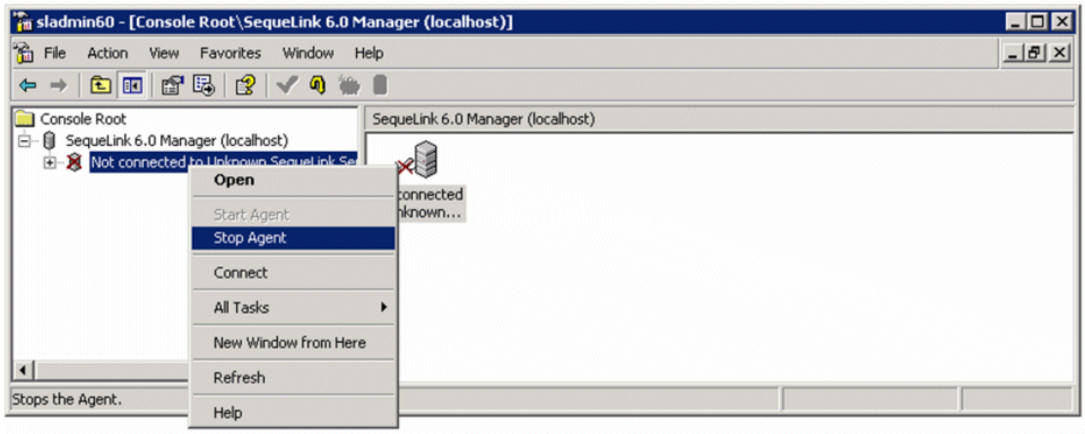

Displaying SequeLink® Service Attributes

Once you are connected to a SequeLink Agent, you can view the following types of attributes:

- Attributes of the SequeLink Agent service
- Attributes of any SequeLink data access service (active and inactive) serviced by the connected SequeLink Agent

SequeLink service attributes are logically grouped into the following categories:

- General
- Advanced
- Logging
- Administration Security
- User Security
- Application Security
- Environment
- Installation parameters
- Workarounds

To view the attributes in a specific category, expand the category by selecting that node or by selecting the + (plus sign) for that node. See Appendix D "SequeLink® Service Attributes" on page 491 for a list and description of SequeLink service attributes for more information.

Refreshing Active Information

You can refresh the active information being viewed in the SequeLink Manager Snap-in, such as active sessions or active services, by using the Refresh button on the SequeLink Manager Snap-in toolbar or by turning on the Auto Refresh option for the current session. The Auto Refresh option can only be used for the Monitor node (and all its subnodes) and the SequeLink Services node. It allows you to specify an interval in seconds to automatically refresh the display.

To turn on the Auto Refresh option:

1. Select the node you want to turn on the Auto Refresh option for, and then, select **View** / **Auto Refresh**.
2. Choose one of the following Auto Refresh options:
   - **Slow** (refreshes every 10 seconds)
   - **Normal** (refreshes every 5 seconds)
   - **Fast** (refreshes every 2 seconds)
   - **Custom** (specify a refresh interval in seconds)

The Auto Refresh option is turned off by default and is not saved when you save your configuration to an .MSC file. To turn off the Auto Refresh option, select **View** / **Auto Refresh**.

Using the SequeLink® Manager Snap-in Toolbar

Table 2-1 lists some important elements of the toolbar and describes the actions they allow you to perform.

Table 2-1. SequeLink Manager MMC Snap-in Toolbar

Item	Description	
Console menu	Commands that perform the following actions:	
	New	Creates a new console.
	Open	Opens a console.
	Save	Saves changes you make to the current console.
	Save as	Saves the current console with another name.

Shortcut Tip: Right-clicking on any item in the console tree displays a menu allowing you to perform the same actions available from the toolbar.

***Table 2-1. SequeLink Manager MMC Snap-in Toolbar** (cont.)*

Item	Description	
	Add/Remove Snap-in	Adds or removes MMC snap-ins to or from the MMC.
	Options	Options that affect how the console can be used.
Window menu	Options that affect the console window.	
Help menu	Accesses online help for the MMC.	
	Creates a new console.	
	Opens a console.	
	Saves changes to the current console.	
	Creates a new window.	
Action menu	Commands that perform actions applicable to the selected object. For example, if an active SequeLink Service is selected, the Stop command is available.	
View menu	Allows you to customize how the console appears. It also allows you to turn on the Auto refresh option for the Monitor nodes and SequeLink Services nodes for the current session.	
Favorites menu	Allows you to add views to your Favorites list.	
	Moves back to the last configurable item in the console tree.	
	Moves forward to the next configurable item in the console tree.	
	Moves up one level in the console tree.	

Shortcut Tip: Right-clicking on any item in the console tree displays a menu allowing you to perform the same actions available from the toolbar.

Table 2-1. SequeLink Manager MMC Snap-in Toolbar *(cont.)*

Item	Description
	Shows or hides the console tree.
	Displays the properties of a SequeLink service.
	Deletes an attribute when you select it while a service attribute is selected.
	Refreshes active information in the console tree.
	Accesses online help for the MMC and for the SequeLink Manager Snap-in.
	Saves all modifications to the configuration file.
	Discards all modifications and reverts to the original configuration file.
	Adds a SequeLink data access service.
	Saves changes to the SequeLink configuration.
	Starts and stops SequeLink services.
Shortcut Tip: Right-clicking on any item in the console tree displays a menu allowing you to perform the same actions available from the toolbar.	

3 Configuring SequeLink® Services Using the SequeLink Manager Snap-in

This chapter contains the following sections to help you configure SequeLink services.

Configuring SequeLink® Services

This chapter describes how to create and manage SequeLink services with the SequeLink Manager MMC Snap-in.

To do this...	See...
Start and stop SequeLink services	"Starting and Stopping SequeLink® Services" on page 65
Create a SequeLink service	"Creating a SequeLink® Service" on page 66
Delete a SequeLink service	"Deleting a SequeLink® Service" on page 71
View service attributes	"Viewing Service Attributes" on page 71
Change a service attribute	"Changing a Service Attribute" on page 72
Add a service attribute	"Adding a Service Attribute" on page 73
Delete a service attribute	"Deleting a Service Attribute" on page 74
Configure monitoring	"Configuring Monitoring" on page 79
Configure event tracing	"Configuring Event Tracing" on page 86

z/OS NOTE: SequeLink services for DB2 on z/OS must be created, started, stopped, and deleted locally using the SequeLink Manager for z/OS. See Chapter 7 "Configuring SequeLink® Services Using the SequeLink® Manager for z/OS" on page 129 for more information about creating and managing SequeLink services on z/OS.

Starting and Stopping SequeLink® Services

NOTE: You can only start and stop the SequeLink Agent locally from the SequeLink Server on which it runs.

1 From the MMC, right-click the **Connected** or **Not Connected** icon and select Start Agent to start (or Stop Agent to stop) the SequeLink Agent.

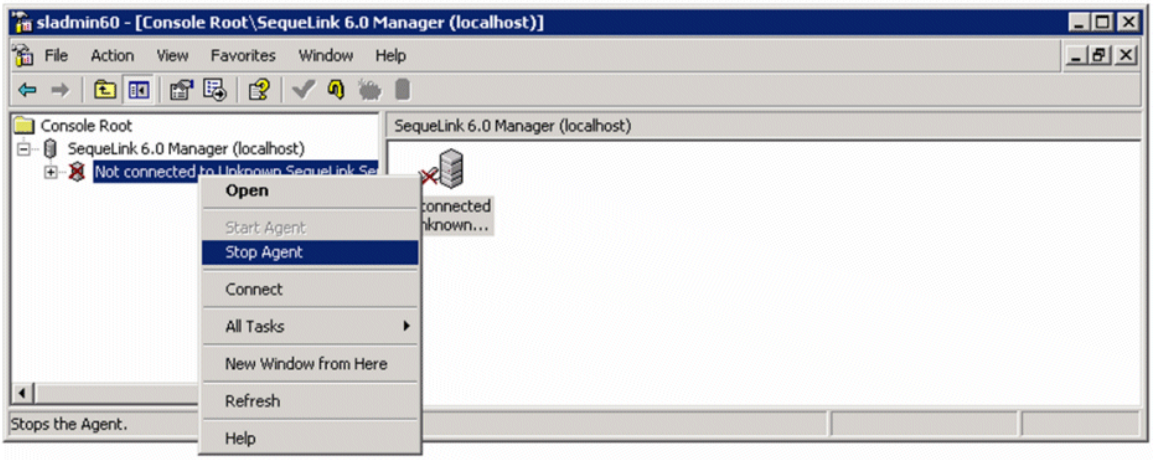

2 You are prompted to confirm that you want to stop or start the service. The following dialog box shows an example of the message that appears when you stop a service.

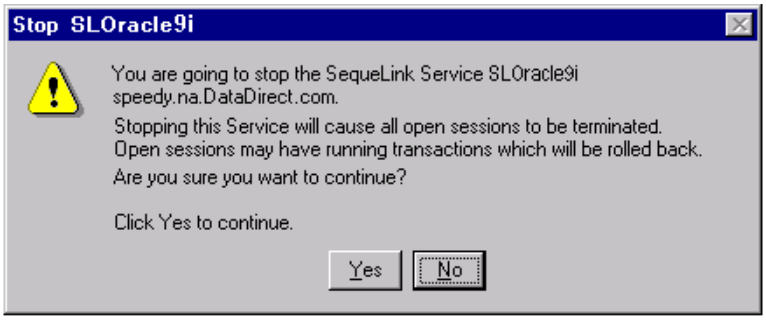

3 Click the **Yes** button to stop or to start the service.

Creating a SequeLink® Service

When you install SequeLink Server, at least one SequeLink data access service, and one or more service templates determined by the SequeLink Server, are installed. You can use the SequeLink Manager to create additional services based on the SequeLink service templates and to modify the default service attributes defined in the service templates:

- UNIX [SequeLink 6.0] Agent service
- z/OS [SequeLink 6.0] Agent service for z/OS
- z/OS [SequeLink 6.0] DB2 service for z/OS
- UNIX [SequeLink 6.0] DB2 UDB LUW service
- UNIX [SequeLink 6.0] DB2 UDB LUW service (enhanced code page support)
- UNIX [SequeLink 6.0] Informix service (32-bit only)
- UNIX [SequeLink 6.0] JDBC Socket service (32-bit only)
- UNIX [SequeLink 6.0] ODBC Socket service
- [SequeLink 6.0] ODBC Socket service (enhanced code page support)
- UNIX [SequeLink 6.0] ODBC Socket service (enhanced code page support - UTF8 encoding)
- UNIX [SequeLink 6.0] ODBC Socket service (enhanced code page support - UTF16 encoding)
- UNIX [SequeLink 6.0] Oracle 9 service (32-bit only)
- UNIX [SequeLink 6.0] Oracle 9 service (enhanced code page support - 32-bit only)
- UNIX [SequeLink 6.0] Oracle 10 service
- UNIX [SequeLink 6.0] Oracle 10 service (enhanced code page support)
- [SequeLink 6.0] SQL Server service
- [SequeLink 6.0] SQL Server service (enhanced code page support)
- UNIX [SequeLink 6.0] Sybase service
- [SequeLink 6.0] Sybase service (enhanced code page support)

UNIX [SequeLink 6.0] Sybase service (enhanced code page support - UTF8 encoding)

UNIX [SequeLink 6.0] Sybase service (enhanced code page support - UTF16 encoding)

NOTE: Unicode ODBC drivers can be written with either UTF-8 or UTF-16 encoding. When configuring SequeLink Server for ODBC Socket, choose the appropriate SequeLink service for the ODBC driver that you will be using. See Chapter 12 "Configuring Transliteration" on page 277 for more information about choosing a template.

z/OS NOTE: SequeLink DB2 services for z/OS must be created with the SequeLink Manager for z/OS. See Chapter 7 "Configuring SequeLink® Services Using the SequeLink® Manager for z/OS" on page 129 for more information.

In most cases, the SequeLink service templates provide a configuration that can be used without any modification. Not all SequeLink service attributes are defined in the templates. See Appendix D "SequeLink® Service Attributes" on page 491 for a list of all SequeLink service attributes.

To create a SequeLink data access service:

1. In the console tree, select the SequeLink Agent to service the new SequeLink data access service.
2. In the Details pane, click the **Create a Service** wizard. The wizard prompts you to choose the type of service to create.

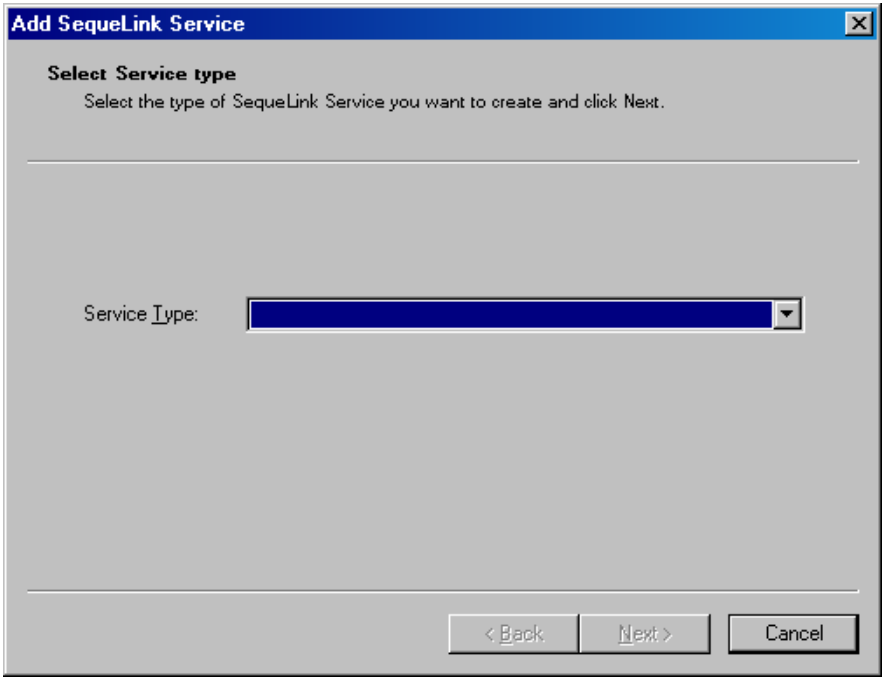

From the Service drop-down list, select the type of SequeLink service you want to create; then, click **Next**.

3 The wizard prompts you for service information.

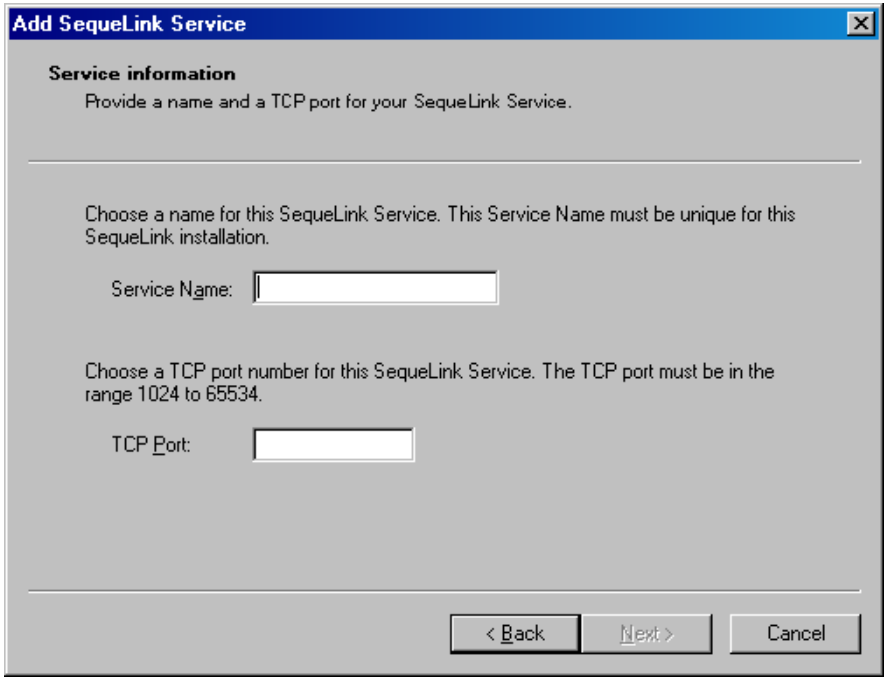

Perform the following actions:

a In the Service Name field, type the service name you want to use for the new service. The service name must be unique (not used by another service).

b In the TCP port field, type the number of the TCP/IP port on which the new service will be listening. The port must be an available port (cannot be used by another service).

c Click **Next**.

4 The wizard prompts you to register the service on the host machine. The default is to register the service, which makes the service information available to the operating system. This is required on Windows if you want to start your service through the SequeLink Manager or the Service Control Manager.

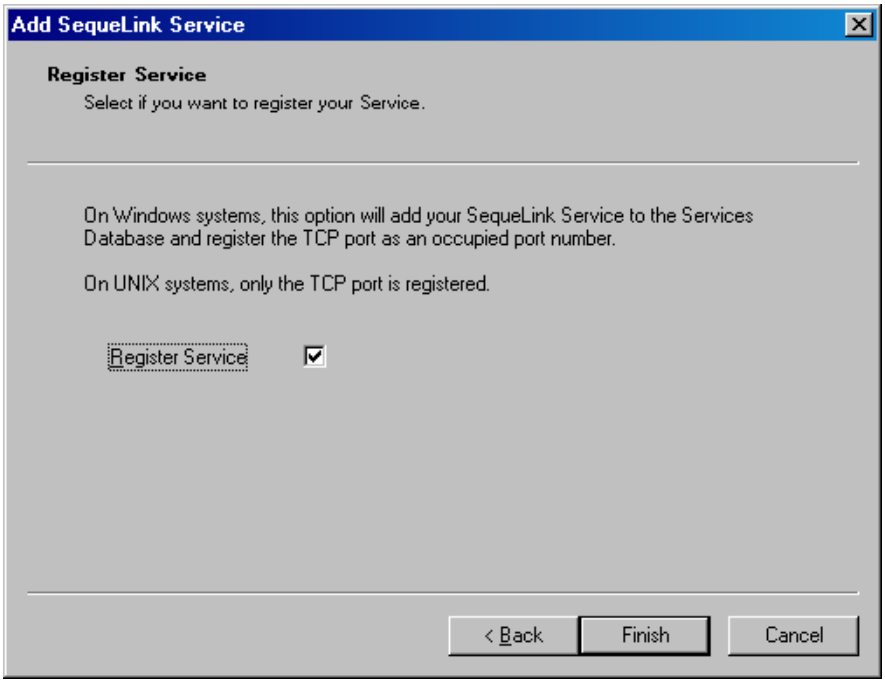

5 Click **Finish** to create the new SequeLink service.

Deleting a SequeLink® Service

NOTE: Before deleting a SequeLink service, you must stop the service you want to delete.

1 Right-click the service you want to delete in the left pane, and select **Delete**.

2 You are prompted to confirm the deletion. Click **OK** to confirm. The service is deleted.

Viewing Service Attributes

To view SequeLink service attributes, select an attribute category from the Service Settings node. The Details pane shows all attributes in that category and their current values. For example, if you select the General category, the ServiceDescription and ServiceConnectInfo attributes are displayed in the Details pane.

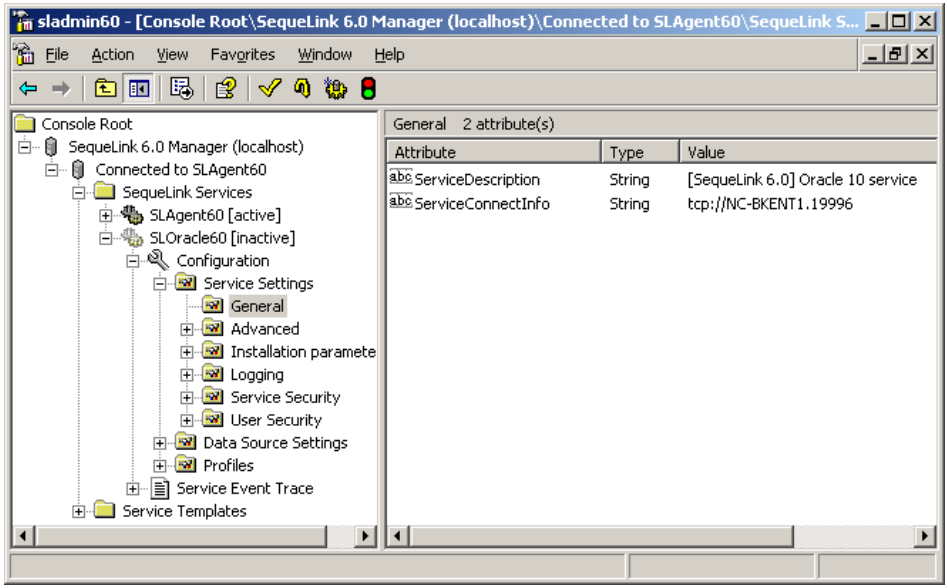

NOTE: Only commonly used attributes are included in the configuration of a newly created SequeLink service. To configure other attributes, you must add the attribute explicitly to your SequeLink configuration.

See Appendix D "SequeLink® Service Attributes" on page 491 for a description of the service attributes.

Changing a Service Attribute

1 Right-click any service attribute and select **Properties**. The properties window for that attribute appears.

For example, to change the ServiceMaxThreads attribute, right-click **ServiceMaxThreads**, and select **Properties**. The ServiceMaxThreads Properties window appears.

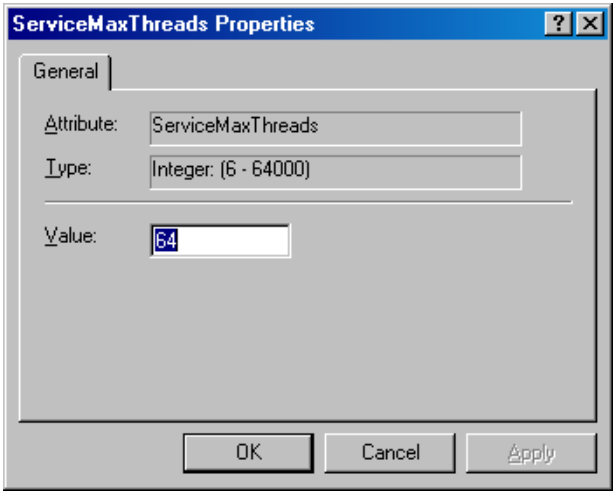

2 Type a new value for the attribute in the Value field, and click **OK**. The attribute is changed.

3 Save the configuration.

NOTE: When you add or change a static attribute, you must restart the SequeLink service before the change takes affect.

Dynamic attributes become effective after the attribute is added or changed and the configuration is saved.

Appendix D "SequeLink® Service Attributes" on page 491 describes the service attributes and specifies whether each attribute is static or dynamic.

Adding a Service Attribute

1 Right-click any service attribute category, and select **New** / **Attribute**. The New Attribute window appears.

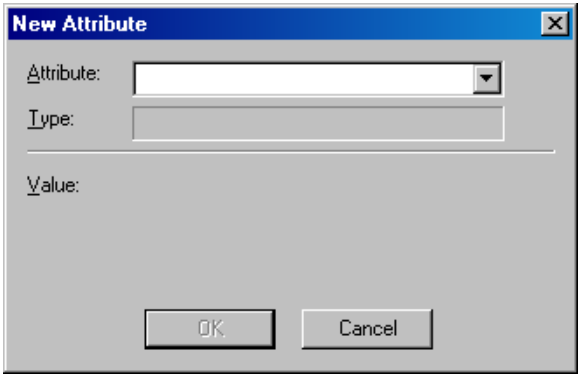

NOTE: If you do not know the attribute category, you can right-click the **Service Settings** node, and select **New** / **Attribute**. In this case, the drop-down list displays all attributes.

2 From the Attribute drop-down list, select the attribute you want to add to the service. The Type field adjusts to show the type of value required.

NOTE: If an attribute is already defined and only one instance of the attribute is allowed, the attribute is not displayed in the drop-down list.

3 In the Value field, type a value for the attribute (or in some cases, select an option), and click **OK**. The attribute is added to the service.

4 Save the configuration.

NOTE: When you add or change a static attribute, you must restart the SequeLink service before the change takes affect. Dynamic attributes become effective after the attribute is added or changed and the configuration is saved.

See Appendix D "SequeLink® Service Attributes" on page 491 for a description of the service attributes.

Deleting a Service Attribute

1 Right-click the SequeLink service attribute you want to delete, and select **Delete**.

2 You are prompted to confirm the deletion. Click **OK** to confirm. The attribute is deleted from the service configuration.

3 Save the configuration.

NOTE: When you delete a static attribute, you must restart the SequeLink service before the change takes affect. Deletion of dynamic attributes becomes effective after the attribute is deleted and the configuration is saved.

See Appendix D "SequeLink® Service Attributes" on page 491 for a description of the service attributes.

Configuring Server Data Sources Using the SequeLink® Manager Snap-in

This section describes how to create and manage server data sources with the SequeLink Manager MMC Snap-in. See "Data Sources" on page 34 for more information about data sources.

To do this...	See...
Create a server data source	"Creating a Server Data Source" on page 76
Delete a server data source	"Deleting a Server Data Source" on page 77
Rename a server data source	"Renaming a Server Data Source" on page 77
View server data source attributes	"Viewing Server Data Source Attributes" on page 77
Add a server data source attribute	"Adding a Server Data Source Attribute" on page 78
Change the default value of a server data source attribute	"Changing the Value of a Server Data Source Attribute" on page 79

Creating a Server Data Source

1 Right-click the **Data Source Settings** node and select **New / Data source**. A new server data source appears in the Details pane.

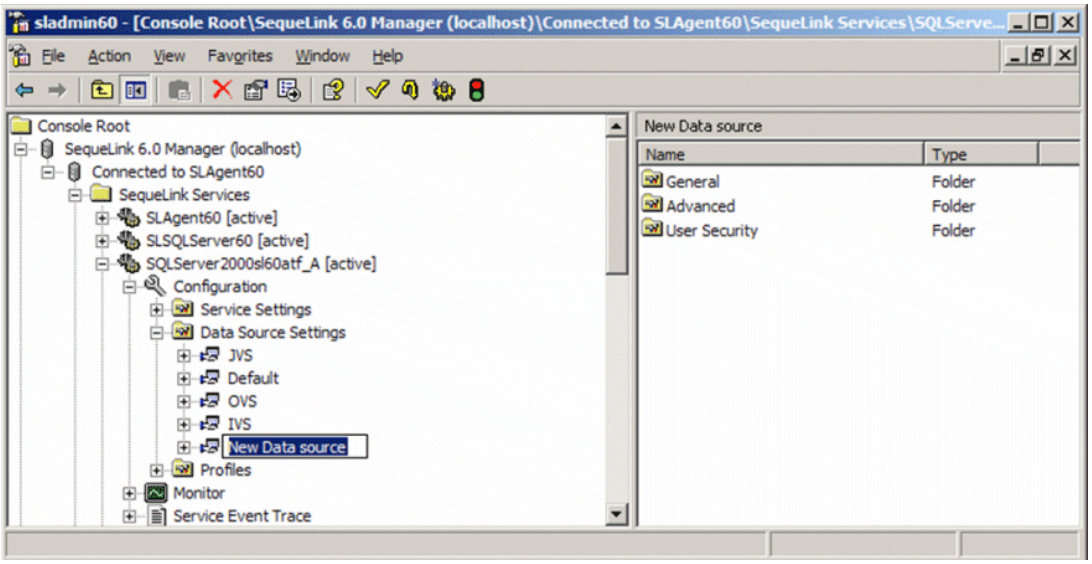

2 When created, the new data source is an editable field. Type the name of the data source and press ENTER.

3 Save the configuration.

When you create a server data source, the attributes for the new server data source are copied from the default data source. This default data source is created when the data access service is created. See "Viewing Server Data Source Attributes" on page 77 for instructions on viewing server data source attributes. See "Changing the Value of a Server Data Source Attribute" on page 79 for instructions on changing server data source attributes.

Deleting a Server Data Source

1 Right-click the server data source you want to delete, and select **Delete**.

2 You are asked to confirm the deletion. To confirm, click **OK**. The server data source is deleted.

NOTE: You cannot delete the Default data source.

3 Save the configuration.

Renaming a Server Data Source

1 Right-click the server data source you want to rename, and select **Rename**. The data source becomes an editable field. Type the name of the data source, and press ENTER.

2 Save the configuration.

Viewing Server Data Source Attributes

To view attributes for a data source, select the server data source. The available attribute categories appear in the Details pane. To view the attributes, select any attribute category. The attributes and their values appear in the Details pane. See Appendix D "SequeLink® Service Attributes" on page 491 for a description of the service attributes.

Adding a Server Data Source Attribute

1 Right-click the data source to which you want to add an attribute, and select **New / Attribute**. The New Attribute window appears.

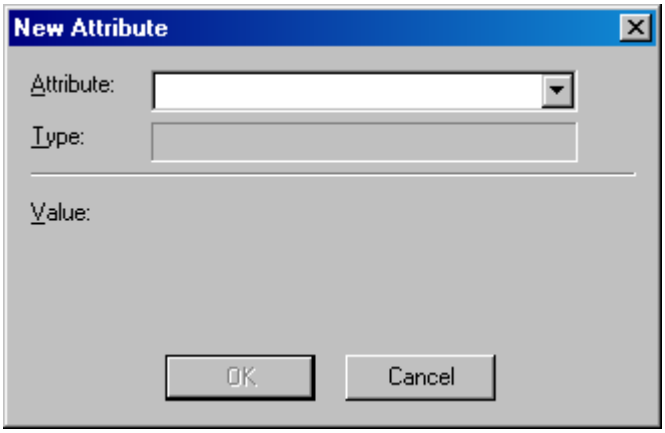

2 In the Attribute drop-down list, select the attribute you want to add to the server data source. The Type field adjusts to show the type of value required.

3 In the Value field, change the default value of the attribute if necessary; then, click **OK**. The attribute is added to the server data source.

4 Save the configuration.

See Appendix D "SequeLink® Service Attributes" on page 491 for a description of the attributes that apply to data sources.

Changing the Value of a Server Data Source Attribute

When you create a server data source, the attributes of the new data source are copied from the default data source.

To change the value of a server data source attribute:

1. Right-click the data source attribute, and select **Properties**. The Properties window for that attribute appears.
2. Type a new value for the attribute in the Value field, and click **OK**. The attribute is changed.
3. Save the configuration.

See Appendix D "SequeLink® Service Attributes" on page 491 for a description of the attributes that apply to data sources.

Configuring Monitoring

SequeLink provides the following levels of monitoring for SequeLink data access services, listed here from highest-level to lowest-level. Some of the monitoring levels are also supported for the SequeLink Agent service.

- **Service monitoring** monitors these activities by service:
  - Statistics of received packets and sent packets
  - Sessions started and statements opened
  - Active statements and sessions
  - Fetched rows and affected rows
  - Transactions

- **Session monitoring** monitors these activities by session within a service:
  - Statistics of received packets and sent packets
  - Statements opened and active statements
  - Fetched rows and affected rows
  - Transactions
  - Information about each session, such as start time, client information (network address, data source used by the client, and type of client), native database session identification, and database user
- **Statement monitoring** monitors these activities by statement within a session:
  - Fetched rows and affected rows
  - SQL statements issued

To enable monitoring at one of the listed levels, higher-level monitoring must be enabled. For example, you cannot monitor Session information unless Service monitoring is enabled. Similarly, you cannot monitor Statement information unless both Service monitoring and Session monitoring are enabled.

UNIX

On Linux/UNIX, both a monitoring and an event trace profile are enabled when you install SequeLink Server.

After installation, you can create a monitoring profile or modify an existing profile.

NOTES:

- You can configure only one monitoring profile for each SequeLink service.
- After changing the connection model for a service, you must delete any monitoring or event profiles for the service. For example, if you changed the connection model from Threadpool to Threadconnection, you must delete the monitoring and event profiles, and create new profiles.

You can integrate SequeLink monitoring with the Windows Performance Monitor tool, which allows you to access monitoring information from the Windows tool.

Creating a Monitoring Profile

1 Select the **Profiles** node for a SequeLink service (beneath the Configuration node). Select **Action / New / Monitoring Profile**. The Monitoring Properties window appears.

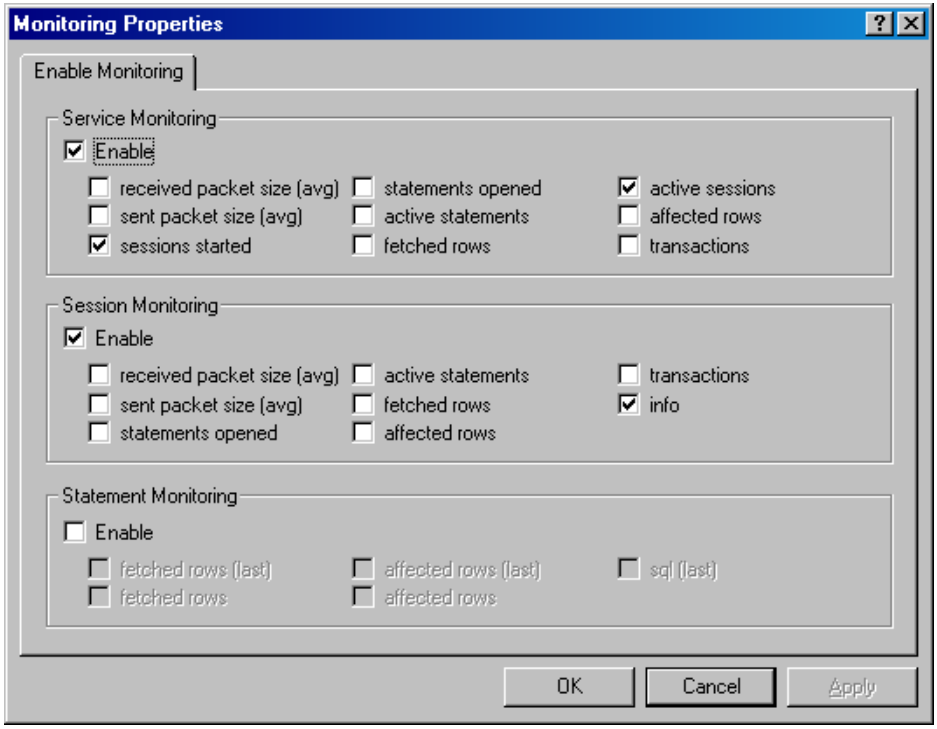

For Windows users: When you create a profile, the Enable Performance Monitoring window appears before the Monitoring window if you have the Windows Performance Monitor tool installed on the same Windows server on which the SequeLink Server and SequeLink Manager Snap-In are installed.

If you want to integrate SequeLink monitoring with the Windows Performance Monitor tool, select the **Enable** check box. Then, click **Next**. The Monitoring Properties window appears.

See "Integrating SequeLink® Monitoring with the Windows Performance Tool" on page 85 for instructions on integrating SequeLink with the Windows Performance Tool.

2 Perform any of the following actions in the Monitoring window:

- To enable a type of monitoring (Service Monitoring, Session Monitoring, and Statement Monitoring), select the **Enable** check box within the appropriate Monitoring group. Remember that to enable monitoring information at a monitoring level, you must enable the higher-level monitoring. Therefore, if you disable Service Monitoring, all lower-level monitoring information is also disabled.
- To enable a property, select the check box beside the property.
- To disable types of monitoring (Service Monitoring, Session Monitoring, and Statement Monitoring), clear the **Enable** check box within the appropriate Monitoring group.
- To disable any property, clear the check box beside the property.

3 Click **OK**.

4 Restart the SequeLink service to activate the new monitoring profile.

Changing an Existing Monitoring Profile

1 Select the **Profiles** node for a SequeLink service (beneath the service's Configuration node). The monitoring profile is displayed in the Details pane. Double-click the monitoring profile. The Monitoring Properties window appears.

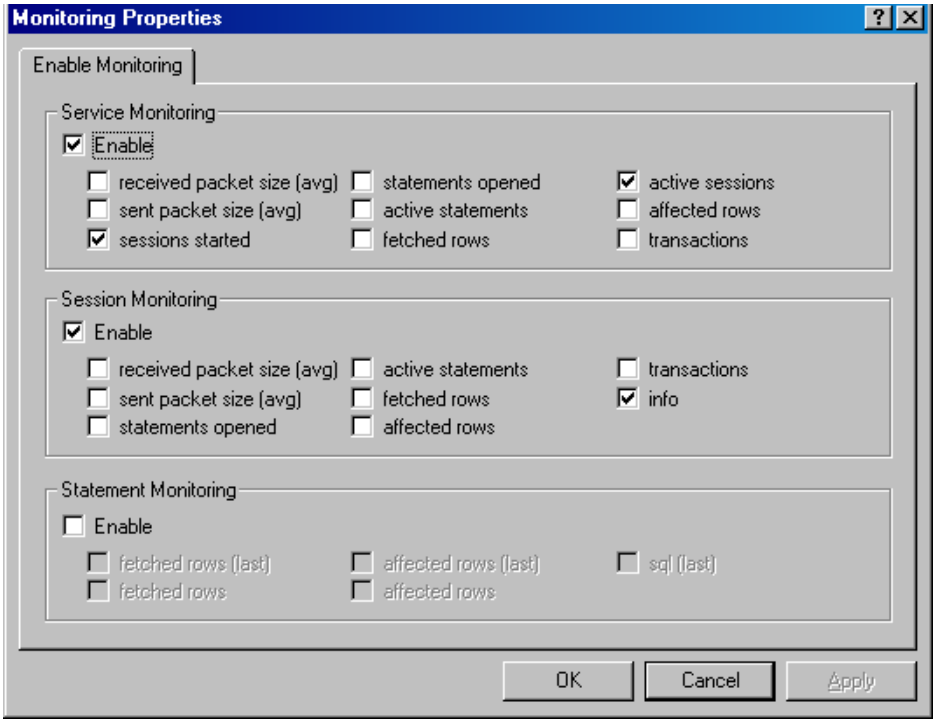

2 Perform any of the following actions in the Monitoring window:

- To enable a type of monitoring (Service Monitoring, Session Monitoring, and Statement Monitoring), select the **Enable** check box within the appropriate Monitoring group. Remember that to enable monitoring information at a monitoring level, you must enable the higher-level monitoring. Therefore, if you disable Service Monitoring, all lower-level monitoring information is also disabled.
- To enable a property, select the check box beside the property.
- To disable types of monitoring (Service Monitoring, Session Monitoring, and Statement Monitoring), clear the **Enable** check box within the appropriate Monitoring group.
- To disable a property, clear the check box beside the property.

3 Click **OK**.

4 Restart the SequeLink service to activate the changed monitoring profile.

Deleting a Monitoring Profile

1 Select the **Profiles** node for a SequeLink service (beneath the service's Configuration node). The existing monitoring profile is displayed in the Details pane.

2 Right-click the monitoring profile, and select **Delete**.

3 You are prompted to confirm the deletion. Click **OK** to confirm. The profile is deleted from the service configuration.

4 Save the configuration.

Integrating SequeLink® Monitoring with the Windows Performance Tool

If you are integrating SequeLink monitoring with the Windows Performance Monitoring tool on the Windows platforms on which the SequeLink Server runs, you must explicitly set the required counters in the Windows registry. When the SequeLink Server installation finishes, the files SWEVPERF.INI and SWEVPERF.H appear in your temporary directory (for example, C:\temp).

To integrate SequeLink monitoring with Windows performance monitoring:

1. Using a command prompt, change the directory to the directory containing the SWEVPERF.INI file (for example, C:\temp).

2. Type the following command; then, press ENTER:

   ```
   lodctr swevperf.ini
   ```

3. If you do not have an active (enabled) monitoring profile that is configured for integration with the Windows Performance Monitor tool, configure one. See "Configuring Monitoring" on page 79 for instructions on configuring monitoring. To activate a new monitoring profile, restart the SequeLink service for which you defined the profile.

   NOTE: If you have a monitoring profile that is not configured for integration with the Windows Performance Monitor Tool, delete that monitoring profile and create one that is configured for integration. Integration must be configured when you create the profile.

4. Start the Windows Performance Monitor tool and select the **Add to Chart** menu item. In the window that appears, select the **SequeLink 6.0 Services** object and select the counters you want to monitor from the Instance list box.

Configuring Event Tracing

Events are generated when the client application accesses data and when specific server activities occur, such as when a service starts or an error occurs. Depending on which SequeLink profiles are active, the information generated by the event is displayed as it occurs in the runtime monitor and is stored persistent in the event trace file.

By default, the event trace file is located in the *installdir*\tracing directory where *installdir* is your SequeLink Server installation directory.

UNIX On Linux and UNIX, both a monitoring and an event trace profile are enabled when you install SequeLink Server.

NOTE: You can configure only one event trace profile for each SequeLink service.

Creating an Event Trace Profile

1 Select the **Profiles** node for a SequeLink service (beneath the service's Configuration node). Select **Action / New / Event Trace Profile**. The Event Tracing Properties window appears.

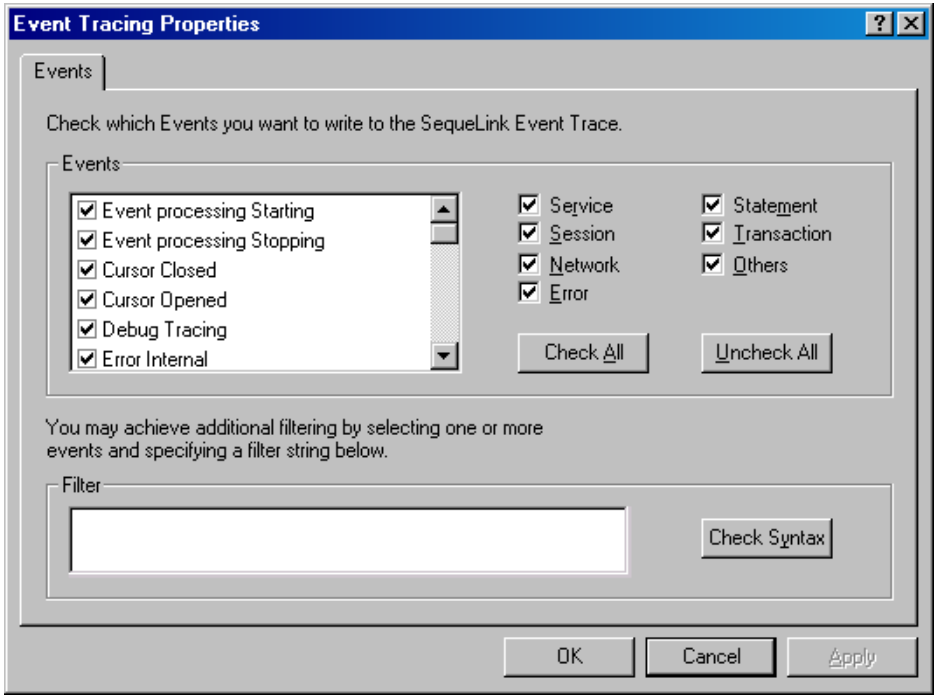

2 Enable and disable events:

- To enable a group of events, select the check box beside the appropriate group name (Service, Session, Network, Error, Statement, Transaction, or Others).
- To enable individual events, select the check box beside the event in the scroll box. See Appendix E "SequeLink® Events" on page 589 for a list and a description of all events.
- To disable a group of events, clear the check box beside the appropriate group name (Service, Session, Network, Error, Statement, Transaction, or Others).

NOTE: Event names that do not start with Service, Session, Network, Error, Statement, or Transaction are Other events (for example, Cursor Closed).

- To disable individual events, clear the check box beside the event in the scroll box.

3. Optionally, you can place a filter on any event or group of events. To add a filter, type the filter in the Filter text box. Some events have a set of attributes. You can place a filter on the attributes of an event. For example, if you want to monitor and trace only authentication events for sessions started by administrators, you could write the following filter for the Session Authenticated event:

```
${Authorization} = "administrator"
```

See Appendix E "SequeLink® Events" on page 589 for an explanation of the filter syntax, and a list and description of the attributes for each event.

4. Click **Next**. The Add Event Trace Profile window appears. Click **Finish** in this window to add the profile.

5. Save the configuration.

6. Restart the SequeLink service to activate the new event trace profile.

Changing an Existing Event Trace Profile

1 Select the **Profiles** node for a SequeLink service (beneath the Configuration node). The existing event trace profile is displayed in the Details pane. Double-click the event trace profile. The Events window appears.

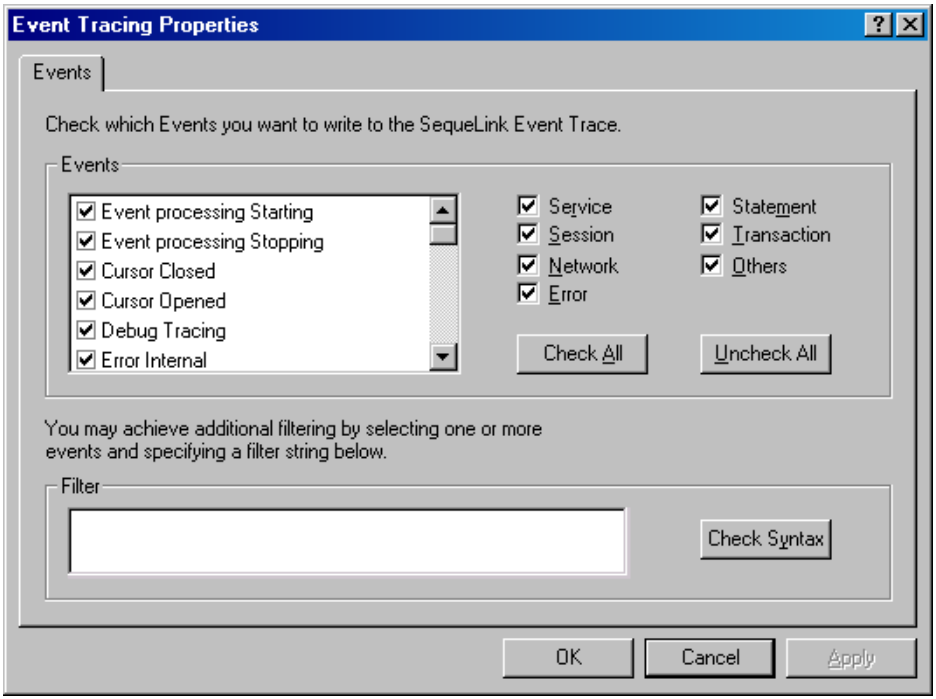

2 Enable and disable events:

- To enable a group of events, select the check box beside the appropriate group name (Service, Session, Network, Error, Statement, Transaction, or Others).
- To enable individual events, select the check box beside the event. See Appendix E "SequeLink® Events" on page 589 for a list of all events and their definition.

- To disable a group of events, clear the check box beside the appropriate group name (Service, Session, Network, Error, Statement, Transaction, or Others).
- To disable individual events, clear the check box beside the event in the scroll box.

3. Optionally, you can place a filter on an event or group of events. Type a filter in the Filter text box. Some events have a set of attributes. You can place a filter on the attributes of an event. For example, if you want to monitor and trace only authentication events for sessions started by administrators, you could write the following filter for the Session Authenticated event:

   ```
   ${Authorization} = "administrator"
   ```

   See Appendix E "SequeLink® Events" on page 589 for an explanation of the filter syntax, and a list and description of the attributes for each event.

4. Click **OK**.
5. Save the configuration.
6. Restart the SequeLink service to activate the new event trace profile.

Deleting an Event Trace Profile

1. Select the **Profiles** node for a SequeLink service (beneath the service's Configuration node). The existing event trace profile is displayed in the Details pane.
2. Right-click the event profile, and select **Delete**.
3. You are prompted to confirm the deletion. Click **OK** to confirm. The event profile is deleted from the service configuration.
4. Save the configuration.

4 Managing Data Access Activity Using the SequeLink® Manager Snap-in

This chapter describes the tasks you perform to manage and monitor data access activity using the SequeLink Manager MMC Snap-in.

To do this...	See...
Kill a session	"Killing a Session" on page 92
View event tracing information	"Viewing Event Tracing Information" on page 92
View details about an active service	"Viewing Details About an Active Service" on page 96
View active sessions and details about an active session	"Viewing Active Sessions and Details About an Active Session" on page 97

Killing a Session

NOTE: To view or kill sessions for a SequeLink service, monitoring must be enabled for the service. See "Configuring Monitoring" on page 79 for information about enabling monitoring for a SequeLink service.

To kill a session:

1. Click the **Monitor / Active sessions** node of the SequeLink service. The Details pane shows a list of active sessions.
2. Right-click the session you want to kill, and select **Kill**.

Viewing Event Tracing Information

You can view all events of a service or filter the events to view. You can also use the Find function to search for a specific event in the events that are displayed in the Details pane.

NOTE: When the service event trace is refreshed, it is possible that the displayed information may no longer be available in the event trace file because the event trace file is circular, meaning that the oldest events are overwritten with the newest events. If this happens, the SequeLink Manager displays `Information not available`.

See "Event Handling" on page 39 for general information about event handling.

To view event trace information:

1 Click the **Service Event Trace** node of a service. The Details pane displays a list of all events, which depends on how event tracing is configured.

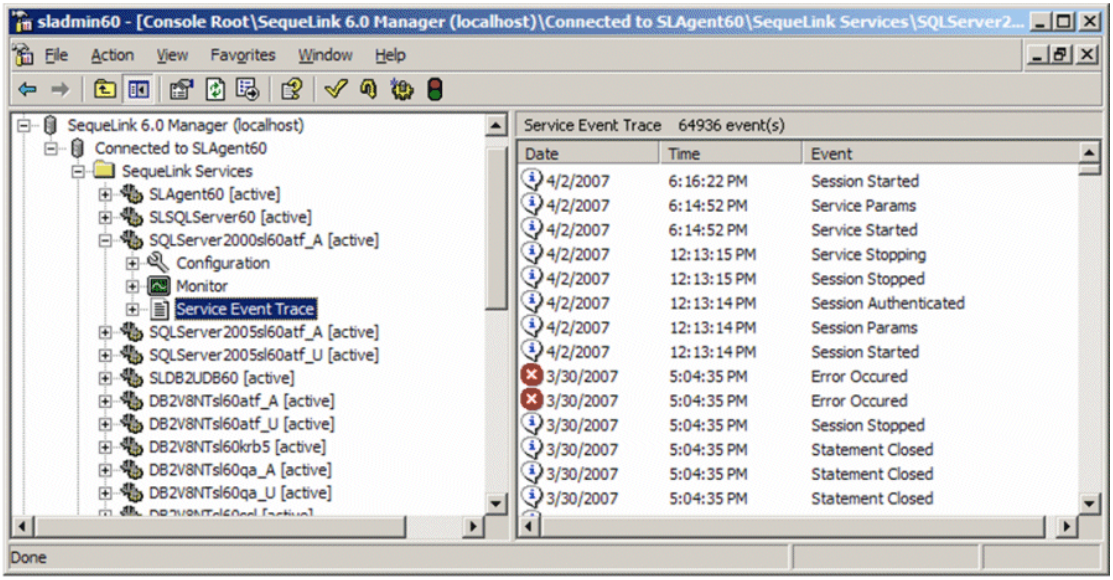

2 Optionally, you can search for a specific event in the Details pane. To do this, right-click the **Service Event Trace** node of a service; then, select **Find Event**. The Find Event window appears.

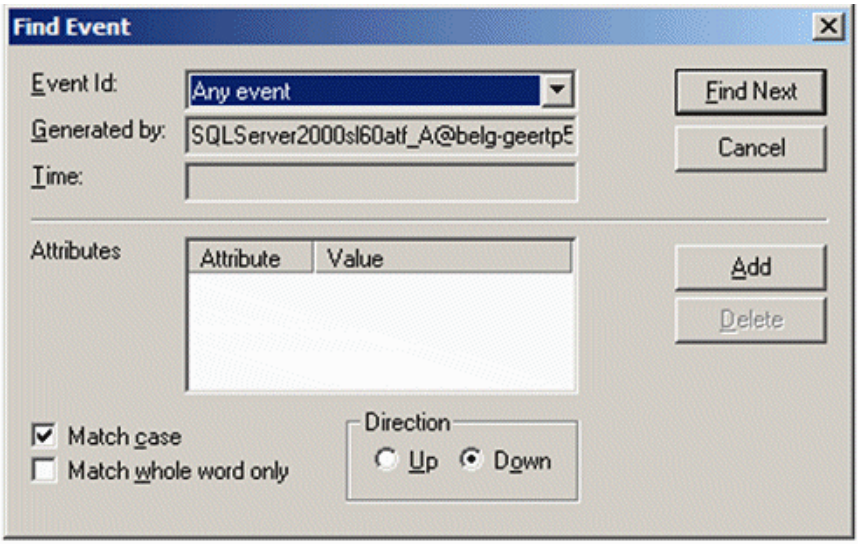

In this window, perform the following actions:

a Select the type of event you want to view from the Event ID drop-down list.

b To search down in the event list, select **Down** in the Direction group. To search up, select **Up**.

c Optionally, you can define search criteria for the event. Click **Add**. The New Attribute window appears. In this window, select the attribute by which you want to search from the Attribute drop-down list. In the Value field, type the search criteria for the attribute. For example, if you want to search for Session Started events from a certain client host (such as 'sales1.company.com'), you would select **Session Started** in the Find Event window, click **Add**, select **ClientInfo** from the Attribute drop-down list, and type `sales1.company.com` in the Value field.

d Click **Find Next** to find the next event.

3 Optionally, you can filter the search to display only specific events in the Details pane. To do this, click the **Service Event Trace** node of a service; then, select **View / Filter**. The Filter window appears.

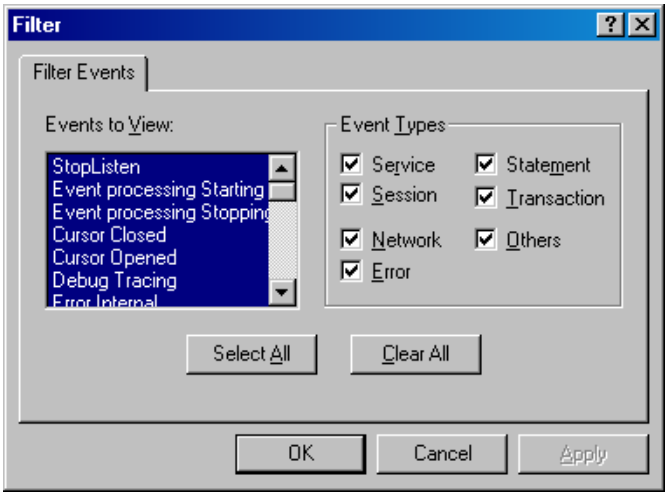

In this window, perform one of the following actions:

- Select the events you want to view from the Events to View list. To select non-adjacent events from the list, hold down the CTRL key while selecting the events. To select adjacent events from the list, hold down the SHIFT key while selecting the events.
- Select the types of events you want to view by selecting the appropriate Event Types check boxes. When you select an event type, the corresponding events are highlighted in the Events to View list.

4 To view information about a specific event, right-click an event in the Details pane, and select **Properties**. The Event Properties window appears.

5 Click the **Attributes** tab to view the attributes for that event.

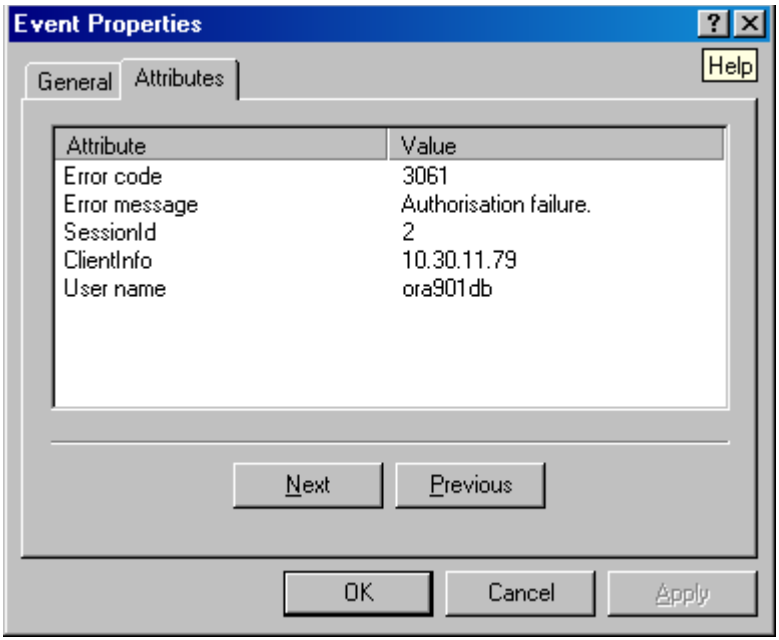

Viewing Details About an Active Service

Select the **Monitor** node of the service for which you want details. The Details pane shows details about the active service. The level of detail that is displayed depends on how the monitoring is configured. See "Configuring Monitoring" on page 79 for more information.

See "Refreshing Active Information" on page 59 for information about refreshing information about active services.

Viewing Active Sessions and Details About an Active Session

1. Select the **Monitor** node of the SequeLink service for which you want to view active sessions.
2. Select the **Active sessions** node. The Details pane shows all active sessions for that SequeLink service.

   NOTE: Optionally, you can sort active sessions that appear in the Details pane based on session name, session type, IP address, or DBMS user. To do this, click the appropriate header category at the top of the Details pane. The active session list is sorted based on the criteria you selected.
3. To view details about an active session, expand the **Active sessions** node. In the console tree in the left pane, select the session for which you want details. The Details pane displays details about that session.

See "Refreshing Active Information" on page 59 for information about refreshing information about active sessions.

5 Using the SequeLink® Manager Command-Line Administrator

This chapter describes how to use the SequeLink Manager Command-Line Administrator, issue SequeLink Manager commands, and lists some commonly used SequeLink Manager commands. See "SequeLink® Server System Administration" on page 29 for general information about the SequeLink Manager.

Using the SequeLink® Manager Command-Line Tool

UNIX

The SequeLink Manager Command-Line Tool runs on Linux, UNIX, and Windows only. It allows you to configure and manage your SequeLink environment remotely from a networked client or locally from a SequeLink Server by issuing commands through the command-line interface.

You can use any of the following methods to issue SequeLink Manager commands:

- Direct
- Batch
- Interactive

z/OS

NOTE: To configure and manage SequeLink services on z/OS or to create z/OS-specific core entities such as DB2 interfaces, use the SequeLink Manager for z/OS. See Chapter 7 "Configuring SequeLink® Services Using the SequeLink® Manager for z/OS" on

page 129 for more information. Monitoring can be performed using any SequeLink Manager.

Direct Command Execution

Using direct command execution, you can invoke one command at a time on the command line. The command syntax for direct execution is:

`executable command_name [parameters]`

where:

`executable` is *installdir*\admin\swcla.exe on Windows and *installdir*/admin/swcla.sh on Linux/UNIX. See "Invoking the SequeLink® Manager Command-Line Tool" on page 103 for the parameters you can specify when invoking the tool.

`command_name` is any SequeLink Manager command (short or long name). See Appendix B "SequeLink® Manager Commands" on page 439 for a list and description of each command.

`parameters` are valid options for the specified command. When using direct command execution to issue a command, you must provide all of a command's required and optional parameters in the correct position (specify the parameters in the order they are documented).

Batch Command Execution

The batch method of executing commands allows you to execute a script file that contains SequeLink Manager commands. The command syntax for batch execution is:

```
executable -i script_file
```

where:

`executable` is *installdir*\admin\swcla.exe on Windows and *installdir*/admin/swcla.sh on Linux/UNIX. See "Invoking the SequeLink® Manager Command-Line Tool" on page 103 for the parameters you can specify when invoking the tool.

`script_file` is the name of the file to execute. This file can contain SequeLink Manager commands only. When using a script file to execute commands, you must provide all of a command's required and optional parameters in the script file (specify the parameters in the order they are documented).

Interactive Command Execution

Using interactive command execution, you do not have to specify all the command parameters on the command line. The command-line tool will prompt you for the required parameters. To start interactive mode, type:

```
executable
```

where:

`executable` is *installdir*\admin\swcla.exe on Windows and *installdir*/admin/swcla.sh on Linux/UNIX. See "Invoking the SequeLink® Manager Command-Line Tool" on page 103 for the parameters you can specify when invoking the tool.

The tool is invoked and a command prompt is displayed, along with copyright text (unless you invoked the tool with the -nologo flag):

```
swcla>
```

At the swcla command prompt, type either a command with all of its parameters or a command name, and the tool will prompt you for the required parameters.

To exit interactive mode, type:

```
exit
```

General Rules

- SequeLink Manager command names are not case-sensitive; however, the command parameter *service_name* is case sensitive.
- If the value of a command parameter contains spaces, the value must be enclosed within single quotes (') or double quotes (").
- If the value of a command parameter contains single or double quotes, use single quotes to quote double quotes and double quotes to quote single quotes.
- The pound sign (#) is a comment character. All text that follows the pound sign on the same line is ignored.
- When issuing commands using the direct or batch method, you must specify all of a command's required and optional parameters in the correct position (specify the parameters in the order they are documented).

Invoking the SequeLink® Manager Command-Line Tool

You invoke the SequeLink Manager Command-Line Tool using the following syntax:

```
executable [-nologo] [-i script_file] [-o output_file] [-e error_file]
[-l | -r host:port] [-uid user_id [-pwd password]] [command]
```

where:

`executable` is *installdir*\admin\swcla.exe on Windows and *installdir*/admin/swcla.sh on Linux/UNIX.

`-nologo` disables the display of the copyright banner.

`-i script_file` specifies a script file to execute.

`-o output_file` specifies the name of a file in which the command-line tool will write all of its output (except errors).

`-e error_file` specifies the name of a file in which the command-line tool will write all its errors (not normal output). All error and normal output can be written to the same file if the file specified for the `-o` and `-e` flags is the same.

`-l` specifies to use the local host configuration.

`-r host:port` specifies to use the remote host configuration. In this case, you must specify the host name and port of the remote server.

`-uid user_id` specifies the user ID to use for local or remote host configuration. This parameter is valid only in combination with the `-l` or `-r` parameters.

`-pwd password` specifies a password to use for local or remote host configuration. This parameter is valid only in combination with the `-l` or `-r` parameters.

command is the name of any SequeLink Manager command and its parameters. See Appendix B "SequeLink® Manager Commands" on page 439 for a list and description of each command.

When invoking the SequeLink Manager Command-Line Tool, you must choose whether you want the tool to connect to a local SequeLink Agent (a SequeLink Agent service running on the same machine as the tool) or a remote SequeLink Agent (a SequeLink Agent service running on another machine). If you have not connected to a SequeLink Agent, the SequeLink Manager Command-Line Tool will return the following message when you issue a command:

```
[SequeLink error 5509] Command not available for current
configuration.
```

You can invoke the SequeLink Manager Command-Line Tool and connect to a local or remote SequeLink Agent after you invoke the tool or when you invoke the tool as shown in the following examples:

Example A: Connecting to a Local or Remote SequeLink Agent After Invoking the Tool

```
SLServer60\admin\swcla.exe
```

This example invokes the tool without specifying whether it will connect to a local or remote SequeLink Agent. You must now explicitly connect to a local or remote SequeLink Agent using the ActivateLocalConfig or ActivateRemoteConfig commands described in Appendix B "SequeLink® Manager Commands" on page 439 before you can enter any other commands.

Example B: Connecting to a Local SequeLink Agent When Invoking the Tool

```
SLServer60\admin\swcla.exe -l
```

This example connects the tool to a local SequeLink Agent. The administrator will be prompted for the user ID and password for the SequeLink administrator. The administrator can now enter any SequeLink Manager command described in Appendix B "SequeLink® Manager Commands" on page 439.

Example C: Connecting to a Remote SequeLink Agent When Invoking the Tool

```
SLServer60\admin\swcla.exe -r
```

This example connects the tool to a remote SequeLink Agent. The administrator will be prompted for the host name and port of the remote server. Then, the administrator will be prompted for the user ID and password for the SequeLink administrator. The administrator can now enter any SequeLink Manager command described in Appendix B "SequeLink® Manager Commands" on page 439.

Displaying Help for a Command

To display help for a command, type the following at a command-line tool prompt:

```
help [long_command_name | short_command_name]
```

For example:

```
help DataSourceCreate
```

You can also display help about the different options you have when invoking the command-line tool. To do this, type the following at a command-line tool prompt:

```
help
```

Commonly Used SequeLink® Manager Commands

This section lists some commonly used SequeLink Manager commands. Both long and short command names are listed. See Appendix B "SequeLink® Manager Commands" on page 439 for a complete list of SequeLink Manager commands.

Starting a SequeLink® Service

Command: ServiceStart | ss

Syntax: `{ServiceStart | ss}` *`service_name`*

Example: `ss SLOracle10`

Stopping a SequeLink® Service

Command: ServiceStop | sst

Syntax: `{ServiceStop | sst}` *`service_name`*

Example: `sst SLOracle10`

Creating a SequeLink® Service

Command: ServiceCreate | sc

Syntax: `{ServiceCreate | sc}` *`service_name service_ID tcp_port`*

Service template IDs can be obtained using the ServiceTemplateList|stl command.

Example: `sc SLOracle10 SL6_Oracle10W 19996`

Deleting a SequeLink® Service

NOTE: Before deleting a SequeLink service, you must stop the service you want to delete.

Command: ServiceDelete | sd

Syntax: `{ServiceDelete | sd}` *`service_name`*

Example: `sd SLOracle10`

Viewing Service Attributes

Command: ServiceInfo | si

Syntax: `{ServiceInfo | si}` *`service_name`*

Example: `si SLOracle10`

Changing a Service Attribute

Command: ServiceAttributeReplace | sar

Syntax: `{ServiceAttributeReplace | sar}` *`service_name attribute_name value`*

Example: `sar SLOracle10 ServiceUser[2] devuser`

Adding a Service Attribute

Command: ServiceAttributeAdd | saa

Syntax: `{ServiceAttributeAdd | saa}` *`service_name attribute_name value`*

Example: `saa SLOracle10 ServiceUser sqlnk`

Deleting a Service Attribute

Command: ServiceAttributeDelete | sad

Syntax: `{ServiceAttributeDelete | sad}` *`service_name attribute_name`*

Example: `sad SLOracle10 ServiceCodePage`

`sad SLOracle10 ServiceUser[2]`

Creating a Server Data Source

Command: DataSourceCreate | dsc

Syntax: `{DataSourceCreate | dsc}` *`service_name data_source_name`*

Example: `dsc SLOracle10 DS_Employees`

Deleting a Server Data Source

Command: DataSourceDelete | dsd

Syntax: `{DataSourceDelete | dsd}` *`service_name data_source_name`*

Example: `dsd SLOracle10 DS_Employees`

Changing a Server Data Source Attribute

Command: DataSourceAttributeReplace | dsar

Syntax: `{DataSourceAttributeReplace | dsar}` *`service_name data_source_name attribute_name value`*

Example: `dsar SLOracle10 DS_Employees DataSourceCurrentCatalog partners`

Adding a Server Data Source Attribute

Command: DataSourceAttributeAdd | dsaa

Syntax: `{DataSourceAttributeAdd | dsaa}` *`service_name data_source_name attribute_name value`*

Example: `dsaa SLOracle10 DS_Employees DataSourceCurrentCatalog employees`

Deleting a Server Data Source Attribute

Command: DataSourceAttributeDelete | dsad

Syntax:

```
{DataSourceAttributeDelete | dsad} service_name
data_source_name attribute_name
```

Example:

```
dsad SLOracle10 DS_Employees DataSourceCurrentCatalog
```

Killing a Session

Command: SessionStop | sess

Syntax:

```
{SessionStop | sess} service_name session_ID
```

Session IDs can be obtained with the ServiceActiveInfo | sai command.

Example:

```
sess SLOracle10 5
```

Viewing Event Tracing Information

Command: EventList | el

Syntax:

```
{eventlist | el} service_name | [remote]file=
event_trace_file_name
[details]
[ [{service | srvc}] |
[{session | sess}] |
[{statement | stmt}] |
[{transaction | trans}] |
[{network | net}] |
[{error | err}] |
[{other | oth}] ] ]
[count=[{ + | - }] {all | number}]
```

```
[offset={begin | end} [{ + | - }]number]
 [query='custom_event_filter_string']
```

Example: **Local host or remote configuration examples:**

```
el SLAgent details

el SLOracle details service count=all offset=10

el SLOracle10 stmt query='${ReturnCode} != 0'

el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLoracle10.trc" count=10

el "remotefile=C:\Program Files\DataDirect\slserver60\
tracing\SLoracle10.trc" service details
```

Offline configuration examples:

```
el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLoracle10.trc"

el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLoracle10.trc" count=10

el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLoracle10.trc" count=-all offset=end

el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLoracle10.trc" offset=5 service session

el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLoracle10.trc" service details
```

Viewing Active Services

Command: ServiceList | sl

Syntax: `{ServiceList | sl}`

Example: `sl`

Viewing Details About an Active Service (Including Active Sessions)

Command: ServiceActiveInfo | sai

Syntax: `{ServiceActiveInfo | sai}` *`service_name`*

Example: `sai SLOracle10`

Viewing Details About an Active Session

Command: SessionInfo | sesi

Syntax: `{SessionInfo | sesi}` *`service_name session_ID`*

Session IDs can be obtained using the ServiceActiveInfo | sai command.

Example: `sesi SLOracle10 5`

6 Using the SequeLink® Manager for z/OS

z/OS You can use the SequeLink Manager for z/OS to configure and manage SequeLink Server for z/OS locally using an ISPF dialog. This chapter describes how to use the SequeLink Manager and the SequeLink Manager Operator Interface.

Starting the SequeLink® Manager for z/OS

1 Start the SequeLink Manager for z/OS. How you start the SequeLink Manager for z/OS depends on whether you allocated the SequeLink ISPF libraries to the TSO session (refer to the *SequeLink Installation Guide*).

- If you allocated the ISPF libraries, continue with Step 2.
- If you did not allocate the ISPF libraries, continue with Step 3.

2 **Starting the SequeLink Manager When ISPF Libraries Are Allocated:**

Type the following command:

```
TSO %SSMC [HLQ(HLQ_prefix) USR(USR_prefix)]
```

where *HLQ_prefix* is the installation dataset prefix, and *USR_prefix* is the server-specific dataset prefix. Then, press ENTER. When using this format to start the SequeLink Manager for z/OS, you must always pass the HLQ parameter. We recommend passing the USR parameter as well. If you do not specify the USR parameter, SequeLink retrieves the last

passed value from your SSMC profile. If neither parameter is passed, and no value can be retrieved from your profile, the default value HLQ is used.

For example:

```
TSO %SSMC HLQ(SQLNK.V6R0M0) USR(COMPANY.DEV)
```

3 **Starting the SequeLink Manager When ISPF Libraries Are Not Allocated:**

Type the following command:

```
TSO EX 'HLQ_prefix.CLIST(SSMC)' '[HLQ(HLQ_prefix)]
[USR(USR_prefix)]'
```

where *HLQ_prefix* is the installation dataset prefix, and *USR_prefix* is the server-specific dataset prefix. Then, press ENTER. You do not have to explicitly specify the HLQ prefix. If it is not specified, the prefix is set to the high-level qualifier of the dataset from which the SSMC CLIST is executed. We recommend passing the USR parameter as well; the parameter will be saved in the user's SSMC profile. If you do not specify the USR parameter, SequeLink retrieves the last passed value from your SSMC profile. If neither parameter is passed, and no value can be retrieved from your profile, the default value HLQ is used.

For example:

```
TSO EX 'SQLNK.V6R0M0.CLIST(SSMC)' 'HLQ(SQLNK.V6R0M0)
USR(COMPANY.DEV)'
```

or

```
TSO EX 'SQLNK.V6R0M0.CLIST(SSMC)' 'USR(COMPANY.DEV)'
```

The HLQ used will be the HLQ from which the SSMC CLIST is fetched.

NOTE: In the first example, the HLQ specified to locate the SSMC CLIST might be different from the HLQ argument passed. This is especially important when administering two different versions of SequeLink. In this case, we recommend that you always specify the HLQ argument explicitly.

4 When the Serverlist dataset (USR_prefix.SERVER) has not been allocated for the environment USR, you are prompted for the necessary information. Press ENTER to continue. The SequeLink Manager for z/OS main menu appears.

```
       DataDirect SequeLink Manager for z/OS - Main menu

 Command ===>

ServerList dataset COMPANY.DEV.SERVERS is not defined yet

Do you want it to be allocated now? YES/NO


    F1=Help      F3=End       F5=View Err F12=Cancel
```

5 If you select NO, you exit the SSMC application. If you select YES, the following panel appears. Enter `VOL(vvvvvv)` where *vvvvvv* is the volume on which you want to allocate the serverlist dataset. This volume will be used to allocate server-specific datasets for SequeLink servers that are added to this environment. Press ENTER to continue.

```
       DataDirect SequeLink Manager for z/OS - Main menu

 Command ===>

Specify VOLUME on which to allocate it?  VOL(vvvvvv) | VOL (SMS)


    F1=Help      F3=End       F5=View Err F12=Cancel
```

6 When you start the SequeLink Manager for z/OS, copyright information appears. Press ENTER to continue. The SequeLink Manager for z/OS ServerList panel appears.

```
DataDirect SequeLink Manager for z/OS - ServerList          Row 1 to 3 of 3



 Allowed actions or commands are:

 o  (S)elect a server      o  (G)enerate JCL
 o  (A)dd a server         o  (O)perator interface
 o  (D)elete a server

    Name     Description
--------------------------------------------------------------------------------

SQLDB2V7     Development SequeLink for DB2v7 (subsys: DB7U)

SQLDB2V8     Development SequeLink for DB2v8 (subsys: DB8A)

SQLDB2V9     Development SequeLink for DB2v9 (subsys: DB9N)

****************************** Bottom of data ********************************

 Command ===>                                                   Scroll > PAGE
```

Working with the SequeLink® Manager for z/OS

The SequeLink Manager for z/OS ISPF panels list items in your SequeLink configuration, allow you to select actions that affect these items, or allow you to view or change attributes.

For example, you can select a server from the SequeLink Manager for z/OS ServerList menu.

```
DataDirect SequeLink Manager for z/OS - ServerList          Row 1 to 3 of 3



 Allowed actions or commands are:

 o  (S)elect a server      o  (G)enerate JCL
 o  (A)dd a server         o  (O)perator interface
 o  (D)elete a server

    Name     Description
--------------------------------------------------------------------------------

SQLDB2V7     Development SequeLink for DB2v7 (subsys: DB7U)

SQLDB2V8     Development SequeLink for DB2v8 (subsys: DB8A)

SQLDB2V9     Development SequeLink for DB2v9 (subsys: DB9N)

****************************** Bottom of data ******************************
```

Using the Server Management Tree

When you select a SequeLink Server, the server management tree appears.

```
                 DataDirect SequeLink Manager for z/OS

 Command ===>                                                  Scroll > CSR

  Management Tree for server SQLDB2V8
  To see a list of valid actions on a node, type '?' beside it.

  Use '/' to expand or collapse tree branches.

  Valid commands are: EXP SAVE REFRESH ERRSTK
-------------------------------------------------------------------------------
SQLDB2V8 (offline)
     - Global Settings
     - DB2 Interface
         - DSN8
      - UID Maps
     - SequeLink Service
```

You can use the server management tree to view or change entities defined within the SequeLink Server, such as:

- Global settings
- DB2 interface
- User ID (UID) maps
- A SequeLink service and its configuration settings

Press the TAB key to navigate to the tree nodes on the server management tree.

To perform an action on a tree node, type the action code at the tree node. To find which actions each node accepts, type `?` beside the node you want information about and press ENTER. For example, when you type `?` beside a SequeLink service, a message appears telling you that the node will accept `D` (Delete), `S` (Select), `/` (Expand or Contract), `+` (Expand), and `-` (Contract).

Using the Command prompt, you can also type the following commands that affect the server management tree:

- `EXP` or `EXPAND` expands the tree so that all nodes become visible.
- `SAVE` saves the current configuration to disk.
- `REFRESH` rebuilds the tree (if the server is online, the monitor data is refreshed also).
- `ERRSTK` displays the error stack.

Using the Function Keys

When working with the SequeLink Manager for z/OS, you can use the function keys listed in Table 6-2.

Table 6-1. SequeLink Manager for z/OS Function Keys

Key	Description
F1	Panel/field help
F2	Split
F3	End with save of changes
F5	Display errors
F7	Scroll up
F8	Scroll down
F9	Swap
F10	Scroll left
F11	Scroll right
F12	Cancel without save of changes or clear validation messages

NOTE: You can use the standard ISPF command FKA ON to switch on the display of the function key labels at the bottom of each panel

Using the SequeLink® Manager for z/OS Operator Interface

You can use the Operator Interface of the SequeLink Manager for z/OS to issue commands to perform administration tasks. See Appendix C "Operator Interface Commands for z/OS" on page 479 for a list of available Operator Interface commands.

Operator Interface Requirements

This section lists the configuration requirements that must be met before using the Operator Interface:

RACF requirements:

- The TSOAUTH and OPERCMDS RACF classes must be activated. Use the command:

  ```
  SETROPTS CLASSACT(TSOAUTH OPERCMDS)
  ```

- The OPERCMDS class must be included in the RACLIST. Use the command:

  ```
  SETROPTS RACLIST(OPERCMDS)
  ```

- The CONSOLE resource must be defined in the TSOAUTH class. Use the command:

  ```
  RDEF TSOAUTH CONSOLE UACC(NONE)
  ```

- The OPERPARM segment of the RACF profile of each user of the Operator Interface must include either AUTH=SYS or AUTH=ALL. For example:

  ALTUSER *user-id* OPERPARM(AUTH(ALL))

  where *user-id* is the user ID of the user of the Operator Interface.

- Each user of the Operator Interface must have the following permissions granted:
  - Permission to use the TSO console command. Use the command:

    ```
    PERMIT CONSOLE CL(TSOAUTH) ID(user-id) ACCESS(READ)
    ```

    where *user-id* is the user ID of the user for who you are granting permissions.
  - Permission to issue operator commands from an MCS-console. Use the command:

    ```
    RDEF OPERCMDS MVS.MCSOPER.user-id UACC(NONE)
    PERMIT MVS.MCSOPER.user-id CL(OPERCMDS) ID(user-id)
    ACCESS(UPDATE)
    ```

    where *user-id* is the user ID of the user for whom you are granting permissions.

    NOTE: Update is required because the Operator Interface issues MODIFY commands.
- Refresh the RACLIST. Use the command:

  ```
  SETROPTS RACLIST(OPERCMDS) REFRESH
  ```

Other requirements:

Make sure that the TSO commands CONSOLE and CONSPROF run as APF-authorized commands. For more information, refer to the information about the IKJTSO*xx* parmlib member in your IBM documentation.

Starting the Operator Interface

You can start the Operator Interface from the ServerList panel. To start the Operator Interface, type `O` beside the SequeLink Server you want to perform operator commands with; then press ENTER. The SequeLink Operator Console panel appears, allowing you to type operator commands at the Operator command ===> prompt.

```
SequeLink Operator Console for ACCT1                        Row 1 to 1 of 1
                                                            Scroll > PAGE
 Enter a command to execute

 Operator command ===>


 Command output:
-------------------------------------------------------------------------------
 Enter a command or press 'F3' to exit...
 ****************************** Bottom of data ******************************
```

Generating JCL

After you create and configure a SequeLink Server for the first time, you must generate JCL to complete the configuration task. The SequeLink Manager customizes JCL obtained from the SequeLink_HLQ.SKELS data set and saves the customized output in a data set named USR_HLQ.*Servername*.CNTL. Both SequeLink_HLQ and USR_HLQ are passed to the SSMC clist when you start the SequeLink Manager for z/OS.

See "Data Sets Created During Installation" on page 126 for a description of SequeLink_HLQ and USR_HLQ.

Generating JCL allows you to easily keep the JCL members synchronized with the actual server configuration.

JCL Members

The following list shows the JCL members that are generated and the tasks they perform. Also, this list notes any tasks you must perform for the JCL member.

- BIND*xx*: These members bind all required SequeLink packages and plans for your DB2 subsystems. Read the comments in this bind job carefully. The tailored job contains a job-step for the defined DB2 interface and bind statements for all data sources that use this interface.

  - BIND7: bind job for DB2 v7
  - BIND8C: bind job for DB2 v8, compatibility mode
  - BIND8N: bind job for DB2 v8, new function mode
  - BIND9: bind job for all DB2 v9 modes.

  NOTE: If data sources are added later, an interface is changed, or the collection-prefixes of a data source are changed, you must regenerate and submit these members.

- CFGPRINT: A summary member you can use for debugging purposes that reports the content of all variables at the time of JCL generation. This member also lists the user and date and time of the last JCL generation.
- EVLDEF: This member creates the event trace file for the server.
- INIT: This member contains operator interface commands that will be executed after the successful startup of the server.

  NOTE: To add or delete operator interface commands you want to execute after the successful startup of the server, edit this member.
- RUNSMF: This member is a sample job that allows SMF records to be printed.
- RUNSRVR: This member is the started task or job that is used to start the server.

  NOTE: If the DB2 interface is changed, you must regenerate this member, and restart the server.
- RUNCLUST: This member prints out the Sysplex registration information.
- RUNWHAT: This member is a diagnostic job that prints all versions of the software components and can be used for debugging. Execute this job only on a specific request from DataDirect Technologies technical support.

Generating the JCL

To generate JCL:

1 Type G beside the SequeLink Server on the ServerList panel; then, press ENTER. The JCL Settings panel appears.

```
DataDirect SequeLink Manager for z/OS - ServerList          Row 1 to 3 of 3
--------------------------------------------------------------------------------
           DataDirect SequeLink Manager for z/OS - JCL Settings

 Change the default settings used for JCL generation:

    Job prefix . . . . . userid
    Job Class  . . . . . A
    Output Class . . . . X
    Accounting Info  . . (acct)
    Generate server as:  1
        1.  Started Task (recommended)
        2.  Normal Job

 Press Enter to continue, F12 to cancel dialog.

 Command ===>
```

2 If you want to change the default settings used for JCL generation, provide the following information:

Job prefix: Type the name of the site-specific job name.

Job Class: Type the name of the site-specific job class.

Output Class: Type the name of the site-specific output class.

Accounting Info: Type your site-specific jobcard account information.

3 Choose whether to generate the JCL for the server as a started task (highly recommended) or as a normal job:

- Type `1` to generate the JCL for the server as a started task; then, press ENTER.
- Type `2` to generate the JCL for the server as a normal job; then, press ENTER.

Data Sets Created During Installation

In previous versions of SequeLink Server for DB2 on z/OS, both the SequeLink installation datasets and the specific SequeLink Server datasets used a high-level qualifier (HLQ) prefix.

This naming convention has been changed to give the SequeLink administrator more control and flexibility over this dataset naming convention:

- Deployment, for example, from development to production, is easier.
- SequeLink Installation libraries are read-only.
- Multiple Server environments, such as Production, Development, and Quality Assurance, can be administered separately. For example:
  - User-specific dataset names can be made unique for each environment
  - The datasets can be allocated on separate volumes.
  - RACF dataset security can be set for each environment.

Two new dataset prefixes have been introduced:

- *HLQ* is the installation dataset prefix. The *HLQ* prefix is required at install time.
- *USR* is the server-specific dataset prefix. The *USR* prefix and the server name (USR.*servername*) are required when a SequeLink Server is added.

The maximum length for each qualifier is 24 characters; the name must be compliant with the operating system dataset naming convention.

The installation libraries are copied to datasets prefixed with the installation data set prefix (HLQ). The SequeLink servers generated and configured with the SequeLink Server Manager for z/OS (SSMC) reside in the datasets prefixed with a server-specific dataset prefix (USR.*servername*), so that datasets related to one SequeLink Server can be grouped, where *USR* specifies the SequeLink Server environment and *servername* is the name of the SequeLink server.

The list of SequeLink servers that will be managed using the SSMC is maintained in the server list dataset. The name consists of the server-specific dataset prefix USR suffixed with '.SERVERS'.

Example:

A company installs SequeLink and configures three SequeLink Servers, one for the production environment and two for the test and development environment. The production server accesses DB2V8 and the testing/development servers access DB2 v8 and DB2 V9.

```
HLQ = SQLNK.V6R0M0
USR1 = COMPANY.PROD
servername1 = SQLDB2V8
USR2 = COMPANY.DEV
servername1 = SQLDB2V8
servername2 = SQLDB2V9
```

This results in the following datasets on your system:

SequeLink installation datasets	■ SQLNK.V6R0M0.CLIST ■ SQLNK.V6R0M0.DBRM ■ SQLNK.V6R0M0.IPE ■ SQLNK.V6R0M0.LINKLIB ■ SQLNK.V6R0M0.LOADLIB ■ SQLNK.V6R0M0.MSGS ■ SQLNK.V6R0M0.PANELS ■ SQLNK.V6R0M0.SKELS ■ SQLNK.V6R0M0.TABLES ■ SQLNK.V6R0M0.SWANDM.INI ■ SQLNK.V6R0M0.README
Serverlist datasets	■ COMPANY.PROD.SERVERS ■ COMPANY.PROD.SQLDB2V8.SWANDM.INI ■ COMPANY.PROD.SQLDB2V8.SWEVLOG ■ COMPANY.PROD.SQLDB2V8.CNTL ■ COMPANY.DEV.SERVERS ■ COMPANY.DEV.SQLDB2V8.SWANDM.INI ■ COMPANY.DEV.SQLDB2V8.SWEVLOG ■ COMPANY.DEV.SQLDB2V8.CNTL ■ COMPANY.DEV.SQLDB2V9.SWANDM.INI ■ COMPANY.DEV.SQLDB2V9.SWEVLOG ■ COMPANY.DEV.SQLDB2V9.CNTL

7 Configuring SequeLink® Services Using the SequeLink® Manager for z/OS

z/OS

This chapter describes the tasks you may need to perform to configure and manage SequeLink services and server data sources locally from a z/OS machine.

NOTES:

- The server does not have to be offline to change its configuration; however, remember that some configuration changes do not take effect until you restart the server. Also, regeneration of the JCL may be necessary, depending on the changes that are made. See "Service Attributes" on page 37 for a description of how attributes are used.
- To delete a service, the server must be offline.

Configuring SequeLink® Servers

This section describes how to create and manage SequeLink Servers locally on z/OS.

To do this...	See...
Create a SequeLink Server	"Creating a SequeLink® Server" on page 130
Delete a SequeLink Server	"Deleting a SequeLink® Server" on page 134

Creating a SequeLink® Server

On z/OS, a SequeLink Server is a started task/job name that corresponds to a single data access service. You can define multiple SequeLink Servers on the same z/OS machine. Each SequeLink Server has SequeLink Agent functionality that is included in the data access service. When you add a SequeLink Server, you are prompted to add a DB2 interface and a service.

To create a SequeLink Server:

1 From the ServerList panel, type `A` at the Command prompt to create a SequeLink Server; then press ENTER. The Add Server panel appears.

```
         DataDirect SequeLink Manager for z/OS - Add server
Command ===>

Complete the following fields for the new server:

 Server name . . . . . . .SQLDB2V8
 Description
 Development SequeLink for DB2v8  (subsys:DB8A)

 Hostname/IP-address
 10.30.14.109
 Portnumber  . . . . . . .26804
```

2 Provide the following information, then press ENTER:

Server name: Type the name of the new SequeLink Server definition. The server name you choose will be the job name or started task name.

Description: Type a description for the new SequeLink Server definition.

Hostname/IP-address: Type the TCP/IP host name or the TCP/IP address of the SequeLink Server definition. The HOME address of the local TCP/IP stack or the standard 127.0.0.1 TCP/IP loopback address will be inserted. In some instances, such as when the TSO session for the SequeLink Manager is

running on another IP stack than the server, you must specify the TCP/IP address of the server machine.

Portnumber: Type the port to be used by the SequeLink service.

NOTE: When you create a server, a server-specific configuration file is created. If such a file already exists, you are prompted to confirm whether you want to overwrite the existing configuration file.

3 You are prompted to supply parameters to the DB2 interface.

```
          DataDirect SequeLink Manager for z/OS - Add DB2 Interface
Command ===>

Enter parameters for the DB2 interface for server SQLDB2V8 : IMPORTANT
after adding the DB2 interface ensure to add the service attribute
MVSServiceDB2InterfaceID.

Interface ID . . . DSN8         DB2 Subsystem ID       DB8N
DB2 Version  . . . V810  (eg.:V810)

Description
DB2 version 8 Development System
DB2 Loadlib  DSN810.SDSNLOAD
DB2 Exitlib  DSN810.DB8N.SDNSEXIT
```

4 Provide the following information; then, press ENTER.

Interface ID: Type a Logical ID identifying the DB2 interface for the SequeLink Server, for example, DSN8. The corresponding service attribute is MVSServiceDB2InterfaceID.

DB2 Subsystem Id: Type the SubsystemId of the DB2 subsystem you want to access, for example, DBR8. The corresponding service attribute is MVSDB2SubsystemName.

DB2 Version: Type the version of DB2 to be used, for example, V810. The corresponding service attribute is MVSDB2Version.

Description: Type a description of the DB2 interface. The corresponding service attribute is MVSDB2RootDescription.

DB2 Loadlib: Type the name of the DB2 load library, for example, DB2.V8R1M0.SDSNLOAD. The corresponding service attribute is MVSDB2LoadLibrary.

DB2 Exitlib: Type the name of the DB2 exit library, for example, DB2.V8R1M0.DB8R.SDSNLOAD. The corresponding service attribute is MVSDB2ExitLibrary.

5 You are prompted to supply parameters to define the service.

```
              DataDirect SequeLink Manager for z/OS - Add service
Command ===>

Complete the following fields for the new service for server ACCT1:

Service name . . . .
Description

Portnumber . . . .  3456
```

6 Provide the following information; then, press ENTER.

```
              DataDirect SequeLink Manager for z/OS - Add service
Command ===>

Complete the following fields for the new service for server SQLDB2V8:

Service name . . . .ACCT1
Description
Accounting

Portnumber . . . .  26804
```

Service name: Type the service name. The corresponding service attribute is ServiceName.

Description: Type a description of the service. The corresponding service attribute is ServiceDescription.

Portnumber: The port number entered in the Server Definition panel is displayed here.

7 The ServerList panel shows the SequeLink Server you just created.

```
 DataDirect SequeLink Manager for z/OS - ServerList          Row 1 to 1 of 1
Command ===>                                                   SCROLL > CSR

Allowed actions or commands are:

 o  (S)elect a server       o  (G)enerate JCL
 o  (A)dd a server         o  (O)perator interface
 o  (D)elete a server

   Name        Description
-------------------------------------------------------------------------------
   SQLDB2V8    Development SequeLink for DB2v8 (subsys: DB8A)
```

Once you have created a SequeLink Server, type `S` beside it to view the server management tree. The server management tree shows the DB2 interface, UID maps, and the SequeLink service and its configuration settings.

See "Using the Server Management Tree" on page 118 for more information about using the server management tree.

Deleting a SequeLink® Server

1 From the ServerList panel, type D beside the SequeLink Server you want to delete; then, press ENTER. The Delete Server panel appears, prompting you to confirm the deletion.

```
          DataDirect SequeLink Manager for z/OS - Delete server

You are about to delete the following server from the list of known
servers.

 Server name . . . . . . : ACCT1
 Description
 DB2 V8 SERVICE USING IBM TCP/IP
 Hostname/IP-address
 10.25.129.07
 Portnumber  . . .       : 8042
 JCL output library  . . : DDTEK.SL60.ACCT1.CNTL
 SequeLink config. file  : DDTEK.SL60.ACCT1.SWANDM.INI

       Are you sure this is what you want to do ?
       2   1.  Yes, go ahead.
           2.  No, don't delete.

       Also delete any data sets used by the server ?
       2   1.  Yes, delete the data sets.
Command ===>
```

2 Confirm or cancel the deletion:

- Type 1 to confirm the deletion; then, press ENTER.
- Type 2 to cancel the deletion; then, press ENTER.

The SequeLink Server is deleted, and you are returned to the list of servers.

Configuring SequeLink® Services

This section describes how to create and manage SequeLink services locally on z/OS.

To do this...	See...
Create a SequeLink service	"Creating a SequeLink® Service" on page 135
Delete a SequeLink service	"Deleting a SequeLink® Service" on page 136
View service attributes	"Viewing SequeLink® Service Attributes" on page 137
Change a service attribute	"Changing a SequeLink® Service Attribute" on page 139
Add a service attribute	"Adding a SequeLink® Service Attribute" on page 141
Delete a service attribute	"Deleting a SequeLink® Service Attribute" on page 143
Configure monitoring	"Configuring Monitoring" on page 144
Configure event tracing	"Configuring Event Tracing" on page 149

Creating a SequeLink® Service

NOTE: You can create only one service for any SequeLink Server. Usually, you will be prompted to create a service while you are creating a SequeLink Server. You can also create or recreate a service *after* you delete an existing service.

To create a SequeLink service:

1 Type `A` beside the SequeLink Services node to add a service to the SequeLink Server; then, press ENTER.

2 The Add service panel or Add DB2 Interface panel appears. Complete the panel as described in "Creating a SequeLink® Server" on page 130.

See "Changing a SequeLink® Service Attribute" on page 139 for instructions on changing SequeLink service attributes.

NOTE: You must restart the SequeLink Server before this change will take effect.

Deleting a SequeLink® Service

1 From the server management tree, type `D` beside the SequeLink service you want to delete; then, press ENTER. The Delete Service panel appears, prompting you to confirm the deletion of the SequeLink service.

```
        DataDirect SequeLink Manager for z/OS - Delete service
Command ===>

You are about to delete the following service from server ACCT1:::::

Service name . . . : ACCT1
Description
ACCT DB2
Portnumber . . . . : 8042

      Are you sure this is what you want to do ?
      2   1.  Yes, go ahead.
          2.  No, don't delete.
```

2 Confirm or cancel the deletion:

- Type `1` to confirm the deletion; then, press ENTER
- Type `2` to cancel the deletion; then, press ENTER.

The SequeLink service is deleted, and you are returned to the server management tree.

NOTE: Create a new service before restarting the SequeLink Server.

Viewing SequeLink® Service Attributes

1 From the server management tree, type `/` beside the Service Settings node of the SequeLink service to expand it. The server management tree shows the attribute categories for the service.

2 Type `S` beside the attribute category for which you want to view attributes; then, press ENTER. The AttributeList panel appears for that category, listing all the attributes and their values.

For example, if you select the Logging category, the following AttributeList panel appears:

```
        DataDirect SequeLink Manager for z/OS - AttributeList
                                                        Row 1 to 2 of 2
                                                              MORE >>>
Service ACCT1
Enter the 'ADD'-command to add an attribute or
perform one of the actions below on a specific attribute

o (S)elect    o (C)hange    o (D)elete    o (?)Help
  Name                            Value
-----------------------------------------------------------------------
  ServiceDebugLogLevel            1
  ServiceDebugLogPath             /tmp
***************************** Bottom of data ****************************

COMMAND ===>                                            SCROLL > PAGE
```

For information about:

- Changing the value of an attribute, see "Changing a SequeLink® Service Attribute" on page 139.
- Adding attributes, see "Adding a SequeLink® Service Attribute" on page 141.
- Deleting an attribute, see "Deleting a SequeLink® Service Attribute" on page 143.

NOTE: To display help about an attribute, type ? beside the attribute; then, press ENTER.

Changing a SequeLink® Service Attribute

1. From the server management tree, type / beside the Service Settings node of a SequeLink service to expand it. The server management tree shows the attribute categories for the service.

2. Type S beside an attribute category to select it; then, press ENTER. The AttributeList panel appears for that category, listing all the attributes and their values.

   For example, if you select the Logging attribute category, the following AttributeList panel appears:

```
            DataDirect SequeLink Manager for z/OS - AttributeList
                                                          Row 1 to 3 of 3
Dialog canceled.                                               MORE >>>

Service ACCT1
Enter the 'ADD'-command to add an attribute or
perform one of the actions below on a specific attribute
o (S)elect    o (C)hange    o (D)elete    o (?)Help

  Name                            Value
-------------------------------------------------------------------------
  ServiceDebugLogLevel            1
  ServiceDebugLogPath             /tmp
  ServiceEventTraceSize           1000000
**************************** Bottom of data ****************************

COMMAND ===>                                                SCROLL > PAGE
```

   NOTE: To display help about an attribute, type ? beside the attribute; then, press ENTER.

3 Type `C` beside the attribute you want to change; then, press ENTER. The Attribute Display panel appears with the cursor positioned at the value field of the attribute.

For example, to change the value of the ServiceDebugLogLevel attribute to 4 (Debug), the following Attribute Display panel appears:

```
                DataDirect SequeLink Manager for z/OS
                                                    Row 1 to 7 of 7

  ServiceDebugLogLevel
  Bitmask currently in effect:

  Use (S) or (/) to select bits to set.

    Bit description
  --------------------------------------------
  / Fatal
    Errors
    Warnings
    Informational
    Debug
    SSP Packet Log
    SSP Requests
  COMMAND ===>                  SCROLL > PAGE
```

4 Type the new value of the attribute; then, press ENTER. You are returned to the AttributeList panel, and the attribute value, if valid, is changed.

5 Press F3 to return to the server management tree.

Adding a SequeLink® Service Attribute

1. From the server management tree, type / beside the Service Settings node of the SequeLink service to expand it. The server management tree shows the attribute categories for the service.

2. Type S beside the attribute category to select it; then, press ENTER. The AttributeList panel appears for that category, listing all the attributes and their values.

   For example, if you select the Logging category, the following AttributeList panel appears:

```
            DataDirect SequeLink Manager for z/OS - AttributeList
                                                          Row 1 to 3 of 3
  Dialog canceled.                                              MORE >>>

  Service ACCT1
  Enter the 'ADD'-command to add an attribute or
  perform one of the actions below on a specific attribute
  o (S)elect    o (C)hange    o (D)elete    o (?)Help

    Name                            Value
-----------------------------------------------------------------------------
    ServiceDebugLogLevel            1
    ServiceDebugLogPath             /tmp
    ServiceEventTraceSize           1000000
  **************************** Bottom of data **************************

  COMMAND ===>                                                SCROLL > PAGE
```

3. Type Add at the command prompt to add an attribute; then, press ENTER. A panel appears listing the attributes you are allowed for that category.

For example, to add an attribute to the Logging category, the following panel appears:

```
  DataDirect SequeLink Manager for z/OS
                              Row 1 to 2 of 2

This is a list of attributes,
you are allowed to add:
Use (S) or (/) to select an attribute.
Use (?) for help on an attribute.

  Attribute
-------------------------------------------
  ServiceEventTraceSize
************** Bottom of data **************


COMMAND ===>                  SCROLL > PAGE
```

NOTE: To display help about an attribute, type `?` beside the attribute; then, press ENTER.

4 Type `S` beside the attribute you want to add; then, press ENTER. The Attribute Display panel appears with the cursor positioned at the value field of the attribute.

NOTE: When you add an attribute, the Value field will display the default value, if a default exists for the attribute. To accept the default, press ENTER.

For example, to add the ServiceEventTraceSize attribute, the following Attribute Value panel appears:

```
                Attribute Value

 Press F1 for help, F3 to leave.

 ServiceEventTraceSize
                ===> 0      ( 10000 - 2000000000)
```

5 Type a value for the attribute or use the default if one is available; then, press ENTER.

6 You are returned to the AttributeList panel. Press F3 to return to the server management tree.

Deleting a SequeLink® Service Attribute

1 From the server management tree, type / beside the Service Settings node of the SequeLink service to expand it. The server management tree shows the attribute categories for the service.

2 Type S beside the attribute category to select it; then, press ENTER. The AttributeList panel appears for that category, listing all the attributes and their values.

For example, if you select the Logging category, the following AttributeList panel appears:

```
          DataDirect SequeLink Manager for z/OS - AttributeList
                                                          Row 1 to 3 of 3
Dialog canceled.                                               MORE >>>

Service ACCT1
Enter the 'ADD'-command to add an attribute or
perform one of the actions below on a specific attribute
o (S)elect    o (C)hange    o (D)elete    o (?)Help

  Name                              Value
-------------------------------------------------------------------------
  ServiceDebugLogLevel              1
  ServiceDebugLogPath               /tmp
  ServiceEventTraceSize             1000000
**************************** Bottom of data ****************************


COMMAND ===>                                                SCROLL > PAGE
```

3 Type D beside the attribute you want to delete; then, press ENTER. The attribute is deleted.

Configuring Monitoring

SequeLink provides the following levels of monitoring for SequeLink data access services, listed here from highest-level to lowest-level. Some of the monitoring levels are also supported for the SequeLink Agent service.

- **Service monitoring** monitors these activities by service:
  - Statistics of received packets and sent packets
  - Sessions started and statements opened
  - Active statements and sessions
  - Fetched rows and affected rows
  - Transactions
- **Session monitoring** monitors these activities by session within a service:
  - Statistics of received packets and sent packets
  - Statements opened and active statements
  - Fetched rows and affected rows
  - Transactions
  - Information about each session, such as start time, client information (network address, data source used by the client, and type of client), native database session identification, and database user
- **Statement monitoring** monitors these activities by statement within a session:
  - Fetched rows and affected rows
  - SQL statements issued

To enable monitoring at one of the listed levels, higher-level monitoring must be enabled. For example, you cannot monitor Session information unless Service monitoring is enabled.

Similarly, you cannot monitor Statement information unless both Service monitoring and Session monitoring are enabled.

See Chapter 4 "Managing Data Access Activity Using the SequeLink® Manager Snap-in" on page 91 for information about using the SequeLink Manager Snap-in to monitor SequeLink service activity.

To configure monitoring locally on the z/OS:

1. From the server management tree of the SequeLink Server, type / beside Configuration node of the SequeLink service for which you want to turn on monitoring; then, press ENTER.

2. Type / beside the Profiles node to expand it, if necessary; then, press ENTER.

   The Profiles node lists all monitoring and event tracing profiles enabled for the SequeLink service. You can perform the following actions:

   - To add a profile, see "Creating a Monitoring Profile" on page 146.
   - To change a profile, see "Changing a Monitoring Profile" on page 148.
   - To delete a profile, see "Deleting a Monitoring Profile" on page 149.

Creating a Monitoring Profile

1 Type `A` beside the Profiles node to add a profile to the SequeLink service; then, press ENTER. The Add Profile panel appears prompting you to select the type of profile to add.

```
                                  Attribute Value

  DataDirect SequeLink Manager for z/OS - Add Profile

Select type of profile to add:

1   1.  Monitor Profile.
    2.  Event Trace Profile.

Press Enter to continue.

Command ===>
```

2 Type 1 to add a monitoring profile; then, press ENTER. The Monitor Profile panel appears.

```
------------------------------ Attribute Value ------------------------------
         DataDirect SequeLink Manager for z/OS - Monitor Profile

 These are the counters and events to be monitored:

 / Enable Service Monitoring:

     received packet size(avg)    statements opened    / active sessions
     sent packet size(avg)        active statements      affected rows
   / sessions started             fetched rows           transactions

 / Enable Session Monitoring:

     received packet size(avg)    active statements      transactions
     sent packet size(avg)        fetched rows         / info
     statements opened            affected rows

   Enable Statement Monitoring:

     fetched rows (last)          affected rows  (last)       sql
     fetched rows                 affected rows

 Command ===>
```

3 Type / beside an option to enable it or clear the / to disable an option. When you are satisfied with your settings, press F3. You are returned to the server management tree.

4 Restart the SequeLink Server to activate the profile you just created.

Changing a Monitoring Profile

1. From the server management tree, type `S` beside the profile you want to change; then, press ENTER. The Monitor Profile panel appears.

```
------------------------------ Attribute Value ------------------------------
         DataDirect SequeLink Manager for z/OS - Monitor Profile

 These are the counters and events to be monitored:

 / Enable Service Monitoring:

     received packet size(avg)    statements opened    / active sessions
     sent packet size(avg)        active statements      affected rows
   / sessions started             fetched rows           transactions

 / Enable Session Monitoring:

     received packet size(avg)    active statements      transactions
     sent packet size(avg)        fetched rows         / info
     statements opened            affected rows

   Enable Statement Monitoring:

     fetched rows (last)          affected rows  (last)       sql
     fetched rows                 affected rows

 Command ===>
```

2. Type / beside an option to enable it or clear the / to disable an option. When you are satisfied with your settings, press ENTER.

3. Restart the SequeLink Server to activate the monitoring profile you just changed.

Deleting a Monitoring Profile

1. From the server management tree, type `D` beside the profile you want to delete; then, press ENTER. The profile is deleted.
2. Restart the SequeLink Server.

Configuring Event Tracing

Events are generated when the client application accesses data and when specific server activities occur, such as when a service starts or an error occurs. Depending on which SequeLink profiles are active, the information generated by the event is displayed as it occurs in the runtime monitor and is stored persistent in the event trace file.

On z/OS, the event trace file must be created using the EVLDEF member from the server's CNTL library.

See Chapter 4 "Managing Data Access Activity Using the SequeLink® Manager Snap-in" on page 91 for information about using the SequeLink Manager Snap-in for event tracing.

To configure event tracing locally on the z/OS:

1. From the server management tree, type `/` beside Configuration node of the SequeLink service for which you want to turn on monitoring; then, press ENTER.
2. Type `/` beside the Profiles node to expand it, if necessary; then, press ENTER.

The Profiles node lists all monitoring and event tracing profiles enabled for the SequeLink service. You can perform the following actions:

- To add a profile, see "Creating an Event Trace Profile" on page 150.
- To change a profile, see "Changing an Event Trace Profile" on page 152.
- To delete a profile, see "Deleting an Event Trace Profile" on page 153.

Creating an Event Trace Profile

1 From the server management tree, type `A` beside the Profiles node to add a profile to the SequeLink service; then, press ENTER. The Add Profile panel appears prompting you to select the type of profile you want to add.

```
                    Attribute Value

     DataDirect SequeLink Manager for z/OS - Add Profile

   Select type of profile to add:

   1   1.  Monitor Profile.
       2.  Event Trace Profile.

   Press Enter to continue.

   Command ===>
```

2 Type 2 to add an event tracing profile; then, press ENTER. The Event Trace Profile panel appears, listing the event tracing options available.

```
                    DataDirect SequeLink Manager for z/OS
                                                        Row 1 to 7 of 26
                                                               MORE >>>

   Event Trace Profile for service ACCT1
   Check which events you want to write to the SequeLink Event Trace.
     Use (S) or (/) to select/unselect an event,
     or use (F) to view or set an event filter.

       Trace  Event
-----------------------------------------------------------------------------
  __   False  Event processing Starting
  __   False  Event processing Stopping
  __   True   Cursor Closed
  __   True   Cursor Opened
  __   True   Debug Tracing
  __   True   Error Internal
  __   True   Error Occured

   COMMAND ===>                                          SCROLL > PAGE
```

3 Type / beside an option to enable it or clear the / to disable an option. When you are satisfied with your settings, press F3.

NOTE: Optionally, you can type F beside an event trace option to define a filter for the option. A panel appears allowing you to set the filter for that event trace option. Type the filter; then, press ENTER. See "Filtering Events" on page 595 for more information about filtering and the syntax of filter statements.

You are returned to the server management tree.

4 Restart the SequeLink Server to activate the profile you just created.

Changing an Event Trace Profile

1 From the server management tree, type S beside the profile you want to change; then, press ENTER. The Event Trace Profile panel appears.

```
                    DataDirect SequeLink Manager for z/OS
                                                          Row 1 to 7 of 26
                                                                 MORE >>>

 Event Trace Profile for service ACCT1
 Check which events you want to write to the SequeLink Event Trace.
   Use (S) or (/) to select/unselect an event,
   or use (F) to view or set an event filter.

     Trace  Event
 ---------------------------------------------------------------------------
__   False  Event processing Starting
__   False  Event processing Stopping
__   True   Cursor Closed
__   True   Cursor Opened
__   True   Debug Tracing
__   True   Error Internal
__   True   Error Occured

 COMMAND ===>                                              SCROLL > PAGE
```

2 Type / beside an event trace option to enable it or disable it; then, press ENTER.

NOTE: Optionally, you can type F beside an event trace option to define a filter for the option. A panel appears allowing you to set the filter for that event trace option. Type the filter; then, press ENTER. See "Filtering Events" on page 595 for more information about filtering and the syntax of filter statements.

You are returned to the server management tree.

3 Restart the SequeLink Server to activate the profile you just changed.

Deleting an Event Trace Profile

1. From the server management tree, type `D` beside the profile you want to delete; then, press ENTER. The profile is deleted.
2. Restart the SequeLink Server.

Configuring Server Data Sources

This section describes how to create and manage server data sources locally using the SequeLink Manager for z/OS.

To do this...	See...
Create a server data source	"Creating a Server Data Source" on page 154
Delete a server data source	"Deleting a Server Data Source" on page 156
View server data source attributes	"Viewing Server Data Source Attributes" on page 158
Add a server data source attribute	"Adding a Server Data Source Attribute" on page 159
Change a server data source attribute	"Changing a Server Data Source Attribute" on page 163
Delete a server data source attribute	"Deleting a Server Data Source Attribute" on page 165

Creating a Server Data Source

When you create a server data source, the attributes for the new server data source are copied from the default data source. Once you create a server data source, you can change any attributes of the new server data source. See "Viewing Server Data Source Attributes" on page 158 for instructions on viewing server data source attributes.

To create a server data source:

1 Type / beside the Datasource Settings node of the SequeLink service to expand it; then, press ENTER. The expanded Datasource Settings node shows the default server data source, which is named Default, and any other server data sources defined for the SequeLink service.

```
                    DataDirect SequeLink Manager for z/OS

Management Tree for server ACCT1
  To see a list of valid actions on a node, type '?' beside it.
  Use '/' to expand or collapse tree branches.
  Valid commands are: EXP SAVE REFRESH ERRSTK
-----------------------------------------------------------------------
            - Logging
            - Others
            - User Security
          - Datasource Settings
            + Default
          - Profiles
            - Monitoring
            - Event Tracing

Command ===>                                          Scroll > PAGE
```

2 Type A beside the Datasource Settings node; then, press ENTER. The Add Datasource panel appears:

```
        DataDirect SequeLink Manager for z/OS - Add datasource
Command ===>

Enter the parameters for the new datasource for service ACCT1 in
server ACCT1:

Datasource name

Description
```

3 Provide the following information; then, press ENTER.

Datasource name: Type the name of the new server data source. The corresponding service attribute is DataSourceName.

Description: Type a description for the server data source. The corresponding service attribute is DataSourceDescription.

4 You are returned to the server management tree, and the panel reminds you to review the attribute settings for the server data source. See "Viewing Server Data Source Attributes" on page 158 and "Changing a Server Data Source Attribute" on page 163 for instructions on viewing and changing server data source attributes.

The server data source you created can be used immediately for incoming connections; you do not need to restart the SequeLink Server.

Adding a UID Map to a SequeLink® Server Data Source (Optional)

UID mapping is the optional mapping of user IDs to alternate user IDs using a UID map. You can specify a UID map for a SequeLink service and another UID map for a SequeLink Server server data source. See "Adding a UID Map to the SequeLink® Service" on page 423 for more information about adding a UID map to a SequeLink Server service.

See "Using UID Mapping" on page 329 for more information about using UID maps.

Deleting a Server Data Source

NOTE: You cannot delete the Default server data source.

1. Type / beside the Datasource Settings node of the SequeLink service to expand it; then, press ENTER. The expanded Datasource Settings node shows the default server data source, which is named Default, and any other server data sources defined for the SequeLink service.

2 To delete a server data source, type D beside the Datasource Settings node; then, press ENTER. The Delete Datasource panel appears prompting you to confirm the deletion.

```
         DataDirect SequeLink Manager for z/OS - Delete datasource

You are about to delete the following datasource from service ACCT1 in
server ACCT1:

Datasource name
ACCTNW
Description
Data source for accounting NW

      Are you sure this is what you want to do ?
      2   1.  Yes, go ahead.
          2.  No, don't delete.

Command ===>
```

3 Confirm or cancel the deletion:

- Type 1 to confirm the deletion; then, press ENTER.
- Type 2 to cancel the deletion; then, press ENTER.

The server data source is deleted, and you are returned to the server management tree.

Viewing Server Data Source Attributes

1. Type / beside the Datasource Settings node of the SequeLink service to expand it; then, press ENTER. The expanded Datasource Settings node shows the default server data source, which is named Default, and any other server data sources defined for the SequeLink service.
2. Type / beside the server data source you want to view attributes for; then, press ENTER. The server data source node expands to show the data source attribute categories.

```
                      DataDirect SequeLink Manager for z/OS

Management Tree for server ACCT1
  To see a list of valid actions on a node, type '?' beside it.
  Use '/' to expand or collapse tree branches.
  Valid commands are: EXP SAVE REFRESH ERRSTK
-----------------------------------------------------------------------
           - Datasource Settings
            + Default
            - ACCTNW
              - Advanced
              - Application security
              - General
              - Others
              - User Security
              - WorkArounds
          - Profiles

Command ===>                                          Scroll > PAGE
```

3 Type S beside any attribute category to view the attributes set for the server data source; then, press ENTER. The AttributeList panel appears, listing all the attributes and their values configured for the server data source.

For example, if you select the Advanced category, the following AttributeList panel appears:

```
         DataDirect SequeLink Manager for z/OS - AttributeList
                                                        Row 1 to 2 of 2
                                                              MORE >>>

DataSource ACCTNW
Enter the 'ADD'-command to add an attribute or
perform one of the actions below on a specific attribute
o (S)elect    o (C)hange    o (D)elete    o (?)Help

  Name                              Value
-----------------------------------------------------------------------
  DataSourceCursorHold              True
  DataSourceTransactionIsolation    Committed
***************************** Bottom of data ******************************


COMMAND ===>                                              SCROLL > PAGE
```

Adding a Server Data Source Attribute

1 Type / beside the Datasource Settings node of the SequeLink service to expand it; then, press ENTER. The expanded Datasource Settings node shows the default server data source, which is named Default, and any other server data sources defined for the SequeLink service.

2 Type / beside the server data source you want to add an attribute for; then, press ENTER. The server data source node expands to show the data source attribute categories.

```
                      DataDirect SequeLink Manager for z/OS

Management Tree for server ACCT1
  To see a list of valid actions on a node, type '?' beside it.
  Use '/' to expand or collapse tree branches.
  Valid commands are: EXP SAVE REFRESH ERRSTK
-------------------------------------------------------------------------
           - Datasource Settings
            + Default
            - ACCTNW
              - Advanced
              - Application security
              - General
              - Others
              - User Security
              - WorkArounds
          - Profiles

Command ===>                                              Scroll > PAGE
```

3 Type `S` beside any attribute category to view the attributes set for the server data source; then, press ENTER. The AttributeList panel appears, listing all the attributes and their values configured for the server data source.

For example, if you select the Advanced category, the following AttributeList panel appears:

```
          DataDirect SequeLink Manager for z/OS - AttributeList
                                                        Row 1 to 2 of 2
                                                              MORE >>>

DataSource ACCTNW
Enter the 'ADD'-command to add an attribute or
perform one of the actions below on a specific attribute
o (S)elect    o (C)hange    o (D)elete    o (?)Help

  Name                            Value
-------------------------------------------------------------------------
  DataSourceCursorHold            True
  DataSourceTransactionIsolation  Committed
**************************** Bottom of data ****************************


COMMAND ===>                                              SCROLL > PAGE
```

4 Type `Add` at the command prompt to add an attribute. A list of attributes you can add for this attribute category appear.

For example, if you choose to add an attribute from the Advanced category, the following panel appears:

```
        DataDirect SequeLink Manager for z/OS
                                Row 1 to 4 of 4

   This is a list of attributes,
   you are allowed to add:

   Use (S) or (/) to select an attribute,
   Use (?) for help on an attribute.
     Attribute
   ---------------------------------------------
     DataSourceArrayFetchMaxBytes
     DataSourceDisableWarnings
     DataSourceSchemaFilterList
     DataSourceTableTypeFilterList
   ************** Bottom of data **************


   COMMAND ===>                  SCROLL > PAGE
```

NOTE: To display help about an attribute, type `?` beside the attribute; then, press ENTER.

5 Type `S` beside the attribute you want to add; then, press ENTER. The Attribute Value panel appears with the cursor positioned on the Value field.

For example, selecting the DataSourceSchemaFilterList attribute causes the following panel to appear:

```
                      Attribute Value
Press F1 for help, F3 to leave.

DataSourceSchemaFilterList

                  ===>
```

6 Type a value for the attribute or use the default if one is available; then, press ENTER.

7 You are returned to the AttributeList panel. Press F3 to return to the server management tree.

Changing a Server Data Source Attribute

1 Type / beside the Datasource Settings node of the SequeLink service to expand it; then, press ENTER. The expanded Datasource Settings node shows the default server data source, which is named Default, and any other server data sources defined for the SequeLink service.

2 Type / beside the server data source to view its attributes; then, press ENTER. The server data source node expands to show the data source attribute categories.

```
                 DataDirect SequeLink Manager for z/OS

Management Tree for server ACCT1
  To see a list of valid actions on a node, type '?' beside it.
  Use '/' to expand or collapse tree branches.
  Valid commands are: EXP SAVE REFRESH ERRSTK
-------------------------------------------------------------------
            - User Security
          - Datasource Settings
            + Default
            - ACCTNW
              - Advanced
              - Application security
              - General
              - Others
              - User Security
              - WorkArounds
          - Profiles

Command ===>                                        Scroll > PAGE
```

3 Type S beside any attribute category to view the attributes set for the data source; then, press ENTER. The AttributeList panel appears, listing all the attributes and their values configured for the server data source.

For example, if you select the Advanced category, the following AttributeList panel appears:

```
          DataDirect SequeLink Manager for z/OS - AttributeList
                                                        Row 1 to 2 of 2
                                                              MORE >>>
DataSource ACCTNW
Enter the 'ADD'-command to add an attribute or
perform one of the actions below on a specific attribute
o (S)elect    o (C)hange    o (D)elete    o (?)Help

  Name                            Value
-------------------------------------------------------------------------
  DataSourceCursorHold            True
  DataSourceTransactionIsolation  Committed
**************************** Bottom of data ****************************

COMMAND ===>                                              SCROLL > PAGE
```

4 Type C beside the attribute you want to change; then, press ENTER. The Attribute Display panel appears with the cursor positioned at the value field of the attribute.

For example, to change the value of the DataSourceCursorHold attribute to FALSE, the following panel appears:

```
                  Attribute Value

   Press F1 for help, F3 to leave.

   DataSourceCursorHold
                 ===> True
```

5 Type the new value of the attribute; then, press ENTER. You are returned to the AttributeList panel, and the attribute value, if valid, is changed.

6 Press F3 to return to the server management tree.

Deleting a Server Data Source Attribute

1 Type / beside the Datasource Settings node of the SequeLink service to expand it; then, press ENTER. The expanded Datasource Settings node shows the default server data source, named Default, and any other server data sources defined for the SequeLink service.

2 Type / beside the server data source to view its attributes; then, press ENTER. The server data source node expands to show the data source attribute categories.

```
                    DataDirect SequeLink Manager for z/OS

Management Tree for server ACCT1
  To see a list of valid actions on a node, type '?' beside it.
  Use '/' to expand or collapse tree branches.
  Valid commands are: EXP SAVE REFRESH ERRSTK
-----------------------------------------------------------------------
            - User Security
          - Datasource Settings
            + Default
            - ACCTNW
              - Advanced
              - Application security
              - General
              - Others
              - User Security
              - WorkArounds
          - Profiles

Command ===>                                          Scroll > PAGE
```

3 Type `S` beside any attribute category to view the attributes set for the data source; then, press ENTER. The AttributeList panel appears, listing all the attributes and their values configured for the server data source.

For example, if you select the Advanced category, the following AttributeList panel appears:

```
            DataDirect SequeLink Manager for z/OS - AttributeList
                                                        Row 1 to 2 of 2
                                                              MORE >>>

DataSource ACCTNW
Enter the 'ADD'-command to add an attribute or
perform one of the actions below on a specific attribute
o (S)elect    o (C)hange    o (D)elete    o (?)Help

  Name                            Value
-----------------------------------------------------------------------
  DataSourceCursorHold            True
  DataSourceTransactionIsolation  Committed
**************************** Bottom of data ****************************


COMMAND ===>                                              SCROLL > PAGE
```

4 Type `D` beside the attribute you want to delete; then, press ENTER. The attribute is deleted. You are returned to the server management tree.

Part 2: Configuring and Managing SequeLink® Clients

This part contains the following chapters:

- Chapter 8 “Configuring the ODBC Client” on page 169 describes the tasks you may need to perform to configure and manage the SequeLink Client *for* ODBC.
- Chapter 9 “Configuring the ADO Client” on page 211 describes the tasks you may need to perform to configure and manage the SequeLink Client *for* ADO.
- Chapter 10 “Configuring the JDBC Client” on page 251 describes the tasks you may need to perform to configure and manage the SequeLink Client *for* JDBC.
- Chapter 11 “Configuring the .NET Client” on page 265 describes the tasks you may need to perform to configure and manage the SequeLink Client *for* .NET.

8 Configuring the ODBC Client

This chapter describes the tasks you may need to perform to configure and manage the SequeLink Client *for* ODBC (the ODBC Client).

About ODBC Connections

You can open an ODBC connection to a SequeLink service by specifying a configuring a ODBC client data source. This chapter explains how to connect to a SequeLink service using client data source.

Using the ODBC Administrator

The first step in setting up an ODBC connection is creating an ODBC data source. You use the ODBC Administrator to create and manage ODBC data sources.

To start the ODBC Administrator, click **Start / Programs**. From the Programs menu, select **DataDirect SequeLink 6.0 Client for ODBC**, and then select the **ODBC Administrator** application. The ODBC Data Source Administrator window appears listing resident data sources.

NOTE: An ODBC Administrator does not exist for Linux/UNIX; you must edit the odbc.ini file using a text editor. See "Configuring ODBC Client Data Sources on Linux and UNIX" on

page 191 for instructions on creating ODBC client data sources for Linux/UNIX.

Configuring ODBC Client Data Sources on Windows

To configure client data sources for the ODBC Client on Windows platforms, you use the ODBC Administrator.

Configuring ODBC User and System Client Data Sources

1 Start the ODBC Administrator. To start the ODBC Administrator, select **Start** / **Programs**. From the Programs menu, select **DataDirect SequeLink 6.0 Client for ODBC** or **DataDirect SequeLink 6.0 Client for ODBC 64-bit**, and then select the **ODBC Administrator** application.

2 Click the **User DSN** tab or the **System DSN** tab to list user or system data sources, respectively.

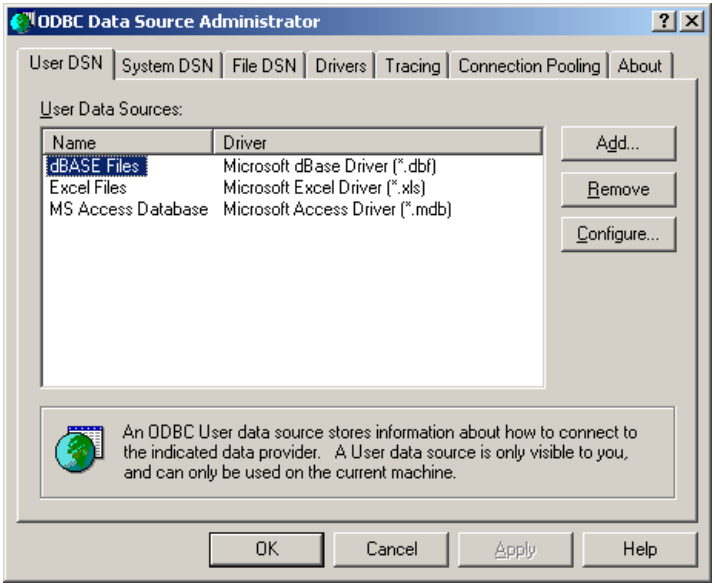

3 To configure a new data source, click the **Add** button. A list of installed drivers appears. Select **DataDirect SequeLink 6.0**; then, click **Finish**.

NOTE: To change an existing data source, select the data source you want to configure and click the **Configure** button.

The DataDirect SequeLink for ODBC Setup window appears.

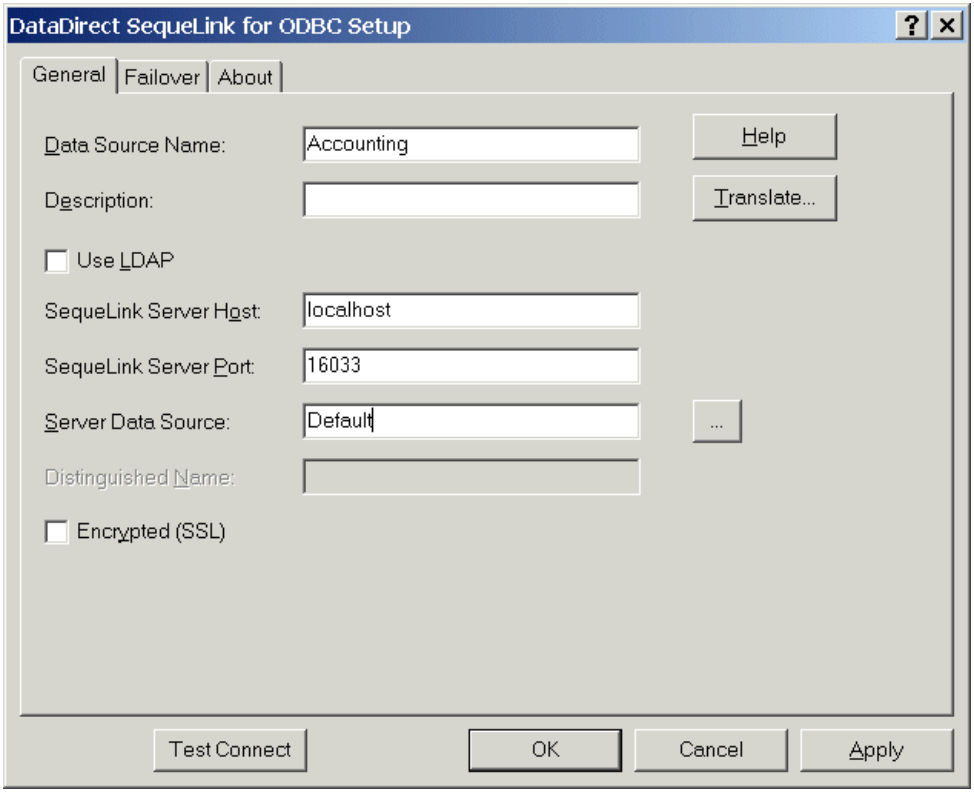

4 On the General tab, provide the following information; then, click **Apply**.

Data Source Name: Type a unique name that identifies this ODBC data source configuration. Examples are `Accounting` or `SequeLink to Oracle Data`.

Description: Optionally, type a description of the data source, for example, `My Accounting Database` or `Accounting Data in Oracle (SSL)`.

SequeLink Server Host: Type the TCP/IP host name of the SequeLink service to which the ODBC Client will connect.

SequeLink Server Port: Type the TCP/IP port the SequeLink service is listening on for connection requests. The port you specify must be the same port that was specified for the

SequeLink service when the SequeLink Server was installed; the default is 19996.

Server Data Source: Type the name of a server data source configured for the SequeLink service to use for the connection, or click the **...** button to select an existing server data source. This field is optional. If a server data source is not specified, the default server data source for that SequeLink service is used.

Use LDAP: To configure the ODBC Client to retrieve connection information from an LDAP directory, select the **Use LDAP** check box. The fields change on the lower half of the screen to accommodate the information required to query an LDAP server for connection information. Provide the following information:

LDAP Server Host: Type the TCP/IP host name of the LDAP server.

LDAP Server Port: Type the TCP/IP port the LDAP server is listening on for connection requests.

Distinguished Name (DN): Type an identifier that uniquely identifies the LDAP entry where the connection information is stored.

Encrypted (SSL): If the remote SequeLink service is configured for Secure Sockets Layer (SSL) encryption, select this check box. If connecting to a SequeLink service enabled for SSL, you must select this check box.

When the check box is cleared (the default), communication between the SequeLink Client and SequeLink Server is not encrypted.

Configuration of encryption is performed on the SequeLink Server. See "Using SSL Encryption" on page 209 for a discussion of encrypting data.

NOTES:

- An ODBC client data source can reference an LDAP directory to retrieve server connection information. See "Retrieving Connection Information from LDAP Directories" on page 428 for more information about retrieving connection information from LDAP directories.
- SSL encryption is not supported for LDAP Servers. The Use LDAP and the Encrypted (SSL) check boxes are mutually exclusive.
- SSL encryption is not supported for SequeLink Server for DB2 for z/OS. Do not select the Encrypted (SSL) check box for a DB2 for z/OS configuration.

Translate: Click **Translate** only if you want to configure an ODBC translator.

NOTE: We strongly recommend that you do not configure an ODBC translator; instead, rely on the native SequeLink transliteration between server and client code pages.

The Select Translator dialog box appears, listing translators specified in the ODBC Translators section of the system information file. Select a translator. When satisfied with your choice, click **OK** to close this dialog box and perform the translation.

5 Optionally, click the **Failover** tab to specify Failover data source settings.

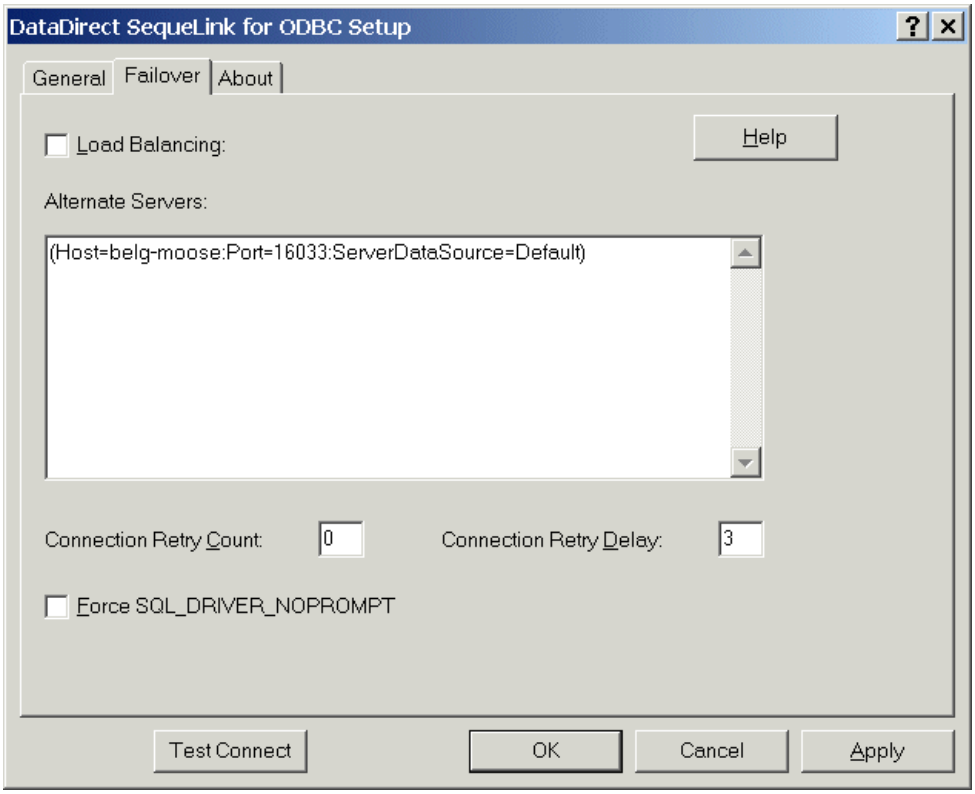

Provide any of the following information; then, click **Apply**.

Load Balancing: Select this check box to allow the driver to use client load balancing in its attempts to connect to primary and alternate database servers. In this case, the driver attempts to connect to the database servers in random order.

If this check box is not selected (the default), client load balancing is not used and the driver connects to each database server based on its sequential order (primary server first, then, alternate servers in the order they are specified).

NOTE: This option has no effect unless alternate servers are defined for the Alternate Servers connection option.

The Load Balancing option is an optional setting that you can use in conjunction with connection failover. See "Configuring Connection Failover for the ODBC Client" on page 197 for a discussion of connection failover and for information about other connection options that you can set for this feature.

Alternate Servers: Type a list of alternate SequeLink servers to which the driver will try to connect if the primary SequeLink server is unavailable, using a string that defines the physical location of each alternate server. Specifying a value for this option enables connection failover for the driver. See "Configuring Connection Failover for the ODBC Client" on page 197 for a discussion of connection failover.

IMPORTANT: If you specified an LDAP server in the LDAP Server Host field, the alternate servers *must* be LDAP servers.

The server name and port are required for each alternate server entry. All of the other required connection information for each alternate server is the same as what is defined for the primary server connection. Currently, the only optional property that can be set for the alternate server is Server Data Source.

The string has the format:

```
(Host=servername1:Port=port1[:ServerDataSource=
serverdatasourcename1], Host=servername2:Port=
port2[:ServerDataSource=serverdatasourcename2],...)
```

For example, the following Alternate Servers value defines two alternate SequeLink servers for connection failover:

```
(Host=server2:Port=19996:ServerDataSource=SDSN2,Host=
server3:Port=19996:ServerDataSource=SDSN3)
```

If you are connecting to an LDAP server, the syntax includes the physical location of the server and the port number:

```
(Host=ld1.foo.com:Port=389,Host=ld2.foo.com:Port=
389,Host=ld3.foo.com:Port=389)
```

Connection Retry Count: Type a value to specify the number of times the driver tries to connect to the primary server and, if configured, to the alternate servers after the initial unsuccessful attempt.

Valid values are integers from 0 to 65535. When set to 0 (the default), the driver does not try to connect after the initial unsuccessful attempt.

If a connection is not established during the retry attempts, the driver returns an error that is generated by the first server to which it tried to connect.

This option and the Connection Retry Delay connection option, which specifies the wait interval between attempts, can be used in conjunction with connection failover.

See "Configuring Connection Failover for the ODBC Client" on page 197 for a discussion of connection failover and for information about other connection options that you can set for this feature.

Connection Retry Delay: Type a value to specify the number of seconds that the driver waits after the initial unsuccessful connection attempt before retrying a connection to the primary server and, if specified, to the alternate servers.

Valid values are integers from 0 to 65535. The default value is 3 (seconds). When set to 0, there is no delay between retries.

NOTE: This option has no effect unless the Connection Retry Count connection option is set to an integer value greater than 0.

This option and the Connection Retry Count connection option, which specifies the number of times the driver tries to connect after the initial unsuccessful attempt, are used in conjunction with connection failover.

See "Configuring Connection Failover for the ODBC Client" on page 197 for a discussion of connection failover and for information about other connection options that you can set for this feature.

Force SQL_DRIVER_NOPROMPT: Select this check box when connection failover or load balancing is enabled. This check box must be also selected if the application cannot change the DriverCompletion argument to SQL_DRIVER_NOPROMPT.

If this check box is not selected (the default), the behavior of the application is not changed.

6 At any point during the configuration process, you can click **Test Connect** to attempt to connect to the data source using the connection properties specified in the driver Setup dialog box. A logon dialog box appears; see "ODBC Connection Dialogs" on page 183 for details.

Note that the information you enter in the logon dialog box during a test connect is not saved.

- If the driver can connect, it releases the connection and displays a Connection Established message. Click **OK**.
- If the driver cannot connect because of an improper environment or incorrect connection value, it displays an appropriate error message. Click **OK**.

NOTE: If you are configuring alternate servers for use with the connection failover feature, be aware that the Test Connect button tests only the primary server, not the alternate servers.

7 Click **OK** or **Cancel**. If you click **OK**, the values you have specified become the defaults when you connect to the data source. You can change these defaults by using this procedure to reconfigure your data source. You can override these defaults by connecting to the data source using a connection string with alternate values.

Configuring ODBC File Client Data Sources

File data sources are data source files stored on a file server. The files are available to any user who can access the server.

To configure ODBC file client data sources:

1. Start the ODBC Administrator by clicking **Start / Programs**. From the Programs menu, select **DataDirect SequeLink 6.0 Client for ODBC** or **DataDirect SequeLink 6.0 Client for ODBC 64-bit**, and then select the **ODBC Administrator** application.
2. Click the **File DSN** tab. The File DSN tab lists any file data sources in the specified directory.

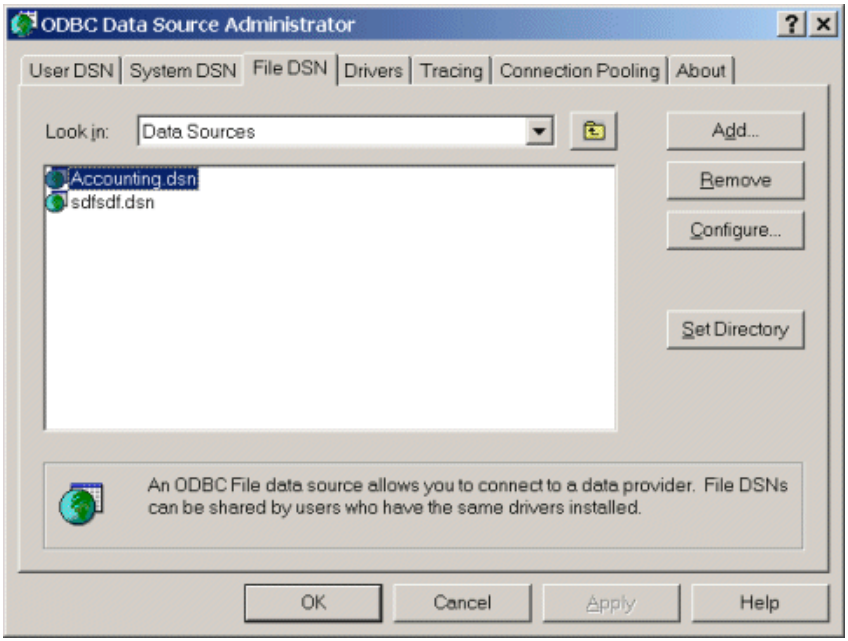

3 To configure a new data source, click the **Add** button. A list of installed drivers appears. Select **DataDirect SequeLink 6.0**; then, perform one of the following actions:

- To configure the file data source to connect directly to a SequeLink Server without retrieving connection information from an LDAP directory, click **OK**. Then, skip to Step 5.
- To configure the file data source to retrieve connection information from an LDAP directory, continue with the next step.

4 Click **Advanced**. The Advanced File DSN Creation Settings window appears.

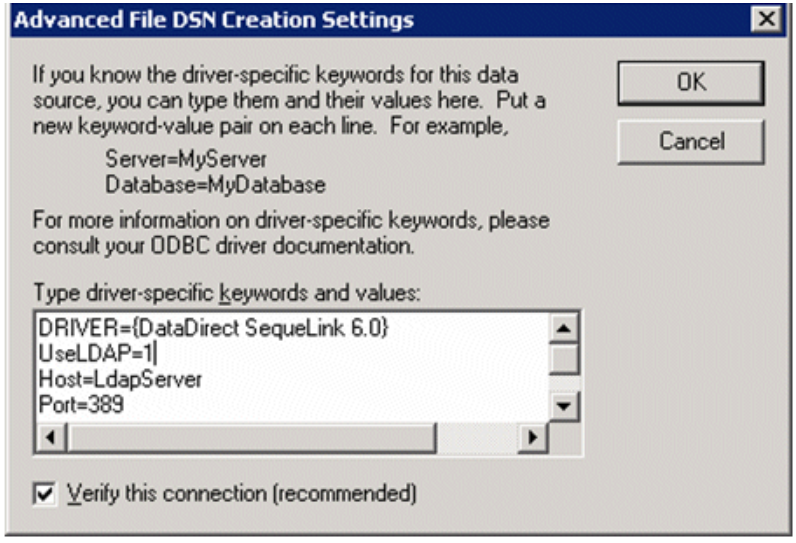

In the Type driver-specific keywords and values scrollable box, type the values required to make the connection, such as the host name and port, as well optional values. For example, type `UseLDAP=1` to configure the ODBC Client to retrieve connection data from an LDAP server, or type `Encrypted=1` when the remote Agent is configured for SSL encryption. Then, click **OK**. You are returned to the list of drivers. Click **Next** and continue with Step 5.

5 The Create New Data Source dialog box appears.

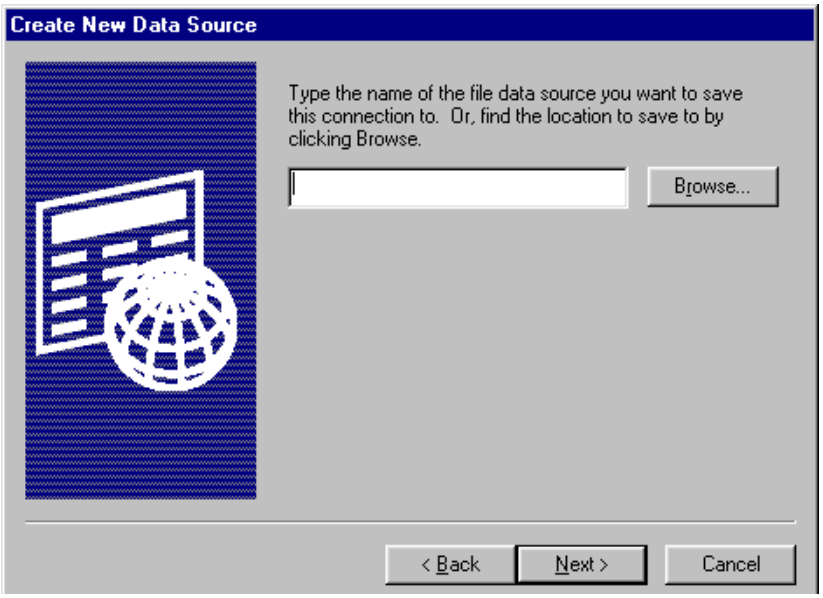

Type the name of the file data source you want to create or click **Browse** to select an existing file data source; then, click **Next**.

6 The Create New Data Source dialog box displays the settings you've configured for this data source.

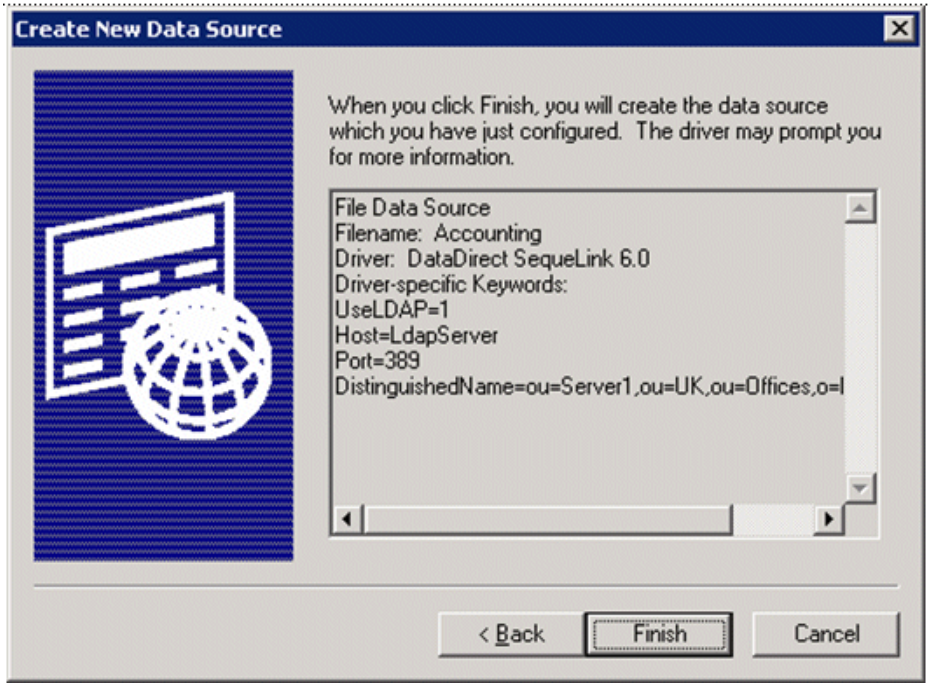

7 Click **Finish** to create the file data source. A series of connection dialogs appear as described in "ODBC Connection Dialogs" on page 183. The file data source will be saved after you enter the correct information in the connection dialog boxes.

ODBC Connection Dialogs

A SequeLink data access connection involves the following stages:

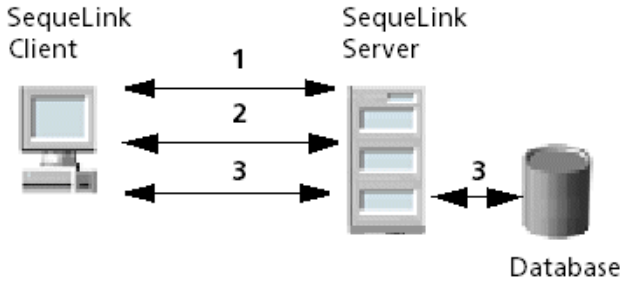

1. A network connection is established.
2. An authentication mechanism is used to establish the identity of the SequeLink Client to the SequeLink Server.
3. Based on information provided by the SequeLink Client application (for example, a database user name and password), a database connection is established.

Stage 1: Establishing a Network Connection

The first stage of the connection process involves establishing a network connection. The dialog box that appears depends on whether the connection has been configured to connect directly to a SequeLink service or to retrieve connection information for the SequeLink service from a centralized LDAP directory.

Connecting Directly to a SequeLink® Service

If the connection has been configured to connect directly to a SequeLink service, the Connect to the SequeLink Server dialog box appears.

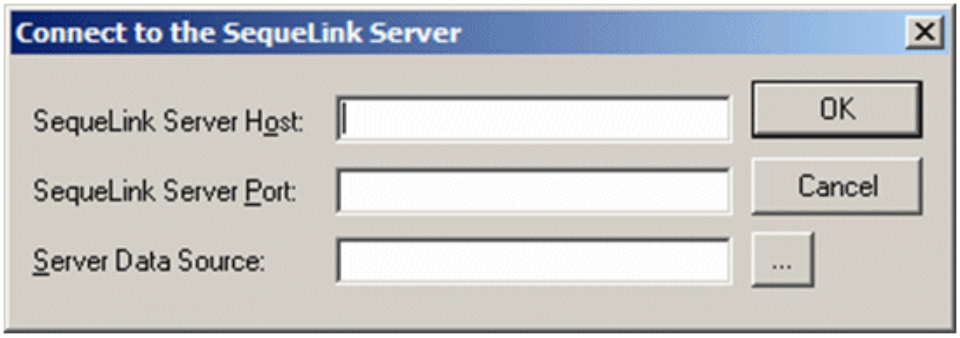

Provide the following information; then, click **OK**.

SequeLink Server Host: Type the TCP/IP host name of the SequeLink service.

SequeLink Server Port: Type the TCP/IP port on which the SequeLink service is listening. A default installation of SequeLink Server uses the port 19996.

Server Data Source: Type the name of a server data source to use for the connection, or select one from the drop-down list. This step is optional. If a server data source is not specified, the default server data source for that service will be used for the connection.

Retrieving Connection Information from an LDAP Directory

If the connection has been configured to connect to an LDAP server to retrieve connection information from an LDAP directory, the Connect to the SequeLink Server dialog box appears.

See "Creating LDAP Entries for SequeLink® Services" on page 430 for information on setting up an LDAP server for SequeLink.

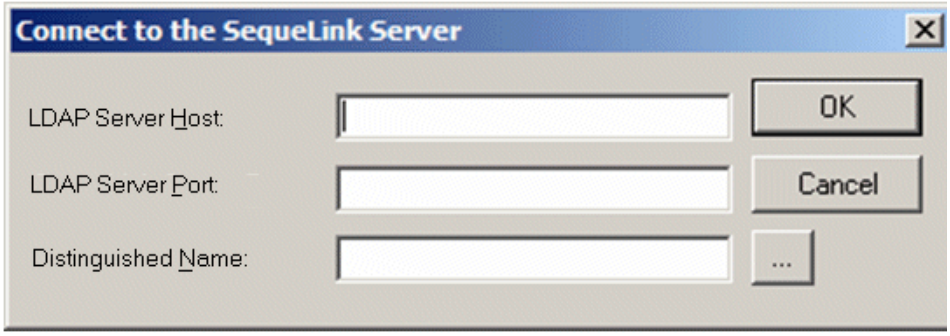

Provide the following information; then, click **OK**.

LDAP Server Host: Type the TCP/IP host name of the LDAP server.

LDAP Server Port: Type the TCP/IP port on which the LDAP server is listening.

Distinguished Name: Type the Distinguished Name (DN) of the LDAP entry.

Stage 2: SequeLink® Server Authentication

The second stage of the connection process involves authentication of the SequeLink Client to the SequeLink Server. The dialog boxes that appear depend on how authentication is configured for the SequeLink service.

- When ServiceAuthMethods=anonymous or ServiceAuthMethods=integrated_nt, no dialog boxes appear.
- When ServiceAuthMethods=OSLogon(*HUID*,*HPWD*) or ServiceAuthMethods=OSLogon(*UID*,*PWD*), the Logon to SequeLink Service dialog box appears.

Provide the following information; then, click **OK**.

Host User Name: Type the host user name.

NOTE: When connecting to a Windows server, you must prefix the host user name with a server name, if authenticating to a local server, or a domain name (for example, SALES\DJONES). If the server name or domain name is omitted, the SequeLink Server will attempt to authenticate the user ID and password with the database account defined for the machine on which the SequeLink Server is running. If this validation fails, the SequeLink Server will attempt to authenticate the user ID and password with the database account defined for the domain of the machine on which the SequeLink Server is running.

Host Password: Type the host password.

- When ServiceAuthMethods=OSLogon(*HUID*,*HPWD*,*NPWD*) or ServiceAuthMethods=OSLogon(*UID*,*PWD*,*NPWD*) and the password is expired, the Password expired. Please specify new password dialog box appears.

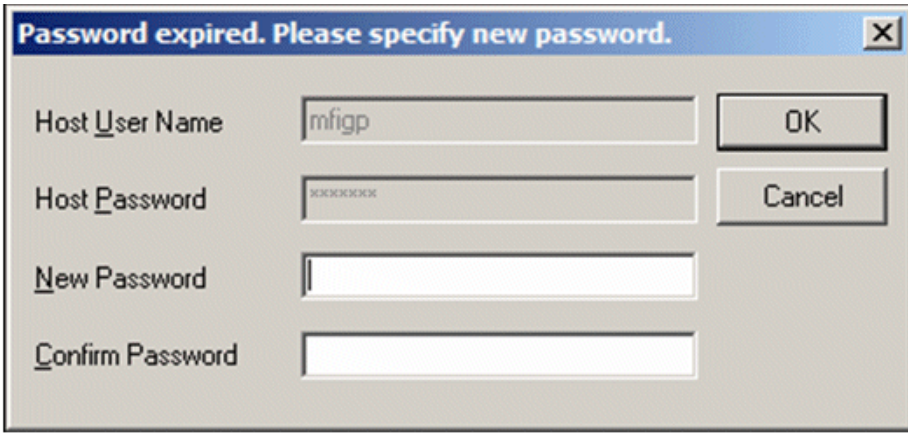

NOTE: If the password is not expired, the Logon to SequeLink Service dialog box appears, prompting only for the host user name and host password.

Provide the following information; then, click **OK**.

Host User Name: Type the host user name.

NOTE: When connecting to a Windows server, you must prefix the host user name with a server name, if authenticating to a local server, or a domain name (for example, SALES\DJONES). If the server name or domain name is omitted, the SequeLink Server will attempt to authenticate the user ID and password with the database account defined for the machine on which the SequeLink Server is running. If this validation fails, the SequeLink Server will attempt to authenticate the user ID and password with the database account defined for the domain of the machine on which the SequeLink Server is running.

Host Password: Type the host password.

New Password: Type the new password to be used by the SequeLink password change mechanism.

Confirm Password: Type the new password again to confirm it.

For more information about configuring authentication, see Chapter 13 "Configuring SequeLink® Security" on page 291.

Stage 3: Data Store Logon

The last stage of the connection process involves logging on the data store. The dialog boxes that appear depend on the data store logon method configured for the SequeLink service:

- When DataSourceLogonMethod=OSIntegrated, no dialog boxes appear.
- When DataSourceLogonMethod=DBMSLogon(*UID*,*PWD*) or DataSourceLogonMethod=DBMSLogon(*DBUID*,*DBPWD*), a data store-specific user name and password are required and the Logon to SequeLink Service dialog box appears.

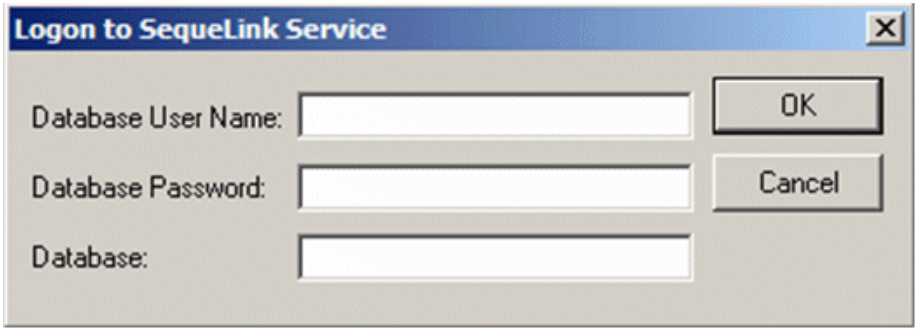

Provide the following information; then, click **OK**.

Database User Name: Type the database logon ID.

Database Password: Type the database password.

Database: Type the name of the database to which you want to connect. This field is disabled when the data store does not recognize the concept of databases.

For more information about configuring authentication, see Chapter 13 "Configuring SequeLink® Security" on page 291 about configuring data store logon methods.

Testing ODBC Connections on Windows

1. On the SequeLink Client, start the ODBC Administrator. To start the ODBC Administrator, select **Start** / **Programs**. From the Programs menu, select **DataDirect SequeLink 6.0 Client for ODBC** or **DataDirect SequeLink 6.0 Client for ODBC 64-bit**, and then select the **ODBC Administrator** application. The ODBC Data Source Administrator window appears listing resident data sources.

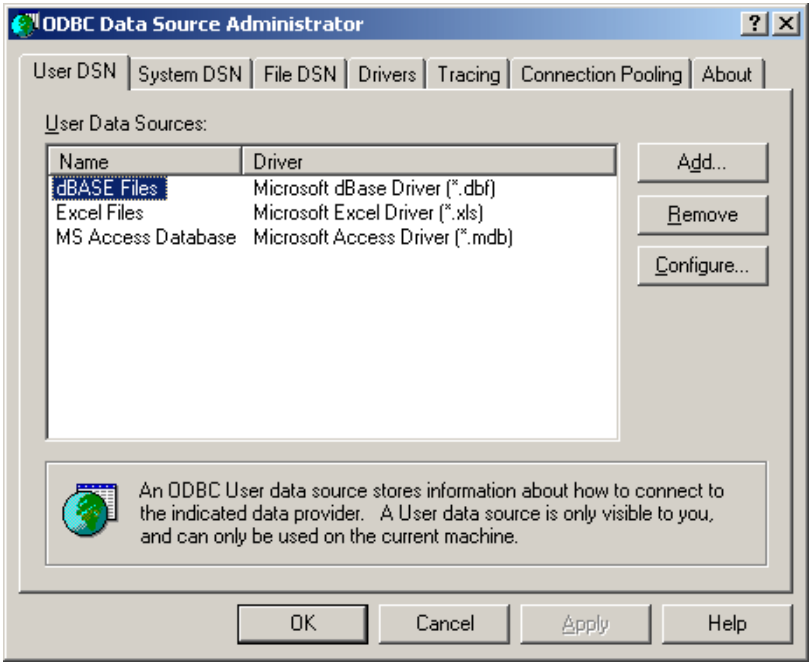

2 Create an ODBC data source as described in "Configuring ODBC User and System Client Data Sources" on page 170 specifying the TCP/IP address and TCP/IP port of the SequeLink service.

3 Click the **Test Connect** button to test the connection. If successful, a dialog appears telling you the connection was successful. You are now ready to start using your ODBC applications with SequeLink.

Configuring ODBC Client Data Sources on Linux and UNIX

For Linux/UNIX, an ODBC Administrator does not exist. This section describes how to configure the odbc.ini file and how to set some required environment variables to use the ODBC Client on Linux/UNIX.

Configuring System Information Files

To configure an ODBC data source for Linux/UNIX, you must edit the system information file, that is, the odbc.ini file (32-bit ODBC Client) or odbc64.ini file (64-bit ODBC Client) using the connect attributes described in "Connecting Using a Connection String" on page 195. The system information file accepts only long names for attributes.

Example: odbc.ini for Solaris

The following code shows an example of an odbc.ini file for a 32-bit ODBC Client installed on a Solaris machine:

```
[ODBC Data Sources]
SALESDB=DataDirect SequeLink 6.0

[SALESDB]
Driver=path_of_installdir/lib/ivslk22.so
Description=DataDirect SequeLink 6.0
Host=
Port=
UseLDAP=0
DistinguishedName=
Encrypted=0
LoadBalancing=0
AlternateServers=
ConnectionRetryCount=0
ConnectionRetryDelay=3

[ODBC]
Trace=0
IANAAppCodePage=4
TraceFile=odbctrace.out
TraceDll=path_of_installdir/lib/odbctrac.so
InstallDir=path_of_installdir
```

where *path_of_installdir* is the path to the ODBC Client installation directory.

Example: odbc64.ini for Solaris

The following code shows an example of the odbc64.ini file for a 64-bit ODBC Client installed on a Solaris machine:

```
[ODBC Data Sources]
SALESDB=DataDirect SequeLink 6.0
```

```
[AccountingDB]
Driver=path_of_installdir/lib64/ivslk22.so
Description=DataDirect SequeLink 6.0
Host=
Port=
UseLDAP=0
DistinguishedName=
Encrypted=0
LoadBalancing=0
AlternateServers=
ConnectionRetryCount=0
ConnectionRetryDelay=3

[ODBC]
Trace=0
IANAAppCodePage=4
TraceFile=odbctrace.out
TraceDll=path_of_installdir/lib64/odbctrac.so
InstallDir=path_of_installdir
```

where *path_of_installdir* is the path to the ODBC Client installation directory.

Setting Environment Variables

You must set several environment variables for the ODBC Client on Linux and UNIX by executing a shell script located in the installation directory.

To execute the shell script:

- If you are using the Bourne or Korn shell, type:
  - `. sqlnk.sh` (32-bit client)
  - `. sqlnk64.sh` (64-bit client)
- If you are using the C shell, type:
  - `source sqlnk.csh` (32-bit client)
  - `source sqlnk64.csh` (64-bit client)

Executing this shell script sets the following environment variables:

ODBCINI	Specifies where the centralized odbc.ini or odbc64.ini file is located.
SQLNK_ODBC_HOME	Specifies the full path of the directory containing the ODBC Client shared libraries.

Executing this shell script also sets the appropriate library search environment variable (LD_LIBRARY_PATH on Solaris and Linux, SHLIB_PATH on HP-UX, or LIBPATH on AIX).

Using a Centralized System Information File

Because Linux and UNIX are multi-user environments, you may want to use a single centralized odbc.ini file controlled by a system administrator. To do this, set the ODBCINI environment variable to point to the fully qualified pathname of the centralized file.

For example:

- In the Bourne or Korn shell, type:

```
ODBCINI=/opt/odbc/system_odbc.ini;export ODBCINI
```

- In the C shell, type:

```
setenv ODBCINI /opt/odbc/system_odbc.ini
```

The odbc.ini file also require an [ODBC] section that includes the InstallDir keyword. The value of the InstallDir keyword must be the path to the directory that contains the /lib and /messages directories.

For example, if you choose the default installation directory for the 32-bit ODBC Client, the following line must be in the [ODBC] section of the odbc.ini file:

```
InstallDir=/usr/slodbc60
```

Connecting Using a Connection String

If you want to use a connection string for connecting to a database, or if your application requires it, you must specify either a DSN (data source name) or a DSN-less connection in the string. The difference is whether you use the DSN= or the DRIVER= keyword in the connection string, as described in the ODBC specification. A DSN connection string tells the driver where to find the default connection information. Optionally, you may specify attribute=value pairs in the connection string to override the default values stored in the data source.

If your application requires a connection string to connect to a data source, you must specify the data source name that tells the driver which data source to use for the default connection information. Optionally, you may specify *attribute=value* pairs in the connection string to override the default values stored in the data source.

The DSN connection string has the form::

```
DSN=data_source_name[;attribute=value[;attribute=value]...]
```

For example, a connection string for SequeLink may look like this:

```
DSN=Accounting;DB=EMP;UID=JOHN;PWD=XYZZY
```

or

```
DSN=Accounting;DB="X:IV;EMP";UID=JOHN;PWD=XYZZY
```

NOTE: If the database name (DB) contains a semicolon (;), you must place the name in quotes, as shown in the preceding example.

The DSN-less connection string specifies a driver instead of a data source. All connection information must be entered in the connection string because there is no data source storing the information.

The DSN-less connection string has the form:

```
DRIVER=[{]driver_name[}][;attribute=value[;attribute=value]
...]
```

NOTE: Empty string is the default value for attributes that use a string value unless otherwise noted.

A DSN-less connection string must provide all necessary connection information:

```
DRIVER=DataDirect SequeLink 6.0;DB=Emp;UID=JOHN;PWD=XYZZY
```

For a list of ODBC connection attributes and their valid values, refer to the *SequeLink Developer's Reference*.

Configuring Connection Failover for the ODBC Client

The ODBC Client can help you make sure that your critical data is available even if the primary database server is unavailable:

- *Connection failover* allows an application to connect to an alternate, or backup, database server if the primary database server is unavailable, for example, because of a hardware failure or traffic overload. See "Specifying Alternate Servers" on page 197 for more information.
- *Connection retry* defines the number of times the driver attempts to connect to the primary server and, if configured, alternate database servers after the initial unsuccessful connection attempt. See "Using Connection Retry" on page 201 for more information.
- *Client load balancing* helps distribute new connections in your environment so that no one server is overwhelmed with connection requests. "Using Client Load Balancing" on page 201 for more information.

Specifying Alternate Servers

To configure connection failover, you specify a list of alternate database servers that are tried at connection time if the primary SequeLink server is not accepting connections. To do this, use the Alternate Servers (AlternateServers) connection option. Connection attempts continue until a connection is successfully established or until all the database servers in the list have been tried once (the default).

On Windows, you can configure a data source to use connection failover on the Failover tab of the driver's Setup dialog box. See

"Configuring ODBC Client Data Sources on Windows" on page 170 for details.

On Linux and UNIX, you can configure a data source to use connection failover by modifying your system information file (odbc.ini). See "odbc.ini File Example" on page 199 and "odbc64.ini File Example" on page 200 for details.

On Linux, UNIX, and Windows, you can use a connection string to direct the driver to use connection failover. See "Connection String Example" for details.

Connection String Example

The following connection string configures the ODBC Client connected to a Linux, UNIX, or Windows server to use connection failover in conjunction with all of its optional features—load balancing, connection retry, and connection retry delay.

```
DSN=MyODBCDSN;AlternateServers=(Host=server1:Port=19996:ServerDataSource=
SDSN1,Host=server2:Port=19996:ServerDataSource=SDSN2,Host=server3:Port=
19996:ServerDataSource=SDSN3);ConnectionRetryCount=4;ConnectionRetryDelay=
5;LoadBalancing=1
```

Specifically, if a successful connection is not established on the ODBC Client's first pass through the list of SequeLink servers, this connection string configures the driver to use two alternate servers as connection failover servers, to attempt to connect four additional times if the initial attempt fails, to wait five seconds between attempts, and to try the primary and alternate servers in a random order.

The additional connection information required for the alternate servers is specified in the SequeLink Server data source MyODBCDSN.

odbc.ini File Example

To configure the 32-bit ODBC Client installed on a Solaris machine to use connection failover in conjunction with some of its optional features in your odbc.ini file, you could set the following connection string attributes:

```
[ODBC Data Sources]
SALESDB=DataDirect SequeLink 6.0

[SALESDB]
Driver=path_of_installdir/lib/ivslk22.so
Description=DataDirect SequeLink 6.0
Host=
Port=
UseLDAP=0
DistinguishedName=
AlternateServers=(Host=server1:Port=19996:ServerDataSource=SDSN1,Host=
server2:Port=19996:ServerDataSource=SDSN2,Host=server3:Port=
19996:ServerDataSource=SDSN3)
...
ConnectionRetryCount=4
ConnectionRetryDelay=5
...
LoadBalancing=0
...
```

Specifically, this odbc.ini configuration tells the ODBC Client to use two alternate servers as connection failover servers, to attempt to connect four additional times if the initial attempt fails, to wait five seconds between attempts, and to try the primary and alternate servers in sequential order (do not use load balancing).

odbc64.ini File Example

To configure the 64-bit ODBC Client installed on a Solaris machine to use connection failover in conjunction with some of its optional features in your odbc.ini file, you could set the following connection string attributes:

```
[ODBC Data Sources]
AccountingDB=DataDirect SequeLink 6.0

[AccountingDB]
Driver=path_of_installdir/lib64/ivslk22.so
Description=DataDirect SequeLink 6.0
Host=
Port=
UseLDAP=0
DistinguishedName=

AlternateServers=(Host=server1:Port=19996:ServerDataSource=SDSN1,Host=
server2:Port=19996:ServerDataSource=SDSN2,Host=server3:Port=
19996:ServerDataSource=SDSN3)
...
ConnectionRetryCount=4
ConnectionRetryDelay=5
...
LoadBalancing=1
...
```

This odbc64.ini configuration tells the ODBC Client to use two alternate servers as connection failover servers, to attempt to connect four additional times if the initial attempt fails, to wait five seconds between attempts, and to try the primary and alternate servers in a random order (use load balancing).

Using Connection Retry

Connection retry defines the number of times the Client attempts to connect to the primary server and, if configured, alternate database servers after the initial unsuccessful connection attempt. Connection retry can be an important strategy for system recovery. For example, suppose you have a power failure in which both the client and the server fails. When the power is restored and all computers are restarted, the client may be ready to attempt a connection before the server has completed its startup routines. If connection retry is enabled, the client application can continue to retry the connection until a connection is successfully accepted by the server.

Connection retry can be used in environments that have only one server or can be used as a complementary feature with connection failover in environments with multiple servers.

Using connection options, you can specify the number of times the driver attempts to connect and the time in seconds between connection attempts. For details on configuring connection retry, see "Connecting Using a Connection String" on page 195.

Using Client Load Balancing

Client load balancing helps distribute new connections in your environment so that no one server is overwhelmed with connection requests. When client load balancing is enabled, the order in which primary and alternate database servers are tried is random.

When Connection Retry is also enabled, the ODBC Client tries to connect to the primary SequeLink server and alternate SequeLink servers in a random order until a successful connection is established. If the connection attempt fails, the driver again randomly selects from the list of servers until all SequeLink servers in the list have been tried or a connection is successfully established.

Connection Failover Attributes

Table 8-1 summarizes the connection attributes that control how connection failover works with the ODBC Client. Refer to the *SequeLink Developer's Reference* for details about configuring each attribute.

Table 8-1. Summary: Connection Failover Attributes for the ODBC Client

Attribute	Characteristic
AlternateServers	List of alternate database servers. An IP address or server name and a port number are required for each server. The ServerDataSource connection attribute is optional.
ConnectionRetryCount	Number of times the driver retries the primary database server, and if specified, alternate servers until a successful connection is established. The initial default is 0.
ConnectionRetryDelay	Wait interval, in seconds, between connection retry attempts when the ConnectionRetryCount attribute is set to a positive integer. The initial default is 3.
LoadBalancing	Sets whether the driver will use client load balancing in its attempts to connect to the list of database servers (primary and alternate). If client load balancing is enabled, the driver uses a random pattern instead of a sequential pattern in its attempts to connect. The initial default is 0 (client load balancing is not used).

Refer to the *SequeLink Developer's Reference* for overviews of connection failover and client load balancing.

Importing and Exporting ODBC Client Data Sources

The SequeLink Data Source SyncTool allows you to export ODBC client data source definitions to data source files and distribute them to multiple end users. The SequeLink Data Source SyncTool provides two user implementations, one for the SequeLink administrator and another for the end user:

- The SequeLink *for* ODBC Data Source SyncTool Administrator is used by the SequeLink administrator to create data source files. It can import and export data sources. This tool should be made available to the SequeLink administrator only.
- The SequeLink *for* ODBC Data Source SyncTool is used by the end user and can import data sources only. It should be installed on every SequeLink Client *for* ODBC.

In addition, you can create a customized, installable image of SequeLink Client *for* ODBC with predefined, site-specific settings, including data source files created with the SequeLink Data Source SyncTool. This customized, installable image is called a *Quick Install image*. For more information about creating Quick Install images, refer to the *SequeLink Installation Guide*.

The window title bar of the SequeLink Data Source SyncTool indicates whether you, or the end user, is performing an export or an import operation. Also, context-sensitive online help is available by clicking **?** on the title bar; then, click the area about which you want more information.

Exporting ODBC Client Data Sources

1. From the SequeLink program manager group, double-click the **ODBC Data Source SyncTool Administrator** icon. The SequeLink *for* ODBC Data Source SyncTool Administrator Welcome window appears.
2. Select the **Manage Data Sources Files** option; then, click **Next**.
3. Select a data source file from the Filename list box, or click **Browse** to find a data source file not listed. The default extension for a data source file is .DSF.

   To create a new data source file, click **New**.
4. Select whether you want to export User or System data sources to the data source file you selected; then click **Next**.
5. Select the data sources you want to export to the data source file.

   NOTE: You cannot export grayed-out data sources, which are data sources that are configured for a previous incompatible version of the ODBC driver.

6 Using the following symbols, verify that the appropriate actions will be performed on the data sources in the data source file; then, click **Next**.

- The data source will remain unchanged.
- The data source will be added to the data source file.
- The data source will be deleted from the data source file.
- The data source will be updated in the data source file.

7 Type a description for the data source file; then, click **Next**. This description will appear when the end user selects this file for importing.

8 Select the mode the end user will use to import these data sources; then, click **Next**.

- *Interactive mode* allows the user to select which data sources will be imported. This mode is not supported by the Quick Install feature; the Quick Install feature supports only data source files created with the Merge or Overwrite options. For instructions on creating Quick Install images, refer to the *SequeLink Installation Guide*.
- *Merge mode* adds or updates all the data sources in the data source file without deleting other data sources.
- *Overwrite mode* adds or updates the data sources in the data source file and deletes any other data sources configured for the ODBC driver.

9 Select the option that will determine how the end user will be able to import the data sources you exported to the data source file; then, click **Next**.

- *Suggest SequeLink User DSN*. When imported, the SequeLink *for* ODBC Data Source SyncTool will suggest to the end user that these data sources be imported as User data sources, but will allow them to be imported as User or System data sources.
- *Suggest SequeLink System DSN*. When imported, the SequeLink *for* ODBC Data Source SyncTool will suggest to the end user that these data sources be imported as System data sources, but will allow them to be imported as User or System data sources.
- *Force SequeLink User DSN*. When imported, the SequeLink *for* ODBC Data Source SyncTool will allow these data sources to be imported as User data sources only.
- *Force SequeLink System DSN*. When imported, the SequeLink *for* ODBC Data Source SyncTool will allow these data sources to be imported as System data sources only.

10 Click **Finish** to quit.

Importing ODBC Client Data Sources

The SequeLink administrator and end user use a different implementation of the SequeLink *for* ODBC Data Source SyncTool to import ODBC data source definitions.

To import ODBC client data sources:

1 From the SequeLink program manager group, double-click the appropriate ODBC SyncTool icon. The Welcome window appears.

2 Select the **Import** option, and click **Next**.

NOTE: If using the SequeLink *for* ODBC Data Source SyncTool Administrator, select the **Import Data Sources** option; then, click **Next**.

3 Select a data source file from the Filename list box, or click **Browse** to find a data source file not listed. The default extension for data source files is .DSF.

4 Indicate whether you want to import the data sources in the data source file you just selected as User or System data sources; then, click **Next**.

5 Verify that the appropriate actions will be performed on the data sources on your local machine; then, click **Next**. Depending on the import mode that was set when the data source file was exported, you may see the following symbols:

- The data source will remain unchanged.
- The data source will be added to your local machine.
- The data source will be deleted from your local machine.
- The data source will be updated to your local machine.

NOTE: Grayed-out data sources are data sources that are configured for a previous incompatible version of the ODBC driver; these data sources will remain unchanged unless you update them in Interactive mode with a data source configured for the current version of the ODBC driver.

6 Click **Finish** to quit.

Authentication

The ODBC driver supports the following methods of authentication:

- **Anonymous**. The SequeLink Server accepts connections from the SequeLink Client without verifying the client's identity.
- **Operating system user ID** and **password**. The SequeLink Server verifies the identity of the SequeLink Client using a user ID and password that must be valid for the platform on which the SequeLink Server is running. If verified, the server accepts the user ID as the identity of the client and permits the connection.
- **Kerberos**. Kerberos authentication uses Kerberos, a trusted third-party authentication service, to verify user identities. Kerberos authentication can take advantage of the user name and password maintained by the operating system to authenticate users to the database. This method requires knowledge of how to configure your Kerberos environment.
- **Integrated NT**. This option is supported for connections between SequeLink Server for Windows servers and ODBC Clients, ADO Clients, and .NET Clients on Windows only. The SequeLink Server verifies the identity of the SequeLink Client using the client's Windows network logon credentials instead of a Windows user ID and password.

For details on configuring authentication for SequeLink, see "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.

Using SSL Encryption

If your SequeLink environment requires greater data privacy than that provided by fixed-key DES, fixed-key 3DES, or byteswap, you can use the Secure Socket Layer (SSL) to encrypt data exchanged between the ODBC Client and the SequeLink Server.

SequeLink supports the use of anonymous ciphers. Anonymous ciphers allow the SSL connection to succeed without proper authentication of the peer by using the DH algorithm. The ODBC Client supports the following cryptographic strong SSL cipher suites:

- TLS_DH_anon_WITH_AES_128_CBC_SHA
- TLS_DH_anon_WITH_AES_256_CBC_SHA

For details on configuring SSL for SequeLink, see "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.

Unicode and Code Page Support

The ODBC driver fully supports the SQL-W functions and Unicode arguments, for example, SQLConnectW. This support enables faster processing of wide-characters and allows binding of the SQL_C_WCHAR output type.

On Windows, SQL-W routines map to UTF-16. On Linux and UNIX, SQL-W routines map to UTF-8 or UTF-16. How the database data types are mapped depends on the database and a number of configuration options. See Appendix F "Internationalization, Localization, and Unicode" on page 597 for more information about internationalization and localization and about important differences in developing applications on Linux and UNIX.

See Appendix D "SequeLink® Service Attributes" on page 491 for information about the SequeLink service attributes that affect configuration.

Refer to the *SequeLink Developer's Reference* for information about data type mappings, and additional information about developing applications on Linux and UNIX.

9 Configuring the ADO Client

This chapter describes the tasks you may need to perform to configure and manage the SequeLink Client *for* ADO (the ADO Client).

About ADO Connections

You can open an ADO connection to a SequeLink service by specifying a configuring an ADO client data source. This chapter explains how to connect to a SequeLink service using client data source.

Using the DataDirect Technologies Configuration Manager

To create and configure data sources for the ADO Client, you use the DataDirect Technologies Configuration Manager.

To start the Configuration Manager, select **Start / Programs**, and select **DataDirect SequeLink 6.0 Client for ADO**. Then, select the **DataDirect Configuration Manager** application.

The Configuration Manager window is divided into two panes. As Figure 9-1 shows, the left pane displays a folder containing defined ADO data sources. When you select a data source, the right pane displays the properties for the selected data source.

Figure 9-1. DataDirect Technologies Configuration Manager

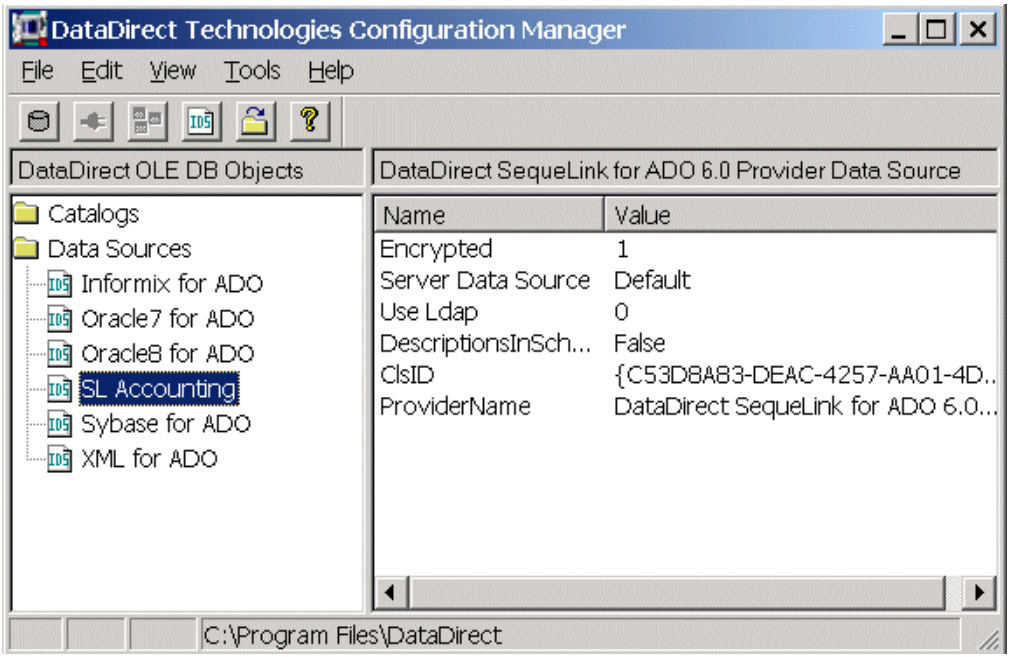

Double-click the **Data Sources** folder to display any existing ADO data sources. The Configuration Manager displays the ADO data sources contained in the current directory, which is shown in the status bar at the bottom of the Configuration Manager. The first time you start the Configuration Manager, the current directory defaults to the \Program Files\DataDirect\slado60 directory.

Working with the DataDirect Technologies Configuration Manager

Table 9-1 summarizes the parts and functions of the Configuration Manager that you use with ADO data sources.

NOTE: Options that are not supported by the ADO provider are disabled in the toolbar and are omitted from this description.

Table 9-1. DataDirect Technologies Configuration Manager: Parts and Functions for ADO Data Sources

Use this element...		To do this...
Toolbar		Create new data sources
		Change the current directory
		View online help
Menu Bar	File	■ Create a new data source
		■ Exit from the DataDirect Configuration Manager
	Edit	■ Delete a data source
		■ Rename a data source
		■ Modify a data source
	View	■ View or hide the toolbar and status bar
		■ Refresh the Configuration Manager

Shortcut Tip: Right-clicking an item in the left pane displays a pop-up menu that allows you to perform the same actions that are available from the toolbar and menu bar.

Table 9-1. DataDirect Technologies Configuration Manager: Parts and Functions for ADO Data Sources *(cont.)*

Use this element...		To do this...
	Tools	■ Change the directory in which to look for data sources ■ Define a Template data source directory ■ Define a Master data source directory
	Help	View online help.
Vertical splitter bar		Adjust the size of the left and right panes.
Status bar		■ Show the current keyboard state, including when NUM LOCK, SCROLL LOCK, and CAPS LOCK are turned on ■ Show the current directory
Shortcut Tip: Right-clicking an item in the left pane displays a pop-up menu that allows you to perform the same actions that are available from the toolbar and menu bar.		

Displaying Data Source Properties

1. Start the Configuration Manager. To start the Configuration Manager, select **Start / Programs**, and select **DataDirect SequeLink 6.0 Client for ADO**. Then, select the **DataDirect Configuration Manager** application.
2. Double-click the **Data Sources** folder to display any existing ADO data sources.

3 Highlight a data source in the list. The properties of the data source display in the right pane. For example, the following figure shows the properties of an ADO data source named SL Accounting displayed in the right pane.

Figure 9-2. DataDirect Technologies Configuration Manager: Displaying Data Source Properties

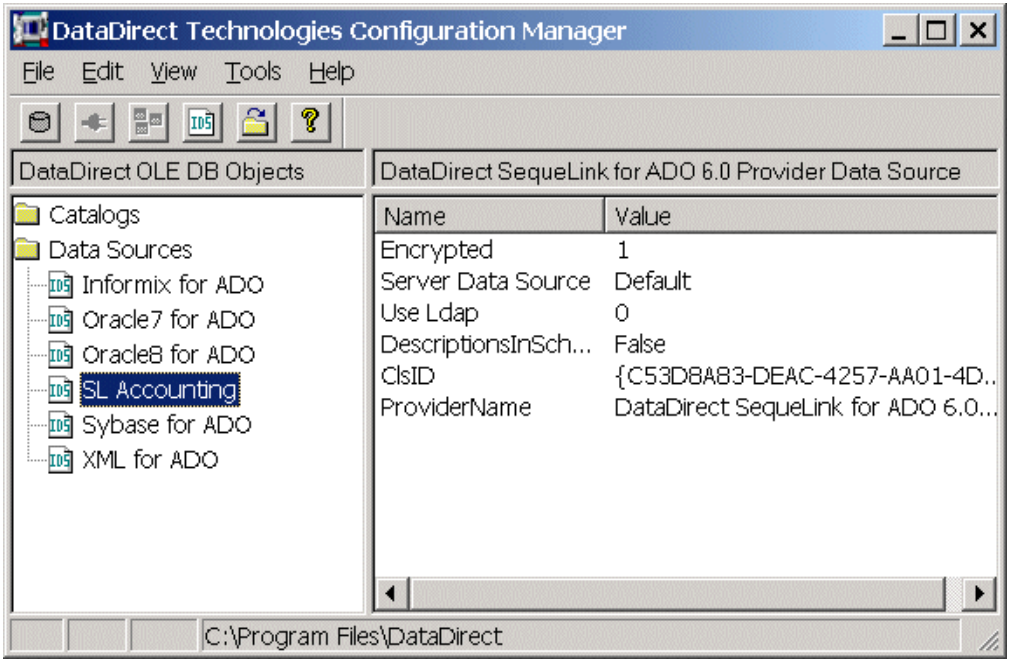

You can right-click a data source in the left pane to display a pop-up menu. The pop-up menu offers the same actions for the item that are available from the Edit menu.

To display a setup window for an existing data source, double-click an ADO data source in the Data Sources folder.

To create a new data source, highlight the **Data Sources** folder; then, select **File / New / Data Source** from the menu bar.

Configuring ADO Client Data Sources

The following sections provide instructions for configuring ADO client data sources:

Creating an ADO Client Data Source

1 Start the DataDirect Configuration Manager. To start the Configuration Manager, select **Start** / **Programs**, and select **DataDirect SequeLink 6.0 Client for ADO**. Then, select the **DataDirect Configuration Manager** application.

2 Select **File** / **New** / **Data Source** from the menu bar. The New Data Source window appears.

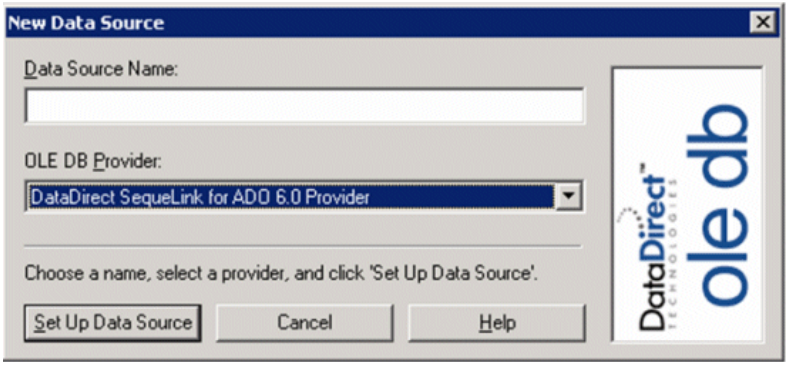

3 Type a name for the data source. All data sources located in the same directory must have unique names. If the name has already been used for another data source, you are prompted to enter a different name.

4 In the DataDirect OLE DB Providers drop-down list, select **DataDirect SequeLink for ADO 6.0 Provider**.

5 Click the **Set Up Data Source** button. The DataDirect SequeLink for ADO 6.0 Provider Setup window appears.

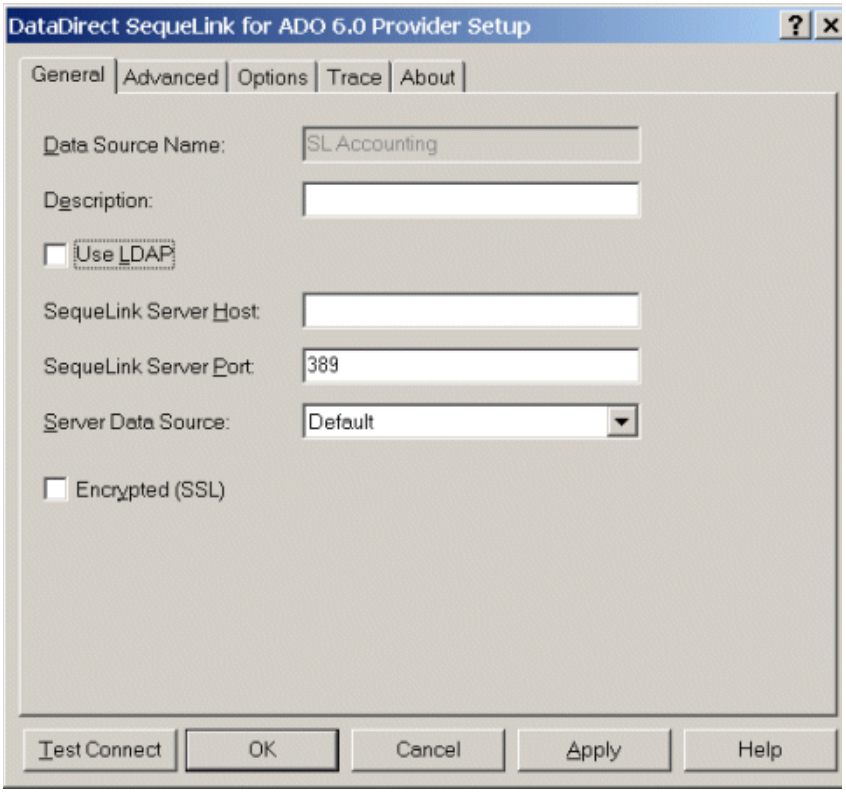

NOTE: The General tab displays only fields that are required for creating a data source. The fields on all other tabs are optional, unless noted otherwise.

Provide the following information:

Data Source Name: This is a read-only field that uniquely identifies this ADO data source configuration. Examples include “Accounting” or “SequeLink to Oracle Data”.

Description: Optionally, type a description of the data source. For example, “My Accounting Database” or “Accounting Data in Oracle”.

SequeLink Server Host: Type the TCP/IP host name of the SequeLink service to which you want the ADO Client to connect. This field is available only if the Use LDAP check box is **not** selected.

SequeLink Server Port: Type the TCP/IP port the SequeLink service is listening on for incoming connection requests. The port you specify must be the same as the one that was specified for the SequeLink service when the SequeLink Server was installed; the default is 19996. This field is available only if the Use LDAP check box is **not** selected.

Server Data Source: Type the name of a server data source configured for the SequeLink service to use for the connection, or select one from the drop-down list. This field is optional. If a server data source is not specified, the default server data source for that SequeLink service will be used for the connection. This field is available only if the Use LDAP check box is **not** selected.

Use LDAP: To configure the ADO Client to retrieve connection information from an LDAP directory, select the **Use LDAP** check box. The fields change on the lower half of the screen to accommodate the information that is required to query an LDAP server for connection information. Provide the following information:

LDAP Server Host: Type the TCP/IP host name of the LDAP server.

LDAP Server Port: Type the TCP/IP port on which the LDAP server is listening for incoming connection requests. If unspecified, the ADO Client will use the default LDAP port 389.

Distinguished Name (DN): Type an identifier that uniquely identifies the LDAP entry where connection information is stored.

See "Retrieving Connection Information from LDAP Directories" on page 428 for more information about retrieving connection information from LDAP directories.

Encrypted (SSL): If the remote SequeLink service is configured for SSL, select the check box. This check box must be selected when connecting to a SequeLink service enabled for SSL.

When the check box is cleared (the initial default), communication between the SequeLink Client and SequeLink Server is not encrypted with SSL.

Configuration of encryption is performed on the SequeLink Server. For more information, see "Data Encryption" on page 299.

NOTES:

- Data encryption with SSL is not supported for LDAP Servers. The Use LDAP and the Encrypted (SSL) check boxes are mutually exclusive.
- SSL encryption is not supported for SequeLink Server for DB2 on z/OS. To support SSL in a DB2 for z/OS environment, use the SequeLink Proxy Server (see Chapter 14 "Configuring the SequeLink® Proxy Server" on page 351).

6 Optionally, click the **Trace** tab to enable tracing options. Specify values on the Trace tab, then, click **Apply**.

NOTE: Settings on this tab apply to all data sources for the ADO data provider. You cannot set the trace options programatically.

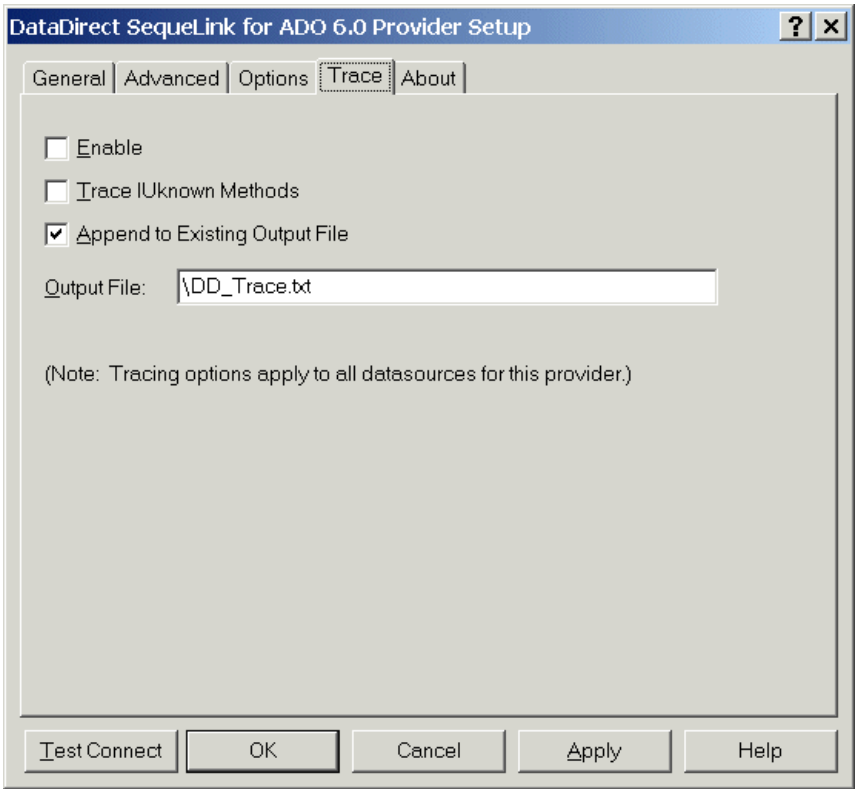

Enable: Select this check box to enable tracing support. By default, the check box is not selected.

Trace IUnknown Methods: Select this check box to enable tracing support of IUnknown methods. By default, the check box is not selected.

Append to Existing Output File: Select this check box to append tracing results to an output file. By default, the check box is selected.

Output File: Type the name of the file to which tracing results will be appended. This file contains the tracing results of all data sources for the ADO data provider.

7 Optionally, click the **Options** tab to add connection attributes (see "Connecting with a Provider String" on page 237 for the values that can be entered on this tab). Values in the Provider Options field are separated by semicolons.

For example, the following string sets values for the Alternate Servers, Connection Retry Count, and Connection Retry Delay connection failover options:

```
Alternate Servers=(Host=server2:Port=19996,Host=
server3:Port=19996,Host=server4:Port=19996);Connection
Retry Count=2;Connection Retry Delay=3
```

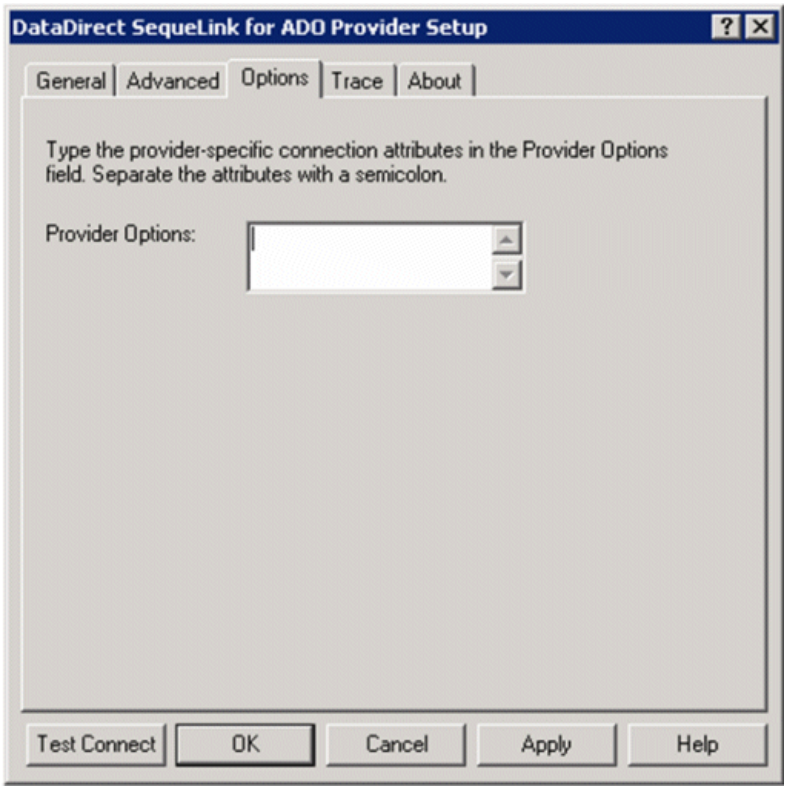

8 At any point during the configuration process, you can click **Test Connect** to attempt to connect to the data source using the connection properties specified in the provider Setup dialog box. A logon dialog box appears; see "Connecting to an ADO Data Source" on page 229 for details.

   Note that the information you enter in the logon dialog box during a test connect is not saved.

   - If the data provider can connect, it releases the connection and displays a Connection Established message. Click **OK**.
   - If the data provider cannot connect because of an improper environment or incorrect connection value, it displays an appropriate error message. Click **OK**.

   NOTE: If you are configuring alternate servers for use with the connection failover feature, be aware that the Test Connect button tests only the primary server, not the alternate servers.

9 Click **OK** or **Cancel**. If you click **OK**, the values you have specified become the defaults when you connect to the data source. You can change these defaults by using this procedure to reconfigure your data source. You can override these defaults by connecting to the data source using a connection string with alternate values.

NOTE: All data sources are saved to the current directory displayed in the Configuration Manager. See "Changing Data Source Directories" on page 225 for instructions on changing the current directory.

Modifying an ADO Client Data Source

To modify the properties of a data source, double-click the data source in the Data Sources folder of the Configuration Manager to display the SequeLink for ADO Provider Setup window. See "Creating an ADO Client Data Source" on page 216 for a description of the fields you can change.

Renaming an ADO Client Data Source

You can rename data sources. You cannot rename or delete the Data Sources folder.

To rename an ADO provider data source:

1. Start the Configuration Manager. To start the Configuration Manager, select **Start** / **Programs**, and select **DataDirect SequeLink 6.0 Client for ADO**. Then, select the **DataDirect Configuration Manager** application.
2. Select the data source you want to rename.
3. Select **Edit** / **Rename**. The data source name becomes an editable field.
4. Type the new name of the data source and press ENTER.

Deleting an ADO Client Data Source

1. Start the Configuration Manager. To start the Configuration Manager, select **Start** / **Programs**, and select **DataDirect SequeLink 6.0 Client for ADO**. Then, select the **DataDirect Configuration Manager** application.
2. Select the data source you want to delete.
3. Select **Edit** / **Delete**.
4. A window appears prompting you to confirm the deletion. Click **Yes** to delete the selected data source.

After you change the current directory, the left pane of the Configuration Manager is automatically refreshed to display the data sources in the new directory. The current directory remains active until you change it again. Any data sources you create are saved to the current directory.

Copying an ADO Client Data Source

Copying a data source can make it easier for you to configure new data sources that use the same properties as existing data sources. When you copy a data source, the copied data source retains all the properties of the original data source. After copying, you can modify the properties of the data source as needed.

To copy a data source:

1. In Windows Explorer, navigate to the directory that contains the data source you want to copy. All ADO provider data sources use .IDS as their file extension. For example, if the data source name appears as TEST in the Configuration Manager, the name of the data source file is TEST.IDS.

   NOTE: The directory location of a data source displayed in the Configuration Manager appears in the status bar at the bottom of the Configuration Manager.

2. Copy the data source to the Windows Explorer clipboard; then, perform one of the following actions:

   - To copy to a different directory, navigate to the directory you want to copy to and paste the data source in that new directory. You can use the same data source name.
   - To copy to the same directory, paste the data source; then, rename the data source to a *unique* name.

3 To display the new data source in the Configuration Manager, perform one of the following actions:

- If you copied the data source to a different directory, make that directory the current directory in the Configuration Manager by selecting **Tools** / **Options** / **Main Data Source Directory**. The new data source appears in the Data Sources folder.
- If you copied the data source to the same directory and renamed the data source, select **View** / **Refresh** in the Configuration Manager. The new data source appears in the Data Sources folder.

Changing Data Source Directories

The Configuration Manager displays the ADO data sources contained in the current directory, which is displayed in the status bar at the bottom of the Configuration Manager. The first time you start the Configuration Manager, the current directory defaults to the ADO Client installation directory.

To change the current directory:

1 Click the **Change main Data Source directory** button on the tool bar.

2 Type the name of the new directory in the Current Directory field, or, click the **Browse** button to select a different directory.

3 Click **OK**.

After you change the current directory, the left pane of the Configuration Manager is automatically refreshed to display the data sources in the new directory. The current directory remains active until you change it again. Any data sources you create are saved to the current directory.

Defining Default Setup Options

The Configuration Manager supports configurable default setup options and override options through the use of a template data source file and a master data source file, respectively.

A template data source file is used by the Configuration Manager to populate values in the fields of the Setup dialog box when a user creates a new data source. By creating a template data source file, you can define the default setup options (default values for newly created data sources). The user can change these default values when setting up a new data source.

A master data source file is used to provide global connection options. The options set in the master data source file override connection options set any other way (for example, by the data source specified by an application or a connection string) when an application is connecting to the database.

Creating a Template Data Source File

You can define template data source files to simplify the creation of data source files. A template data source file allows you to define the default setup options for SequeLink data providers. The Configuration Manager supplies these values in the Setup dialog box when a user creates a new data source. The user can change these default values when setting up a new data source.

To create a template data source file:

1. Create a directory in which to store the template data source file.

   IMPORTANT: The template data source directory cannot be the same as the directory for other data sources.

2 In the Configuration Manager window, select **Tools / Options / Template Data Source Directory**. Specify the directory that you created in Step 1; then, click **OK**.

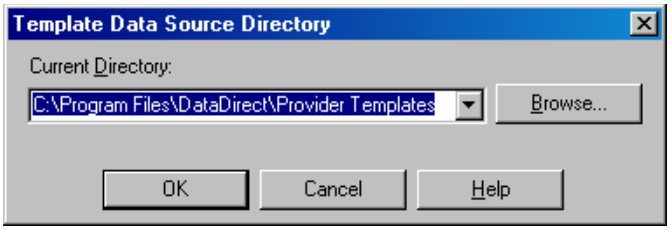

3 Select **Tools / Options / Main Data Source Directory**. Specify the template directory; then, click **OK**. This sets the template directory as the location in which to create the template data source file.

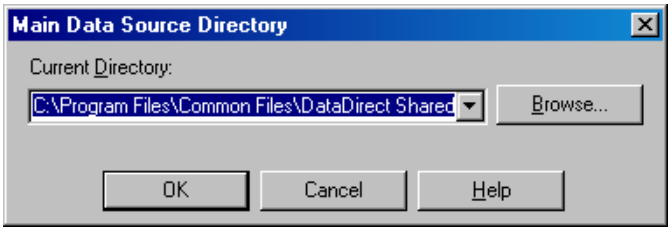

4 Create a data source, defining the values that will be most commonly used. This will be your template data source file for the specified data provider.

5 Select **Tools / Options / Main Data Source Directory**. Specify the directory that contains your data sources; then, click **OK**.

Creating a Master Data Source File

You can define a master data source file that overrides connection options set any other way. This allows you to control the way that users connect to the database.

During connection, the Main data source directory is checked for a data source, and connection values are retrieved. If a Master data source directory exists, it is then checked for the same data

source. The connection settings for user data sources will be overridden by the master data source file.

To create a master data source file:

1 Create a directory in which to store the master data source file.

IMPORTANT: The master data source directory cannot be the same as the directory for template data sources or any other data provider data sources.

2 In the Configuration Manager window, select **Tools** / **Options** / **Master Data Source Directory** and specify the directory that you created in Step 1. The master data source file will be used at connection time.

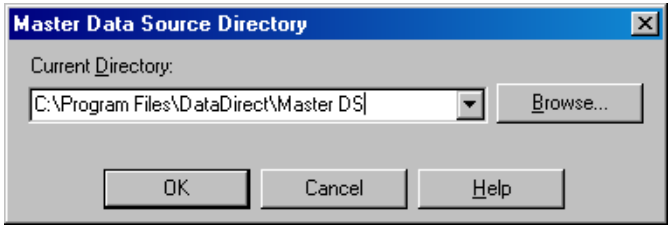

3 Select **Tools** / **Options** / **Main Data Source Directory**. Specify the master data source directory; then, click **OK**. This sets the master data source directory as the location in which to create the master data source file.

4 Create one or more data sources. The data sources in this directory will be your master data source files for the specified data providers.

5 Select **Tools** / **Options** / **Main Data Source Directory**. Specify the directory that contains your data sources; then, click **OK**.

Connecting to an ADO Data Source

You can connect to a data source using a Connection window, or using a provider string. For information about connecting using an ADO provider string, refer to the *SequeLink Developer's Reference*.

Testing ADO Connections

The ADO provider opens a Connection window when you perform either of the following actions:

- You request a connection to an ADO provider from within your data consumer, and your data consumer requests the ADO provider to prompt for missing connection parameters.
- You click **Test Connect** in an ADO provider setup window to test the connection to a data source you have set up.

See "ADO Connection Dialogs" on page 230 for more information about ADO connection dialogs that may appear.

ADO Connection Dialogs

When your data consumer requests the ADO provider to prompt for missing connection parameters and an ADO data source has not been specified, the DataDirect SequeLink for ADO Provider Connection dialog box appears.

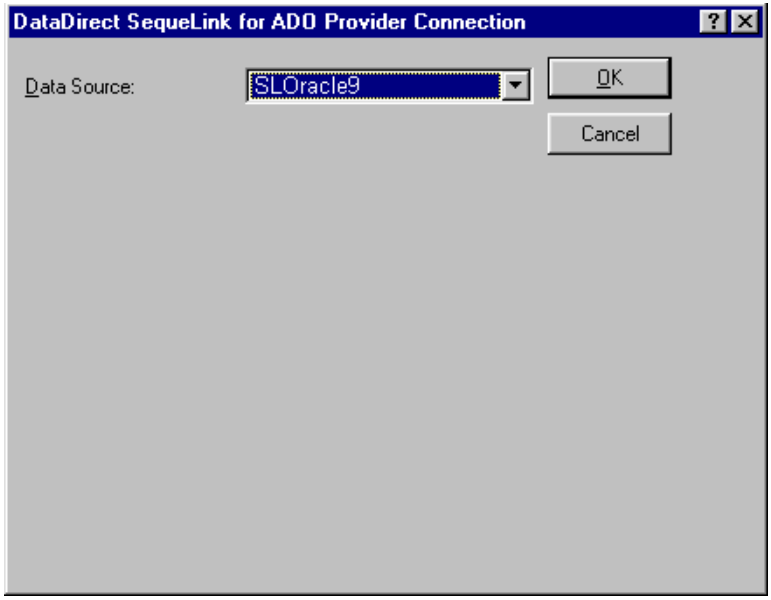

Select the data source that you want to use from the drop-down list. If you do not want to specify a data source name, select **None** from the drop-down list. In some cases, the data source name may be supplied automatically. Then, click **OK**.

The other connection dialogs that may appear involve prompting for information required to make a SequeLink data access connection.

A SequeLink data access connection involves the following stages:

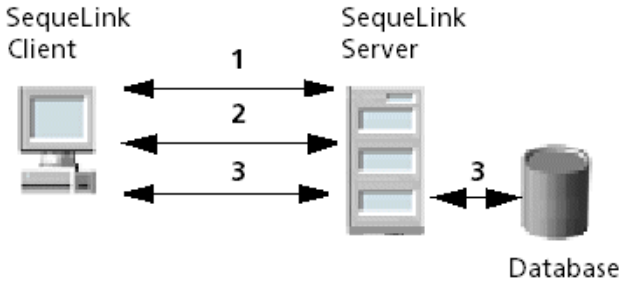

1 A network connection is established.

2 An authentication mechanism is used to establish the identity of the SequeLink Client to the SequeLink Server.

3 Based on information provided by the SequeLink Client application (for example, a database user name and password), a database connection is established.

Stage 1: Establishing a Network Connection

The first stage of the connection process involves establishing a network connection. The dialog that appears depends on whether the connection has been configured to connect directly to a SequeLink service or to retrieve connection information for the SequeLink service from a centralized LDAP directory.

Connecting Directly to a SequeLink® Service

If the connection has been configured to connect directly to a SequeLink service, the Connect to the SequeLink Server dialog box appears.

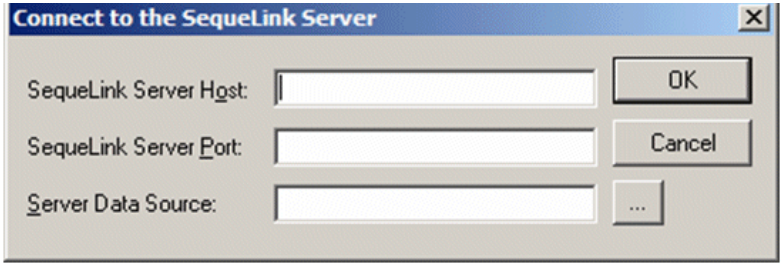

Provide the following information; then, click **OK**.

SequeLink Server Host: Type the TCP/IP host name of the SequeLink service.

SequeLink Server Port: Type the TCP/IP port on which the SequeLink service is listening. A default installation of SequeLink Server uses the port 19996.

Server Data Source: Type the name of a server data source to use for the connection or click the **...** button to select an existing data source. This step is optional. If a server data source is not specified, the default server data source for that service will be used for the connection.

Retrieving Connection Information from an LDAP Directory

If the connection has been configured to connect to an LDAP server to retrieve connection information from an LDAP directory, the Connect to the SequeLink Server dialog box appears.

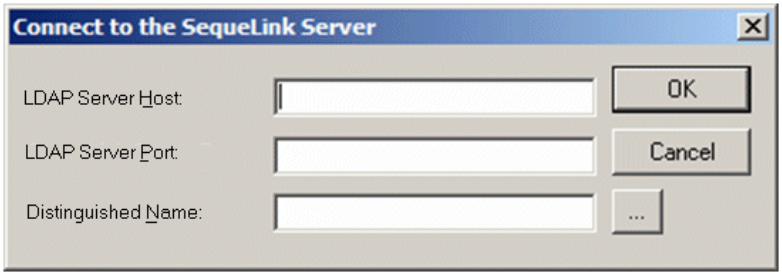

Provide the following information; then, click **OK**.

LDAP Server Host: Type the TCP/IP host name of the LDAP server.

LDAP Server Port: Type the TCP/IP port on which the LDAP server is listening.

Distinguished Name: Type the Distinguished Name (DN) of the LDAP entry.

See "Creating LDAP Entries for SequeLink® Services" on page 430 for information on setting up an LDAP server for SequeLink.

Stage 2: SequeLink® Server Authentication

The second stage of the connection process involves authentication of the SequeLink Client to the SequeLink Server. The dialog boxes that appear depend on how authentication is configured for the SequeLink service.

- When ServiceAuthMethods=anonymous or ServiceAuthMethods=integrated_nt, no dialogs appear.

- When ServiceAuthMethods=OSLogon(*HUID*,*HPWD*) or ServiceAuthMethods=OSLogon(*UID*,*PWD*), the Logon to SequeLink Service dialog box appears.

Provide the following information; then, click **OK**.

Host User Name: Type the host user name.

NOTE: When connecting to a Windows server, you must prefix the host user name with a server name, if authenticating to a local server, or a domain name (for example, SALES\DJONES). If the server name or domain name is omitted, the SequeLink Server will attempt to authenticate the user ID and password with the database account defined for the machine on which the SequeLink Server is running. If this validation fails, the SequeLink Server will attempt to authenticate the user ID and password with the database account defined for the domain of the machine on which the SequeLink Server is running.

Host Password: Type the host password.

- When ServiceAuthMethods=OSLogon(*HUID*,*HPWD*,*NPWD*) or ServiceAuthMethods=OSLogon(*UID*,*PWD*,*NPWD*) and the password is expired, the Password expired. Please specify new password dialog box appears.

NOTE: If the password is not expired, the previous dialog appears. You are only prompted for the Host User Name and Host Password.

Provide the following information; then, click **OK**.

Host User Name: Type the host user name.

NOTE: When connecting to a Windows server, you must prefix the host user name with a server name, if authenticating to a local server, or a domain name (for example, SALES\DJONES). If the server name or domain name is omitted, the SequeLink Server will attempt to authenticate the user ID and password with the database account defined for the machine on which the SequeLink Server is running. If this validation fails, the SequeLink Server will attempt to authenticate the user ID and password with the database account defined for the domain of the machine on which the SequeLink Server is running.

Host Password: Type the host password.

New Password: Type the new password to be used by the SequeLink password change mechanism.

Confirm Password: Type again the new password to confirm it.

See Chapter 13 “Configuring SequeLink® Security” on page 291 for more information about configuring authentication.

Stage 3: Data Store Logon

The last stage of the connection process involves logging on the data store. The dialogs that appear depend on the data store logon method configured for the SequeLink service:

- When DataSourceLogonMethod=OSIntegrated, no dialogs appear.
- When DataSourceLogonMethod=DBMSLogon(*UID*,*PWD*) or DataSourceLogonMethod=DBMSLogon(*DBUID*,*DBPWD*), a data store-specific user name and password are required and the Logon to SequeLink Service dialog box appears.

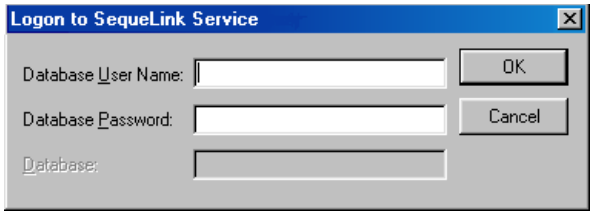

Provide the following information; then, click **OK**.

Database User Name: Type the database logon ID.

Database Password: Type the database password.

Database: Type the name of the database to which you want to connect. This field is disabled when the data store does not recognize the concept of databases.

See Chapter 13 “Configuring SequeLink® Security” on page 291 for more information about configuring data store logon methods.

Connecting with a Provider String

Once a data source is defined through the DataDirect Configuration Manager and the SequeLink *for* ADO Provider Setup Assistant, your application can connect directly to that data source. You can override the current settings for the data source when you connect using a *provider string*.

A provider string contains *attribute=value* pairs that control various aspects of the data provider's connection and interaction with the database. When an application names a specific data source to connect to, the application can also pass the data provider a provider string of *attribute=value* pairs. The data provider uses the values in the provider string instead of the default values defined for the data source in the system information.

Using provider strings allows application developers to configure connections for users programmatically and ensures that users have the optimum settings for working with the provider and database. Any values a user has set for a data source through the DataDirect Configuration Manager are overridden by corresponding values in the provider string for the current session only.

The provider string sets the DBPROP_INIT_PROVIDERSTRING initialization property and has the form:

```
"attribute=value;attribute=value;"
```

You can specify the `attribute=value` pairs on the Options tab of the DataDirect SequeLink 6.0 ADO Provider Setup window.

Refer to the *SequeLink Developer's Reference* for a list of ADO connection attributes.

Configuring Connection Failover and Client Load Balancing for the ADO Client

The ADO Client can help you make sure that your critical data is available even if the primary database server is unavailable:

- *Connection failover* allows an application to connect to an alternate, or backup, database server if the primary database server is unavailable, for example, because of a hardware failure or traffic overload. See "Using Connection Failover" on page 238 for more information.
- *Connection retry* defines the number of times the driver attempts to connect to the primary server and, if configured, alternate database servers after the initial unsuccessful connection attempt. See "Using Connection Retry" on page 239 for more information.
- *Client load balancing* helps distribute new connections in your environment so that no one server is overwhelmed with connection requests. "Using Client Load Balancing" on page 241 for more information.

Using Connection Failover

Connection failover allows an application to connect to an alternate, or backup, database server if the primary database server is unavailable, for example, because of a hardware failure or traffic overload. Connection failover ensures that the data on which your critical ADO applications depend is always available.

To configure connection failover, you *must* specify a list of alternate database servers that are tried at connection time if the primary SequeLink server is not accepting connections. To do this, specify the Alternate Servers connection option on the Options tab of the Setup dialog box or in the connection string. (See

"Configuring ADO Client Data Sources" on page 216 for details.) Connection attempts continue until a connection is successfully established or until all the database servers in the list have been tried once (the default).

IMPORTANT: If you specify an LDAP server and also enable connection failover, the alternate servers must be LDAP servers. The data provider will access the servers using the LDAP protocol.

Optionally, you can specify the following additional connection failover features:

- The number of times the ADO data provider attempts to connect to the primary and alternate SequeLink servers after the initial unsuccessful connection attempt. By default, the ADO data provider does not retry. See "Using Connection Retry" on page 239 for more information.
- The wait interval, in seconds, between attempts to connect to the primary and alternate database servers. The default interval is 3 seconds. See "Using Connection Retry" on page 239 for more information.
- Whether the ADO data provider will use client load balancing in its attempts to connect to primary and alternate SequeLink servers. If load balancing is enabled, the ADO data provider uses a random pattern instead of a sequential pattern in its attempts to connect. The default value is not to use load balancing. See "Using Client Load Balancing" on page 241 for more information.

Using Connection Retry

Connection retry defines the number of times the driver attempts to connect to the primary SequeLink Server and, if configured, alternate SequeLink Servers after the initial unsuccessful connection attempt. Connection retry can be an

important strategy for system recovery. For example, suppose you have a power failure in which both the SequeLink Client and the SequeLink Server fail. When the power is restored and all computers are restarted, the SequeLink Client may be ready to attempt a connection before the SequeLink Server has completed its startup routines. If connection retry is enabled, the client application can continue to retry the connection until a connection is successfully accepted by the SequeLink Server.

Connection retry can be used in environments that have only one server or can be used as a complementary feature with connection failover in environments with multiple SequeLink Servers.

Using connection options, you can specify the number of times the driver attempts to connect and the time in seconds between connection attempts.

Connection String Example

The following connection string configures the ADO data provider connected to a Linux, UNIX, or Windows server to use connection failover in conjunction with all of its optional features—load balancing, connection retry, and connection retry delay.

```
Data Source=MyADODSN;User Id=test;Password=secret;Alternate Servers=(Host=
server2:Port=19996:Server Data Source=SDSN2,Host=server3:Port=19996:Server
Data Source=SDSN3);Connection Retry Count=4;Connection Retry Delay=5;Load
Balancing=1
```

Specifically, if a successful connection is not established on the SequeLink ADO data provider's first pass through the list of SequeLink servers, this connection string configures the data provider to use two alternate servers as connection failover servers, to attempt to connect four additional times if the initial attempt fails, to wait five seconds between attempts, and to try the primary and alternate servers in a random order.

The additional connection information required for the alternate servers is specified in the SequeLink Server data source SDSN3.

Using Client Load Balancing

Client load balancing helps distribute new connections in your environment so that no one server is overwhelmed with connection requests. When client load balancing is enabled, the order in which primary and alternate database servers are tried is random.

First, SequeLink Server B is tried (1). Then, SequeLink Server C may be tried (2), followed by a connection attempt to SequeLink Server A (3). In contrast, if client load balancing were not enabled in this scenario, each SequeLink Server would be tried in sequential order, primary server first, then each alternate SequeLink server based on its entry order in the alternate servers list.

Connection Failover Properties

Table 9-2 summarizes the connection properties that control how connection failover works with the ADO Client. Refer to the *SequeLink Developer's Reference* for details about configuring each connection attribute.

Table 9-2. Summary: Connection Failover Connection Attributes for the ADO Client

Connection Attribute	Characteristic
Alternate Servers	List of alternate database servers. An IP address or server name and a port number are required for each server. The Server Data Source connection property is optional.
Connection Retry Count	Number of times the data provider retries the primary database server, and if specified, alternate servers until a successful connection is established. The initial default is 0.
Connection Retry Delay	Wait interval, in seconds, between connection retry attempts when the ConnectionRetryCount property is set to a positive integer. The initial default is 3.
Load Balancing	Sets whether the data provider will use client load balancing in its attempts to connect to the list of database servers (primary and alternate). If client load balancing is enabled, the data provider uses a random pattern instead of a sequential pattern in its attempts to connect. The initial default is false (client load balancing is not used).

Refer to the *SequeLink Developer's Reference* for overviews of connection failover and client load balancing.

Importing and Exporting ADO Client Data Sources

The SequeLink Data Source SyncTool allows you to export ADO client data source definitions to data source files and distribute them to multiple end users. The SequeLink Data Source SyncTool provides two user implementations, one for the SequeLink administrator and another for the end user:

- The SequeLink *for* ADO Data Source SyncTool Administrator is used by the SequeLink administrator to create data source files. It can import and export data sources. This tool should be made available to the SequeLink administrator only.
- The SequeLink *for* ADO Data Source SyncTool is used by the end user and can import data sources only. It should be installed on every SequeLink Client *for* ADO.

In addition, you can create a customized, installable image of SequeLink Client *for* ADO with predefined, site-specific settings, including data source files created with the SequeLink Data Source SyncTool. This customized, installable image is called a *Quick Install image*. For more information about creating Quick Install images, refer to the *SequeLink Installation Guide*.

The window title bar of the SequeLink Data Source SyncTool indicates whether you, or the end user, is performing an export or an import operation. Also, context-sensitive online help is available by clicking **?** on the title bar; then, clicking the area about which you want more information.

Exporting ADO Client Data Sources

1. From the SequeLink program manager group, double-click the **ADO Data Source SyncTool Administrator** icon. The SequeLink *for* ADO Data Source SyncTool Administrator Welcome window appears.
2. Select the **Manage Data Sources Files** option; then, click **Next**.
3. Select a data source file from the Filename list box, or click **Browse** to find a data source file not listed. The default extension for data source files is .OSF.

   To create a new data source file, click **New**.
4. Select whether you want to export User or System data sources to the data source file you selected; then click **Next**.
5. Select the data sources you want to export to the data source file.

   NOTE: You cannot export grayed-out data sources, which are data sources that are configured for a previous incompatible version of the ADO provider.
6. Using the following symbols, verify that the appropriate actions will be performed on the data sources in the data source file:

   - The data source will remain unchanged.
   - The data source will be added to the data source file.
   - The data source will be deleted from the data source file.
   - The data source will be updated in the data source file.
7. Type a description for the data source file; then, click **Next**. This description will appear when the end user selects this file for importing.

8 Select the mode the end user will use to import these data sources; then, click **Next**.

- *Interactive mode* allows the user to select which data sources will be imported. This mode is not supported by the Quick Install feature; the Quick Install feature supports only data source files created with the Merge or Overwrite options. For instructions on creating Quick Install images, refer to the *SequeLink Installation Guide*.
- *Merge mode* adds or updates all the data sources in the data source file without deleting other data sources.
- *Overwrite mode* adds or updates the data sources in the data source file and deletes any other data sources configured for the ADO provider.

9 Select the option that will determine how the end user will be able to import the data sources you exported to the data source file; then, click **Next**.

- *Suggest SequeLink User DSN*. When imported, the SequeLink *for* ADO Data Source SyncTool will suggest to the end user that these data sources be imported as User data sources, but will allow them to be imported as User or System data sources.
- *Suggest SequeLink System DSN*. When imported, the SequeLink *for* ADO Data Source SyncTool will suggest to the end user that these data sources be imported as System data sources, but will allow them to be imported as User or System data sources.
- *Force SequeLink User DSN*. When imported, the SequeLink *for* ADO Data Source SyncTool will allow these data sources to be imported as User data sources only.
- *Force SequeLink System DSN*. When imported, the SequeLink *for* ADO Data Source SyncTool will allow these data sources to be imported as System data sources only.

10 Click **Finish** to quit.

Importing ADO Client Data Sources

The SequeLink administrator and end user use a slightly different implementation of the SequeLink *for* ADO Data Source SyncTool to import ADO data source definitions.

To import ADO client data sources:

1. From the SequeLink program manager group, double-click the appropriate ADO SyncTool icon. The Welcome window appears.
2. Select the **Import** option; then, click **Next**.

   NOTE: If using the SequeLink *for* ADO Data Source SyncTool Administrator, select the **Import Data Sources** option, and click **Next**.
3. Select a data source file from the Filename list box, or click **Browse** to find a data source file not listed. The default extension for data source files is .OSF.
4. Verify that the appropriate actions will be performed on the data sources on your local machine; then, click **Next**. Depending on the import mode that was set when the data source file was exported, you may see the following symbols:

   - The data source will remain unchanged.
   - The data source will be added to your local machine.
   - The data source will be deleted from your local machine.
   - The data source will be updated to your local machine.

   NOTE: Grayed-out data sources are data sources that are configured for a previous incompatible version of the ADO provider; these data sources will remain unchanged unless you update them in Interactive mode with a data source configured for the current version of the ADO provider.
5. Click **Finish** to quit.

Authentication

The ADO provider supports the following methods of authentication:

- **Anonymous**. The SequeLink Server accepts connections from the SequeLink Client without verifying the client's identity.
- **Operating system user ID** and **password**. The SequeLink Server verifies the identity of the SequeLink Client using a user ID and password that must be valid for the platform on which the SequeLink Server is running. If verified, the server accepts the user ID as the identity of the client and permits the connection.
- **Kerberos**. Kerberos authentication uses Kerberos, a trusted third-party authentication service, to verify user identities. Kerberos authentication can take advantage of the user name and password maintained by the operating system to authenticate users to the database. This method requires knowledge of how to configure your Kerberos environment.
- **Integrated NT**. This option is supported for connections between SequeLink Server for Windows servers and ODBC Clients, ADO Clients, and .NET Clients on Windows only. The SequeLink Server verifies the identity of the SequeLink Client using the client's Windows network logon credentials instead of a Windows user ID and password.

For details on configuring authentication for SequeLink, see "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.

Using SSL Encryption

If your SequeLink environment requires greater data privacy than that provided by fixed-key DES, fixed-key 3DES, or byteswap, you can use the Secure Socket Layer (SSL) to encrypt data exchanged between the ADO Client and the SequeLink Server.

SequeLink supports the use of anonymous ciphers. Anonymous ciphers allow the SSL connection to succeed without proper authentication of the peer by using the DH algorithm. The ADO Client supports the following cryptographic strong SSL cipher suites:

- TLS_DH_anon_WITH_AES_128_CBC_SHA
- TLS_DH_anon_WITH_AES_256_CBC_SHA

For details on configuring SSL for SequeLink, see "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.

Transliterating Character Data

The ADO Client uses the standard service template to transliterate character data from the code page the SequeLink service is using to the code page of the client application/system. In some environments, the ServiceCodePage service attribute must also be set.

On Linux/UNIX/Windows, the ADO Client uses the OS setting for ServiceCodePage to transliterate character data, in addition to the standard service template. This option enables support for

single-byte character sets (SBCS) and multi-byte character sets (MBCS). However, Unicode is not supported.

For DB2 for z/OS, the ADO Client uses only the standard service template to transliterate character data. This approach enables support for single-byte character sets, multi-byte character sets, or Unicode.

See Appendix D "SequeLink® Service Attributes" on page 491 for information about the SequeLink service attributes that affect configuration. For information about data type mappings, refer to the *SequeLink Developer's Reference*.

10 Configuring the JDBC Client

This chapter describes the tasks you perform to configure and manage the SequeLink Client *for* JDBC (the JDBC Client).

The JDBC Client runs on 32-bit and 64-bit platforms. No changes are required to existing applications to enable them to run on 64-bit platforms.

About JDBC Connections

You can open a JDBC connection to a SequeLink service by specifying a JDBC connection URL or configuring a JDBC client data source. This section explains how to connect to a SequeLink service using connection URLs.

For information about configuring JDBC client data sources and a list of JDBC connection attributes and their valid values, refer to the *SequeLink Developer's Reference*.

Specifying JDBC Driver Connection URLs

The connection URL format depends on whether you are using SSL encryption. See Chapter 14 "Configuring the SequeLink® Proxy Server" on page 351 for more information about SSL encryption and certificate-based authentication.

NOTE: SequeLink Server for DB2 for z/OS does not support SSL. To use SSL encryption in a DB2 for z/OS environment, use the SequeLink Proxy Server.

If not using SSL encryption over the SequeLink Proxy Server, the connection URL format is:

```
jdbc:sequelink://hostname:port[;key=value]...
```

If using SSL encryption over the SequeLink Proxy Server, the connection URL format is:

```
jdbc:sequelink:ssl://hostname:port[;key=value]...
```

where:

hostname	is the TCP/IP address or TCP/IP host name of the SequeLink server to which you are connecting. NOTE: Untrusted applets cannot open a socket to a machine other than the originating host. See "Using the SequeLink® Proxy Server" on page 351 for more information about untrusted applets.
port	is the TCP/IP port on which the SequeLink server is listening. A default installation of SequeLink Server uses the port 19996.
key=value	specifies connection properties. Refer to the *SequeLink Developer's Reference* for a list of connection properties and their valid values.

Connection URL Examples

The following examples show some typical JDBC driver connection URLs:

```
jdbc:sequelink://sequelinkhost:19996;
```

```
jdbc:sequelink://189.23.5.25:19996;user=john;
password=whatever
```

```
jdbc:sequelink://189.23.5.132:19996;databaseName=stores7
```

```
jdbc:sequelink://189.23.5.68:19996;databaseName=pubs;
HUser=john;HPassword=whatever
```

```
jdbc:sequelink://sequelinkhost:4006;
databaseName=pubs;DBUser=john;DBPassword=whatever
```

```
jdbc:sequelink:ssl://mysecurehost:9500;
cipherSuites=SSL_DH_anon_WITH_RC4_128_MD5
```

```
jdbc:sequelink:ssl://mysecurehost:9502;
cipherSuites=SSL_DHE_RSA_WITH_DES_CBC_SHA;
certificateChecker=CheckAgainstCertificateFromJar
```

The preceding examples do not show the user and password connection properties. Typically, these properties are specified in the connection properties stored in the java.util.Properties object, which is supplied as a parameter to the getConnection method.

Connecting Using LDAP and JNDI

The JDBC Client connection settings can be stored in an LDAP repository using Java Naming and Directory Interface (JNDI) using javaReferenceAddress.

The following example shows a JNDI entry for SequeLink:

```
dn: cn=DB2V8 on belg-keidis,ou=JVS,ou=Arne,ou=USERS,o=Development
cn: DB2V8 on belg-keidis
objectClass: top
objectClass: javaContainer
objectClass: javaObject
objectClass: javaNamingReference
javaReferenceAddress: #0#description#
javaReferenceAddress: #1#portNumber#6007
javaReferenceAddress: #2#serverName#belg-keidis
javaReferenceAddress: #3#databaseName#jvs
javaReferenceAddress: #4#serverDataSource#Default
javaReferenceAddress: #5#encrypted#1
javaFactory: com.ddtek.jdbcx.sequeLink.SequeLinkDataSourceFactory
javaClassName: com.ddtek.jdbcx.sequelink.SequeLinkDataSource
```

Configuring Connection Failover for the JDBC Client

Connection failover allows an application to connect to an alternate, or backup, database server if the primary database is unavailable, for example, because of a hardware failure or traffic overload. Connection failover ensures that the data on which your critical JDBC applications depend is always available.

You can customize the JDBC Client for connection failover by configuring a list of alternate database servers that are tried if the primary server is not accepting connections. Connection

attempts continue until a connection is successfully established or until all of the alternate database servers have been tried the specified number of times.

Specifying Primary and Alternate Servers

Connection information for primary and alternate servers can be specified using either a connection URL through the JDBC Driver Manager or a JDBC data source. For example, the following connection URL for the JDBC Client specifies connection information for the primary and alternate servers using a connection URL:

```
jdbc:sequelink://server1:19996;serverDataSource=SDSN1;User=test;
Password=secret;AlternateServers=(server2:19996;serverDataSource=SDSN2,
server3:19996;serverDataSource=SDSN3)
```

In this example:

```
...server1:19996;serverDataSource=SDSN1...
```

is the part of the connection URL that specifies connection information for the primary server. Alternate servers are specified using the AlternateServers property. For example:

```
...;AlternateServers=(server2:19996;serverDataSource=SDSN2,
server3:19996;serverDataSource=SDSN3)
```

Similarly, the same connection information for the primary and alternate servers specified using a JDBC data source would look like this:

```
SequeLinkDataSource mds = new SequeLinkDataSource();
mds.setDescription("My SequeLinkDataSource");
mds.setServerName("server1");
mds.setPortNumber(19996);
mds.setServerDataSource("SDSN1");
mds.setUser("test");
```

```
mds.setPassword("secret");
mds.setAlternateServers(server2:19996;serverDataSource=SDSN2,
server3:19996;ServerDataSource=SDSN3)
```

In this example, connection information for the primary server is specified using the ServerName, PortNumber, and serverDataSource properties. Connection information for the alternate servers is specified using the AlternateServers property.

Using the AlternateServers Property

Connection information for alternate servers is specified using the AlternateServers property with either a connection URL or a JDBC data source. The value of the AlternateServers property is a string that has the format:

```
(servername1:port1[;serverDataSource=SDSN1;][,servername2:
port2[;serverDataSource=SDSN2]...)
```

where:

servername1 is the IP address or server name of the first alternate database server, *servername2* is the IP address or server name of the second alternate database server, and so on. The IP address or server name is required for each alternate server entry.

port1 is the port number on which the first alternate database server is listening, *port2* is the port number on which the second alternate database server is listening, and so on. The port number is required for each alternate server entry.

serverDataSource=SDSN2 is the server data source name for the first alternate server, *serverDataSource=SDSN3* is the server data source name for the second alternate server, and so on. Currently, the only optional property that can be set for the alternate server is serverDataSource. For example:

```
jdbc:sequelink://server1:19996;serverDataSource=SDSN1;User=test;Password=
secret;AlternateServers=(server2:19996;serverDataSource=SDSN2)
```

Specifying Connection Retry

Connection retry allows the JDBC Client to retry connections to the primary database server, and if specified, alternate servers until a successful connection is established. You use the ConnectionRetryCount and ConnectionRetryDelay properties to enable and control how connection retry works. For example:

```
jdbc:sequelink://server1:19996;serverDataSource=SDSN1;User=test;Password=
secret;AlternateServers=(server2:19996;serverDataSource=SDSN2,
server3:19996;serverDataSource=SDSN3);ConnectionRetryCount=2;
ConnectionRetryDelay=5
```

In this example, if a successful connection is not established on the JDBC Client's first pass through the list of database servers (primary and alternate), the JDBC Client retries the list of servers in the same sequence twice (`ConnectionRetryCount=2`). Because the connection retry delay has been set to five seconds (`ConnectionRetryDelay=5`), the JDBC Client waits five seconds between retry passes.

Specifying Load Balancing

Client load balancing helps distribute new connections in your environment so that no one server is overwhelmed with connection requests. When client load balancing is enabled, the order in which primary and alternate database servers are tried is random.

When Connection Retry is also enabled, the JDBC Client tries to connect to the primary SequeLink server and alternate SequeLink servers in a random order until a successful connection is established. If the connection attempt fails, the driver again randomly selects from the list of servers until all SequeLink servers in the list have been tried or a connection is successfully established.

Connection Failover Properties

Table 10-1 summarizes the connection properties that control how connection failover works with the JDBC Client. Refer to the *SequeLink Developer's Reference* for details about configuring each property.

Table 10-1. Summary: Connection Failover Properties for the JDBC Client

Property	Characteristic
AlternateServers	List of alternate database servers. An IP address or server name and a port number are required for each server. The serverDataSource connection property is optional.
ConnectionRetryCount	Number of times the driver retries the primary database server, and if specified, alternate servers until a successful connection is established. The initial default is 0.
ConnectionRetryDelay	Wait interval, in seconds, between connection retry attempts when the ConnectionRetryCount property is set to a positive integer. The initial default is 3.
LoadBalancing	Sets whether the driver will use client load balancing in its attempts to connect to the list of database servers (primary and alternate). If client load balancing is enabled, the driver uses a random pattern instead of a sequential pattern in its attempts to connect. The initial default is false (client load balancing is not used).

Refer to the *SequeLink Developer's Reference* for overviews of connection failover and client load balancing.

Testing JDBC Connections

This section describes how to test your connection with DataDirect Test™ for JDBC (DataDirect Test). For more information about using DataDirect Test, refer to the *SequeLink Developer's Reference*.

To connect with the JDBC Client using DataDirect Test for JDBC:

1 Start DataDirect Test as a Java application or applet.

- As a Java application on Windows: Run the TestForJDBC.bat or TestForJDBC14.sh file located in the testforjdbc directory.

UNIX

- As a Java application on Linux/UNIX: Run the TestForJDBC.sh or or TestForJDBC14.sh shell script located in the testforjdbc directory.

NOTE: Use TestForJDBC14.sh with JDK 1.4 or higher.

2 From the DataDirect Test Welcome window, click the **Press Here To Continue** button. The DataDirect Test window appears.

3 Select **Driver** / **Register Driver**. DataDirect Test prompts you for the JDBC driver you want to load.

4 In the Please Supply a Driver URL field, make sure that the following driver is specified; then, click **OK**.

```
com.ddtek.jdbc.sequelink.SequeLinkDriver
```

5 Select **Connection / Connect To DB**. The Select A Database window appears with a list of default JDBC driver connection URLs.

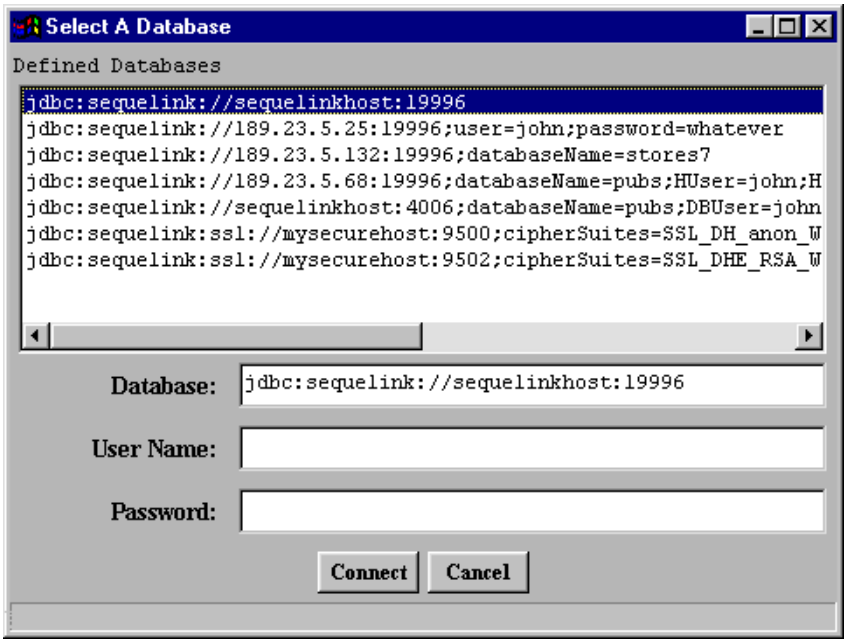

6 Select one of the default JDBC driver connection URLs. In the Database field, modify the default values of the connection URL appropriately for your environment.

7 In the User Name and Password fields, type the required user and password connection properties; then, click the **Connect** button. Refer to the *SequeLink Developer's Reference* for information about JDBC connection properties.

8 If the connection was successful, the Connection window shows the `Connection Established` message in the JDBC/Database Output scroll box.

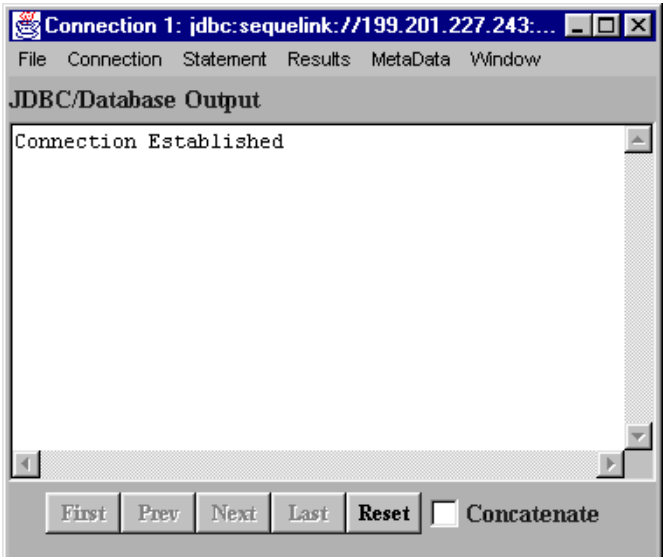

If the connection was successful, you can start using your JDBC applications with SequeLink.

If the connection was unsuccessful, you are returned to the DataDirect Test window. The `getConnection() Failed:` message appears in the JDBC/Database Output scroll box. If your connection failed, refer to the *SequeLink Troubleshooting Guide and Reference*.

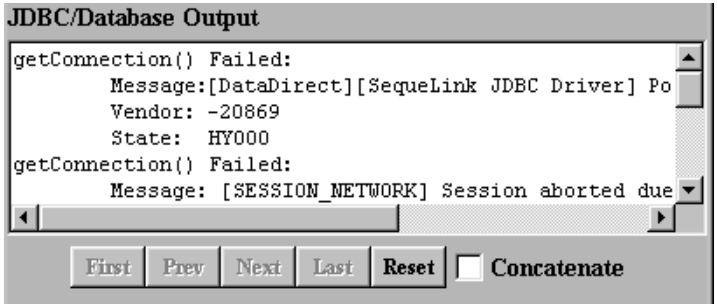

Authentication

The JDBC driver supports the following methods of authentication:

- **Anonymous**. The SequeLink Server accepts connections from the SequeLink Client without verifying the client's identity.
- **Operating system user ID** and **password**. The SequeLink Server verifies the identity of the SequeLink Client using a user ID and password that must be valid for the platform on which the SequeLink Server is running. If verified, the server accepts the user ID as the identity of the client and permits the connection.
- **Kerberos**. Kerberos authentication uses Kerberos, a trusted third-party authentication service, to verify user identities. Kerberos authentication can take advantage of the user name and password maintained by the operating system to authenticate users to the database. This method requires knowledge of how to configure your Kerberos environment.
- **Integrated NT**. This option is supported for connections between SequeLink Server for Windows servers and ODBC Clients, ADO Clients, and .NET Clients on Windows only. The SequeLink Server verifies the identity of the SequeLink Client using the client's Windows network logon credentials instead of a Windows user ID and password.

For details on configuring authentication for SequeLink, see "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.

Using SSL Encryption

If your SequeLink environment requires greater data privacy than that provided by fixed-key DES, fixed-key 3DES, or byteswap, you can use the Secure Socket Layer (SSL) to encrypt data exchanged between the JDBC Client and the SequeLink Server.

SequeLink supports the use of anonymous ciphers. Anonymous ciphers allow the SSL connection to succeed without proper authentication of the peer by using the DH algorithm. The JDBC Client supports the following cryptographic strong SSL cipher suites:

- TLS_DH_anon_WITH_AES_128_CBC_SHA
- TLS_DH_anon_WITH_AES_256_CBC_SHA

The JDBC Client adds a direct encryption mechanism using the `encrypted` option. For example:

```
jdbc:sequelink://mysecurehost:9503;encrypted=1
```

This type of encryption is directly configured on the SequeLink Server 6.0 and does not require the use of SequeLink Proxy Server.

NOTE: In previous releases of SequeLink, encryption for the JDBC Client was configured through the SequeLink Proxy Server. The SequeLink Proxy Server is still used for certificate-based authentication for JDBC applications (see Chapter 14 “Configuring the SequeLink® Proxy Server” on page 351). The SequeLink Proxy Server is also used to provide SSL encryption when using SequeLink in a DB2 for z/OS environment.

For details on configuring SSL for SequeLink, see "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.

Code Page Support

All code pages supported by the database server are now available from the SequeLink Client *for* JDBC.

See Appendix D "SequeLink® Service Attributes" on page 491 for information about the SequeLink service attributes that affect configuration. For information about data type mappings, refer to the *SequeLink Developer's Reference*.

11 Configuring the .NET Client

This chapter describes the tasks you perform to configure and manage the SequeLink Client *for* .NET (the .NET Client).

Specifying Connection Properties

You can modify a connection by specifying connection attributes in the connection string. Optionally, you may specify attribute= value pairs in the connection string to override the default values stored in the data source. You can specify long or short names in the connection string, which has the form:

"*attribute=value;attribute=value;attribute=value*"

For example, in the following example, the connection string can be used to connect to the DB2 host hal:

```
"Host=hal;Port=19998;User Id=test01;Password=test01;Database=test;"
```

NOTE: All connection string attribute names are case-insensitive. For example, Password is the same as password. However, the values of some attributes, such as User ID and Password, may be case-sensitive.

Refer to the *SequeLink Developer's Reference* for information about specifying attributes through the .NET Client's Connection object.

Testing .NET Connections

Once the .NET Client is installed, you can connect from your application to your database with a connection string.

The following example illustrates connecting to the underlying Sybase database using the .NET Client from an application developed in Visual Studio .NET. If you are connecting using a different ADO.NET data provider or connecting from the command line, the specific details vary.

If you are using Visual Studio .NET:

1 In the Solution Explorer, right-click **References**; then, click **Add Reference**.

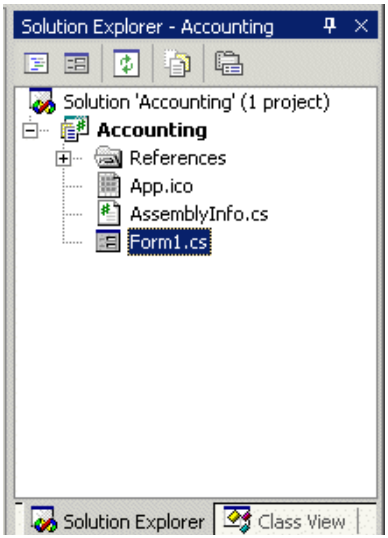

2 Select the DataDirect Technologies SequeLink .NET Client in the component list.

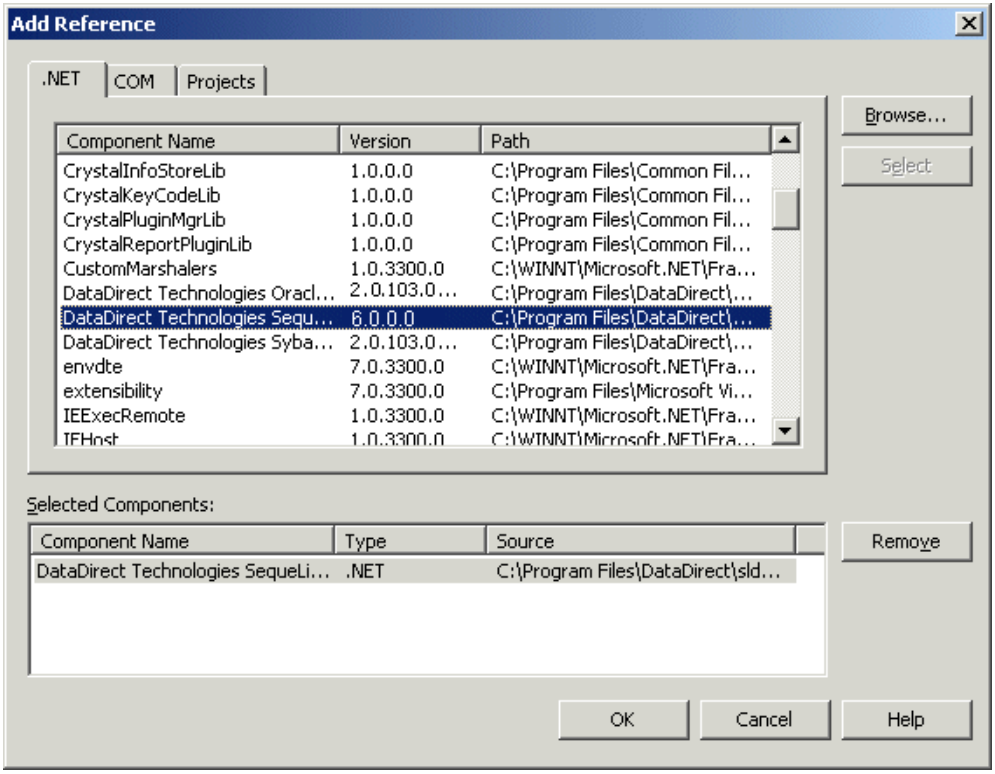

3 Click **OK**. The Solution Explorer now includes a reference for the .NET Client.

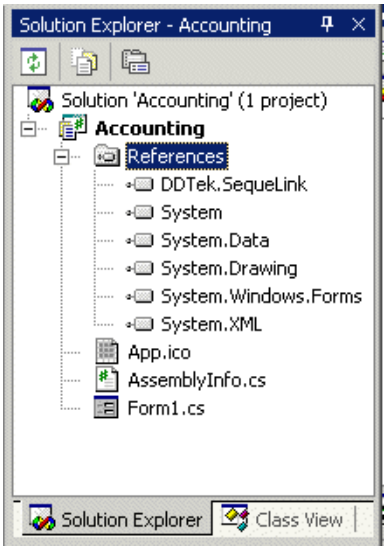

4 Check the beginning of your application. If the data provider's namespace is not present, add it, as shown in the following code fragments:

C#

```
// Access SequeLink Server
using System.Data;
using DDTek.SequeLink;
```

Visual Basic

```
' Access SequeLink Server
Imports System.Data
Imports DDTek.SequeLink
```

5 Add the connection information for your server and exception handling code, as shown in the following C# code fragment. (Refer to the *SequeLink Developer's Reference* for the connection string options.)

C#

```
DBConn = new SequeLinkConnection("Host=sydney;
   Port=19998;User ID=test01;Password=test01;
   Database=test");

try
{
   DBConn.Open();
   Console.WriteLine ("Connection successful!");
}
// Display any exceptions
catch (SequeLinkException ex)
{
   // Connection failed
   Console.WriteLine (ex.Message);
   return;
}
```

Visual Basic

```
Dim Conn As New SequeLinkConnection("Host=bowhead;
Port=19996;User ID=test01;Password=test01;
Database=test")

Try
  'Open the connection
  Conn.Open()
  MessageBox.Show("Connection successful!")

Catch SLex As SequeLinkException
  'Connection failed
  MessageBox.Show(SLex.Message, "SequeLink Exception")

End Try
```

6 Close the connection.

C#

```
// Close the connection
DBConn.Close();
```

Visual Basic

```
'Close the connection
Conn.Close()
```

For information about using connection strings and connection string options, refer to the *SequeLink Developer's Reference*.

Using Connection Failover with the .NET Client

The .NET Client can help you make sure that your critical data is available even if the primary database server is unavailable:

- *Connection failover* allows an application to connect to an alternate, or backup, database server if the primary database server is unavailable, for example, because of a hardware failure or traffic overload. See "Specifying Primary and Alternate Servers" on page 271 for more information.
- *Connection retry* defines the number of times the data provider attempts to connect to the primary server and, if configured, alternate database servers after the initial unsuccessful connection attempt. See "Using Connection Retry" on page 271 for more information.
- *Client load balancing* helps distribute new connections in your environment so that no one server is overwhelmed with connection requests. See "Using Client Load Balancing" on page 272 for more information.

Specifying Primary and Alternate Servers

Connection failover allows an application to connect to an alternate, or backup, database server if the primary database is unavailable, for example, because of a hardware failure or traffic overload.

You can customize the .NET data provider for connection failover by configuring a list of alternate database servers that are tried if the primary server is not accepting connections. Connection attempts continue until a connection is successfully established or until all of the alternate database servers have been tried the specified number of times.

The following C# code fragment includes a connection string that configures the .NET Client to use connection failover:

```
SequeLinkConnection Conn = new SequeLinkConnection();
Conn = new SequeLinkConnection("Host=server1;port=19996;ServerDataSource=
SDSN1:User ID=test;Password=secret;Alternate Servers=(Host=server2:Port=
19996:ServerDataSource=SDSN2,Host=server3:Port=19996:ServerDataSource=
SDSN3);Connection Timeout=60"
```

Specifically, this connection string configures the .NET Client to use two alternate servers as connection failover servers.

Using Connection Retry

Connection retry defines the number of times that the data provider attempts to connect to the primary, and, if configured, alternate database servers after the first unsuccessful connection attempt. Connection retry can be an important strategy for system recovery. For example, suppose you have a power failure scenario in which both the client and the server fail. When the power is restored and all computers are restarted, the client may be ready to attempt a connection before the server has

completed its startup routines. If connection retry is enabled, the client application can continue to retry the connection until a connection is successfully accepted by the server.

Connection retry can be used in environments that only have one server or can be used as a complementary feature in connection failover in environments with multiple servers.

The following C# code fragment includes a connection string that configures the .NET Client to use connection failover in conjunction with connection retry and connection retry delay:

```
SequeLinkConnection Conn = new SequeLinkConnection();
Conn = new SequeLinkConnection("Host=server1;port=19996;ServerDataSource=
SDSN1:User ID=test;Password=secret;Alternate Servers=(Host=server2:Port=
19996:ServerDataSource=SDSN2,Host=server3:Port=19996:ServerDataSource=
SDSN3);Connection Retry Count=4;Connection Retry Delay=5;Connection Timeout=
60"
```

Specifically, this connection string configures the .NET Client to use two alternate servers as connection failover servers, to attempt to connect four additional times if the initial attempt fails, and to wait five seconds between attempts. Each connection attempt lasts for 60 seconds.

Using Client Load Balancing

Client load balancing helps distribute new connections in your environment so that no one server is overwhelmed with connection requests. When client load balancing is enabled, the order in which primary and alternate database servers are tried is random.

When Connection Retry is also enabled, the .NET Client tries to connect to the primary SequeLink server and alternate SequeLink servers in a random order until a successful connection is established. If the connection attempt fails, the .NET Client again randomly selects from the list of servers until all SequeLink

servers in the list have been tried or a connection is successfully established.

The following C# code fragment includes a connection string that configures the .NET Client to use connection failover in conjunction with all of its optional features—load balancing, connection retry, and connection retry delay:

```
SequeLinkConnection Conn = new SequeLinkConnection();
Conn = new SequeLinkConnection("Host=server1;port=19996;ServerDataSource=
SDSN1:User ID=test;Password=secret;Alternate Servers=(Host=server2:Port=
19996:ServerDataSource=SDSN2,Host=server3:Port=19996:ServerDataSource=
SDSN3);Connection Retry Count=4;Connection Retry Delay=5;Load Balancing=true;
Connection Timeout=60"
```

Specifically, this connection string configures the .NET Client to use two alternate servers as connection failover servers, to attempt to connect four additional times if the initial attempt fails, to wait five seconds between attempts, and to try the primary and alternate servers in a random order. Each connection attempt lasts for 60 seconds.

Connection Failover Properties

Table 11-1 summarizes the connection string options that control how connection failover works with the .NET Client. Refer to the *SequeLink Developer's Reference* for details about configuring each property.

Table 11-1. Summary: Connection Failover Connection String Options for the .NET Client

Connection String Option	Characteristic
Alternate Servers	List of alternate database servers. An IP address or server name and a port number are required for each server. The Server Data Source connection string option is optional.

Table 11-1. Summary: Connection Failover Connection String Options for the .NET Client *(cont.)*

Connection String Option	Characteristic
Connection Retry Count	Number of times the data provider retries the primary database server, and if specified, alternate servers until a successful connection is established. The initial default is 0.
Connection Retry Delay	Wait interval, in seconds, between connection retry attempts when the ConnectionRetryCount property is set to a positive integer. The initial default is 3.
Load Balancing	Sets whether the data provider will use client load balancing in its attempts to connect to the list of database servers (primary and alternate). If client load balancing is enabled, the data provider uses a random pattern instead of a sequential pattern in its attempts to connect. The initial default is false (client load balancing is not used).

Refer to the *SequeLink Developer's Reference* for overviews of connection failover and client load balancing.

Code Page Support

The .NET Client supports Unicode as specified in the .NET Framework SDK. Effectively, this means that the .NET Client uses Unicode UTF-16 encoding to represent characters.

All code pages supported by the database server are available from the .NET Client.

See Appendix D "SequeLink® Service Attributes" on page 491 for information about the SequeLink service attributes that affect configuration. For information about data type mappings, refer to the *SequeLink Developer's Reference*.

Part 3: Configuring SequeLink® in Your Environment

This part contains the following chapters:

12 Configuring Transliteration

Transliteration is the transformation of text from one character set to another. In SequeLink, transliteration is defined by a combination of factors, which include which service template and SequeLink Client are used, as well as database-specific configuration settings.

See Appendix F "Internationalization, Localization, and Unicode" on page 597 for more information about internationalization.

Choosing a Service Template

When you install SequeLink Server, at least one SequeLink data access service, and one or more service templates determined by the SequeLink Server, are installed. You can use the SequeLink Manager to create additional services based on the SequeLink service templates and to modify the default service attributes defined in the service templates:

- UNIX [SequeLink 6.0] Agent service
- z/OS [SequeLink 6.0] Agent service for z/OS
- z/OS [SequeLink 6.0] DB2 service for z/OS
- UNIX [SequeLink 6.0] DB2 UDB LUW service
- UNIX [SequeLink 6.0] DB2 UDB LUW service (enhanced code page support)
- UNIX [SequeLink 6.0] Informix service (32-bit only)
- UNIX [SequeLink 6.0] JDBC Socket service (32-bit only)
- UNIX [SequeLink 6.0] ODBC Socket service

[SequeLink 6.0] ODBC Socket service (enhanced code page support)

[SequeLink 6.0] ODBC Socket service (enhanced code page support - UTF8 encoding)

[SequeLink 6.0] ODBC Socket service (enhanced code page support - UTF16 encoding)

[SequeLink 6.0] Oracle 9 service (32-bit only)

[SequeLink 6.0] Oracle 9 service (enhanced code page support - 32-bit only)

[SequeLink 6.0] Oracle 10 service

[SequeLink 6.0] Oracle 10 service (enhanced code page support)

[SequeLink 6.0] SQL Server service

[SequeLink 6.0] SQL Server service (enhanced code page support)

[SequeLink 6.0] Sybase service

[SequeLink 6.0] Sybase service (enhanced code page support)

[SequeLink 6.0] Sybase service (enhanced code page support - UTF8 encoding)

[SequeLink 6.0] Sybase service (enhanced code page support - UTF16 encoding)

In most cases, the SequeLink service templates provide a configuration that can be used without any modification. However, some service templates provide additional transliteration capabilities. You must select the service template that meets the transliteration requirements of your application.

NOTES:

- If you want enhanced code page support for the SequeLink Server for ODBC Socket or SequeLink Server for Sybase, use the UTF-16 encoding version of the template.
- SequeLink Server for DB2 for z/OS has only one service template. Unicode is enabled for all SequeLink Clients.

Table 12-1 lists the transliteration environments that the service templates support on Linux/UNIX/Windows.

Table 12-1. Selecting a Service Template on Linux/UNIX/Windows

Service Template	Transliteration Environment
Standard service template Automatically sets ServiceCodePage= Default (see "ServiceCodePage" on page 566)	■ ODBC and ADO Clients: The database code page is a 7-bit or 8-bit code page and matches the code page of the client application. ■ JDBC and .NET Clients: The database code page is a 7-bit or 8-bit code page and has a binary-compatible transliteration to UTF-16. ■ Operating system (OS) and database code pages use a single-byte or multi-byte character set.*
Service template with enhanced code page support Automatically sets ServiceCodePage= "Database" (see "ServiceCodePage" on page 566)	The database contains Unicode data or the operating system and database code pages differ.* If you want enhanced code page support for the SequeLink Server for ODBC Socket or SequeLink Server for Sybase, we recommend using the UTF-16 encoding version of the template.

* Not available for SequeLink Server for Informix, SequeLink Server for DB2 UDB on z/OS USS, or SequeLink Server for JDBC Socket.

Choosing the Transliteration Mode

SequeLink supports three transliteration modes:

- Default transliteration mode (see "Default Transliteration Mode")
- OS transliteration mode, which is used when the operating system (OS) and database code page are single-byte or multi-byte (see "OS Transliteration Mode" on page 281)
- Database transliteration mode, which is used to support Unicode databases or when the operating system and database code pages differ (see "Database Transliteration Mode" on page 283)

Default Transliteration Mode

Default transliteration mode can be used with all SequeLink Clients and with SequeLink Servers on Linux, UNIX, and Windows. To use Default transliteration mode, set ServiceCodePage=Default.

Use Default transliteration mode *only* when you created your service with the standard template. Do not use Default transliteration mode with the service template with enhanced code support.

- For the ODBC and ADO Clients, the administrator defines this mode when the database code page is a 7-bit or 8-bit code page and matches the code page of the client application.
- For the JDBC and .NET Clients, the administrator defines this mode when the database code page is a 7-bit or 8-bit code page and has a binary-compatible transliteration to UTF-16. In other words, it is possible to convert from the single-byte database code page to UTF-16, by setting either the higher or lower byte to zero.

OS Transliteration Mode

OS transliteration mode can be used when the operating system (OS) and database code pages use a single-byte or multi-byte character set. OS transliteration mode is useful for European and Asian languages. To use OS transliteration mode, set ServiceCodePage=OS.

Use OS transliteration mode *only* when you created your service with the standard template. Do not use this mode with the service template with enhanced code support.

To configure operating system environment settings:

- On Linux and UNIX, set locale settings for the service, for example, using the LC_ALL environment variable. Consult your system administrator for the supported values for the LC_ALL environment variable, and to find out which locale settings are supported on your system.
- On Windows, set regional options on the Control Panel.

Table 12-2 provides the database transliteration settings to use when configuring a SequeLink service.

Table 12-2. Defining Database-Specific Settings for OS Transliteration Mode

Database	Requirements
DB2	No additional configuration required.
Informix	The code page is set through the service template to a default value: CLIENT_LOCALE.
Oracle	Setting and determining the operating system code page is operating system-dependent. The operating system code page, the database code page, and the code page configured through the NLS_LANG environment variable must be compatible.
SQL Server	No additional configuration required.
Sybase	Set the connection option of the DataSourceSYBConnectOptions service attribute so that it refers to the same (or compatible) code page as the operating system code page. See "DataSourceSybConnectOptions" on page 538 for details.
JDBC Socket	No additional configuration required.
ODBC Socket	Refer to the documentation of the backend ODBC driver for information on passing the operating system code page to the driver as the application code page. On Linux and UNIX, configure the DataDirect Driver Manager with the correct IANAAppCodePage. You can use the ivcheckcp utility to determine the IANAAppCodePage setting. For information about using this utility, refer to the *SequeLink Troubleshooting Guide and Reference*.

Database Transliteration Mode

Database transliteration mode is used when connecting to Unicode databases or when the operating system and database code pages differ. To use Database transliteration mode, create a service using the service template with enhanced code page support, which automatically sets the service attributes to the correct settings for database transliteration mode.

See Table 12-3 for database-specific transliteration settings.

Table 12-3. Defining Database-Specific Settings for Database Transliteration Mode

Database	Configuration Requirements
DB2 for Linux/UNIX /Windows	No additional configuration required.
DB2 for z/OS	No additional configuration required.
Oracle	Configure the NLS_LANG environment variable to be compatible with the database code page.
SQL Server	No additional configuration required.
Sybase	No additional configuration required.
ODBC Socket	Refer to the documentation of the backend ODBC driver for settings that impact or enable the SQL-W function behavior. On Linux/UNIX, configure the DataDirect Driver Manager with the correct IANAAppCodePage. You can use the ivcheckcp utility to determine the IANAAppCodePage setting. For information about using this utility, refer to the *SequeLink Troubleshooting Guide and Reference*.
NOTE: Database transliteration mode is not supported with SequeLink Servers for Informix or JDBC Socket, and cannot be used with SequeLink Client for ADO.	

Configuring Database Transliteration Settings

The following sections provide database-specific guidelines and configuration examples for determining and setting transliteration settings.

General Guidelines for Transliteration

The methods for configuring transliteration are database-specific. This section provides general guidelines for transliteration.

DB2

1. Determine the code page of the DB2 database you want to connect to:

   `db2 get database configuration for database name`

   - If the database code page is 1208 (database code set is UTF-8), use the service template [SequeLink 6.0] DB2 for LUW service (enhanced code page support).
   - For all other database code pages, continue at Step 2.

2. Use the service template: [SequeLink 6.0] DB2 for LUW service.
3. If transliteration is required, set ServiceCodePage to OS. No additional configuration is required.

Microsoft SQL Server

Select either the SQL Server or SQL Server (enhanced code page) Service Template. See "Choosing a Service Template" on page 277 for more information.

For example, use the [SequeLink 6.0] SQL Server (enhanced code page support) Service Template if you want to support national character set data types in your application.

No additional configuration is required.

Oracle

See "Setting Up a SequeLink® Service for Oracle" on page 288 for an example of configuring an Oracle service.

To configure transliteration to support Unicode and national database character sets:

1. Create a service using a service template with enhanced code page support. See "Choosing a Service Template" on page 277 for more information.
2. Determine the code page of your Oracle database. For example, execute `SELECT` *`parameter`*`,` *`value`* `FROM nls_database_parameters` where *`parameter`*= 'NLS_CHARACTERSET'.
3. Set the value of the ServiceEnvironmentVariable service attribute to define the NLS_LANG environment variable to the code page of the Oracle database. For example, ServiceEnvironmentVariable=ITALIAN_WE8MSWIN1252.

Sybase

To configure transliteration in the Sybase environment, you must first identify the installed Sybase Server default character set (refer to the Sybase documentation for information about sp_default_charset). Then select the appropriate service template.

	If the character set is...	And the platform is...	The service template is...
UNIX	Not Unicode	Windows/ Linux/UNIX	[SequeLink 6.0] Sybase service
	Unicode (UTF-8)	Windows	[SequeLink 6.0] Sybase service (enhanced code page support) No additional configuration is required.
UNIX	Unicode (UTF-8)	Linux/UNIX	[SequeLink 6.0] Sybase service (enhanced code page support UTF-8) No additional configuration is required.
UNIX	Not Unicode or UTF-8	Linux/UNIX	[SequeLink 6.0] Sybase service (enhanced code page support UTF-16) No additional configuration is required.

If transliteration is required, set ServiceCodePage=OS and use the DataSourceSYBConnectOptions service attribute to specify the code page (see "DataSourceSybConnectOptions" on page 538).

Transliteration Scenarios

This section provides scenarios for setting up transliteration on Unicode and non-Unicode databases:

- "Setting Up a SequeLink® Service for DB2" on page 287
- "Setting Up a SequeLink® Service for Oracle" on page 288
- "Setting Up a SequeLink® Service for Microsoft SQL Server" on page 288
- "Setting Up a SequeLink® Service for Sybase" on page 289

Setting Up a SequeLink® Service for DB2

In this scenario, we set up a SequeLink service to access a DB2 Unicode database on Windows:

1. Check the database code page to confirm the database code set and database type:

   ```
   execute "get db cfg for accountdb"
   ```

Database territory	US
Database code page	1208
Database code set	UTF-8
Database country/region code	1

   The values for the database code set and code page indicate a DB2 Unicode database.

2. To add a service to the SequeLink Server for DB2, use the [SequeLink 6.0] DB2 UDB LUW service (enhanced code page support).

3. Start up the [SequeLink 6.0] DB2 UDB LUW service (enhanced code page support).

Setting Up a SequeLink® Service for Oracle

In this scenario, we set up a SequeLink service to access an Oracle9i Unicode database on Linux (en_US).

1. Determine the code page of the Oracle database. Execute

   ```
   SELECT parameter, value FROM nls_database_parameters
   WHERE parameter='NLS_CHARACTERSET'
   ```

2. Because our database is a Unicode database, add a service with service template [SequeLink 6.0]Oracle 10 service (enhanced code page support).

3. Set the NLS_LANG environment variable to the code page of your Oracle database. For example:

   ```
   ServiceEnvironmentVariable NLS_LANG=
   ENGLISH_UNITED KINGDOM.UTF8
   ```

4. Save the new settings in the SequeLink configuration file.

5. Start the SequeLink service for Oracle.

Setting Up a SequeLink® Service for Microsoft SQL Server

In this scenario, we set up a SequeLink service for SQL Server on Windows with English (United States). Because client and server operating systems have the same regional settings, no transliteration is required.

Add a service with service template [SequeLink 6.0] SQL Server service. The service template sets the ServiceCodePage service attribute to Default, which disables transliteration.

Setting Up a SequeLink® Service for Sybase

In this scenario, we set up a SequeLink service for Sybase on a Solaris (en_US) to access a Greek, non-Unicode, database.

1. Determine the installed Sybase server default character set (sp_default_charset). Execute:

   ```
   execute sp_default_charset
   ```

2. Because this is not a Unicode database, add a service with service template [SequeLink 6.0] Sybase service.

3. Set the ServiceCodePage service attribute to OS.

4. Look up the Sybase default character set in Table D-11 "Values for DataSourceSybConnectOptions" on page 539 to determine which DataSourceSYBConnectOptions values to use. For example:

   ```
   Sybase charset: iso88597
   codepage: ISO8859-7
   IANAAppCodePage: 10
   ```

   In this scenario, you add IANAAppCodePage=10 to the DataSourceSYBConnectOptions service attribute as follows:

   ```
   DataSourceSYBConnectOptions : IANAAppCodePage=10
   ```

5. Save the new settings in the SequeLink configuration file.

6. Start up the SequeLink Service for Sybase.

13 Configuring SequeLink® Security

This chapter offers an overview of the security options provided by SequeLink and describes how to configure SequeLink security for Linux, UNIX, and Windows and z/OS platforms. See Chapter 14 "Configuring the SequeLink® Proxy Server" on page 351 for information about configuring SequeLink Proxy Server, which provides certificate-based authentication on the Java platform.

About SequeLink® Security

SequeLink supports security mechanisms for the following purposes:

- Verification of a user by the SequeLink Server. The **Authentication** security mechanism allows the SequeLink Server to verify the identity of the user.
- Defining the types of requests that are accepted by the server. The **Authorization** security mechanism controls whether the user can send data access requests and administrative (SequeLink Manager) requests. Server configuration settings determine whether the server can accept the requests.

- Connection to a data store using the following security mechanisms:
  - **Data Store Logon** controls whether a user who is connected to the SequeLink Server can connect to the data store.
  - **Application IDs** control whether a client application can connect to the data store. This mechanism adds a layer of security on top of Data Store Logon.
  - **TCP/IP Location Filters** control whether a client application can connect to the data store based on the TCP/IP network identifier from which the connection request orginates.
  - **Terminal Security** is supported for connections to SequeLink Servers on z/OS. It controls whether the client application requesting access to the SequeLink data access service has permission to access it based on the TCP/IP address (terminal ID) originating the request.
- Defining the types of SQL statements accepted by the data store. The **ReadOnly** security mechanism controls whether the data store connection is read-only.
- The privacy of the data being transmitted. The **data privacy** security mechanism ensures that data transmitted between the client and server is kept private using data scrambling methods and encryption using SSL. SSL is supported for the ODBC Client, JDBC Client, and ADO Client.

Authentication

Authentication allows the SequeLink Server to verify the identity of the SequeLink Client when the client connects to the SequeLink Server. If authentication fails, the SequeLink Client disconnects from the server.

You must set an authentication method separately for users who send data access requests and users who send SequeLink Manager requests. For example, you may want to use an operating system user ID and password for administrative activities and Kerberos for data access activities.

Depending on the combination of client and server platforms involved in the connection, SequeLink supports the following authentication methods:

- **Anonymous**. The SequeLink Server accepts connections from any SequeLink Client without verifying the client's identity.
- **Operating system user ID** and **password**. The SequeLink Server verifies the identity of the SequeLink Client using a user ID and password that must be valid for the platform on which the SequeLink Server is running. If verified, the server accepts the user ID as the identity of the client and permits the connection.
- **Kerberos**. Kerberos authentication uses Kerberos, a trusted third-party authentication service, to verify user identities. Kerberos authentication can take advantage of the user name and password maintained by the operating system to authenticate users to the database. This method requires knowledge of how to configure your Kerberos environment.
- **Integrated NT**. This option is supported for connections between SequeLink Server for Windows servers and ODBC Clients, ADO Clients, and .NET Clients on Windows only. The SequeLink Server verifies the identity of the SequeLink Client using the client's Windows network logon credentials instead of a Windows user ID and password.

Although a user may be able to connect to the SequeLink Server, the user does not automatically have access to the database that the SequeLink Server services.

For instructions on configuring authentication:

- On Linux, UNIX, and Windows, see "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.
- On z/OS, see "Configuring SequeLink® Security for z/OS" on page 322.

Authorization

After the SequeLink Server has authenticated the client, SequeLink verifies that the client is authorized to perform data access activities or SequeLink Manager activities. SequeLink supports authorization for data access requests and for SequeLink Manager requests. You configure the authorization for the two types of requests separately. Authorization options depend on your SequeLink Server platform.

For instructions on configuring authorization:

- On Linux, UNIX, and Windows, see "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.
- On z/OS, see "Configuring SequeLink® Security for z/OS" on page 322.

Data Store Logon

Once a connection is established, authentication is complete, and the type of requests the server will accept has been established, a connection from the SequeLink Server to the data store can be established using either of the following methods:

- Specifying data store logon information (a valid DBMS user ID and password). This is the default for Windows and UNIX.
- Allowing the database to inherit the logon user ID that was established during the authentication process. For example, if

using Kerberos authentication, use this method to allow Kerberos to use the operating system user ID and password.This method must be used for z/OS, but can be used for Linux, UNIX, and Windows also.

For details on configuring data store logon, see "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.

Application IDs

Application IDs are alphanumeric strings passed by a SequeLink Client that identify the client application to a SequeLink service that has been configured to accept connections only from specific application IDs.

Application IDs add another layer of security for the connection to the data store beyond that provided by the Data Store Logon security mechanism. Data Store Logon allows all users of client applications to access the data store if they meet the qualifications set by Data Store Logon. Using application IDs, you can restrict connections to the data store to only those client applications that identify themselves to the SequeLink Server through an application ID.

On Windows platforms, application IDs can be specified explicitly by the client application or they can be automatically generated by the ODBC Client or the ADO Client. The advantage of using application IDs generated by the ODBC Client or ADO Client is the application itself does not need to contain the application ID; however, you must specify in the client application that you want to turn on the automatic generation of application IDs. The application ID is generated using the sha-1 hashing algorithm, resulting in a 160-bit hash value.

See "Using Application IDs to Restrict User Access" on page 335 for more information.

TCP/IP Location Filters

Using TCP/IP network identifiers, such as TCP/IP host names (for example, burner.ddtek.com) or a range of TCP/IP addresses (for example, 192.16.*.*), TCP/IP location filters allow you to specify which clients can access a SequeLink data access service or SequeLink agent service.

When you create a filter, the IP address can contain wild card characters to indicate that any decimal number in that location is considered valid. Use ? for a single number or * for multiple numbers. Use \ as the escape character.

The following address and name formats are supported:

Client TCP/IP host name	`burner.ddtek.com`
Client TCP/IP domain names through the use of a wild card	■ `192.16.2.*` ■ `*.ddtek.com` ■ `belg?.progress.com`
Client TCP/IP address	`127.0.45.1??`
Client TCP/IP address range	`192.16.*.*`

For more information about configuring TCP/IP location filters, see "Configuring TCP/IP Location Filters" on page 342.

Terminal Security

z/OS When terminal security is enabled, through activating the RACF TERMINAL security class, the SequeLink Server verifies that the client application requesting access to the SequeLink data access service has permission to access it based on the TCP/IP address (terminal ID) originating the request. You can use terminal security to make sure that:

- Only specific TCP/IP addresses can be used by specific users to connect to the SequeLink Server.

- Only specific groups of users can use specific TCP/IP addresses to connect to the SequeLink Server. For example, you may want to make sure that a user ID associated with an application running on an application server can only log on to the SequeLink Server from a specific TCP/IP address.

See "Activating Terminal Security" on page 325 for more information about activating terminal security.

Terminal security is controlled by activating the RACF TERMINAL security class instead of setting a service attribute.

ReadOnly

SequeLink allows you to configure the types of SQL statements the data store connection will accept:

- Select statements only (makes the connection read-only)
- Select statements and Stored Procedures
- All SQL statements
- Read-only settings of the database

The service attribute that controls this functionality is DataSourceReadOnly.

Data Privacy

SequeLink provides data scrambling to ensure the privacy of data. In addition, you can use data encryption to provide a more secure transmission of data across the network. For example, you may want to use data encryption in the following scenarios:

- You have offices that share confidential information over an intranet.
- You send sensitive data, such as credit card numbers, over a database connection.

- You need to comply with government or industry privacy and security requirements.

NOTE: Data encryption may adversely affect performance because of the additional overhead (mainly CPU usage) required to encrypt and decrypt data.

Data Scrambling

Data scrambling ensures that no cleartext messages are transmitted between the client and server over the network. SequeLink provides the following implementations of data scrambling:

- **Fixed-key DES** operates using a 56-bit key.
- **Fixed-key 3DES** operates using a 168-bit key.
- **Byte swapping** means that bytes of data are randomly swapped to scramble data. Different encoded mappings are used for different sessions.

Data scrambling does not provide the same level of security as data encryption and is not enabled by default.

NOTE: Even if you choose not to use a data scrambling method, user IDs and passwords are *never* sent as cleartext.

See "Configuring Data Privacy" on page 313 for more information about configuring data scrambling.

Data Encryption

Secure Sockets Layer (SSL) is an industry-standard protocol for sending encrypted data over database connections. SSL secures the integrity of your data by encrypting information and providing client/server authentication.

SequeLink supports SSL for the following types of data transfers:

- Between a SequeLink Client and a SequeLink Server on Linux/UNIX/Windows. SequeLink uses SSL for data encryption. For an SSL connection to be successful, both the SequeLink Server and SequeLink Client must be configured for SSL encryption. If a SequeLink Client that is not configured for SSL attempts to connect to a SequeLink Server configured for SSL, the SequeLink Server rejects the connection request and returns the following error message: `TCP/IP, connection reset by peer`.
- Between a SequeLink Client *for* JDBC and the SequeLink Proxy Server. SequeLink uses SSL for data encryption and authentication.

See "Configuring Data Privacy" on page 313 for details on configuring SSL for data transfers between a SequeLink Client and SequeLink Server. See Chapter 14 "Configuring the SequeLink® Proxy Server" on page 351 for details on configuring SSL for data transfers between the SequeLink Client *for* JDBC and the SequeLink Proxy Server.

Configuring SequeLink® Security on Linux, UNIX, and Windows

This section describes how to configure authentication, authorization, and data store logon for SequeLink for Linux, UNIX, and Windows.

Configuring Authentication for Data Access Activities

Set the ServiceAuthMethods service attribute to one or more of the following values:

- **Anonymous authentication**: ServiceAuthMethods=anonymous.
- **User ID and password authentication**: ServiceAuthMethods=OSLogon(*HUID*,*HPWD*), ServiceAuthMethods=OSLogon(*UID*,*PWD*)
- **User ID and password authentication**: ServiceAuthMethods=OSLogon(*UID*,*PWD*,*NPWD*), or ServiceAuthMethods=OSLogon(*HUID*,*HPWD*,*NPWD*).

  NOTE: The NPWD parameter of OSLogon allows you to change the password.
- **Kerberos authentication**: ServiceAuthMethods=kerberos. See "Configuring Kerberos Authentication" on page 303 for details.
- **Integrated NT authentication**: ServiceAuthMethods=integrated_nt (Windows only).

You can configure multiple authentication mechanisms for data access activities. For example, if you configure ServiceAuthMethods=integrated_nt and

ServiceAuthMethods=OSLogon(*UID*,*PWD*), the SequeLink Server uses the Integrated NT authentication mechanism and does not require an ODBC or ADO Client connecting from a Windows workstation to provide user ID and password information. SequeLink Clients on Linux/UNIX or a JDBC application still must provide a valid Windows user ID and password.

Configuring Authentication for Administrative Activities

To configure authentication for a SequeLink agent service, set the ServiceAdminAuthMethods service attribute to one or more of the following values:

- **Anonymous authentication**: ServiceAdminAuthMethods=anonymous.
- **User ID and password authentication**: ServiceAdminAuthMethods=OSLogon(*UID*,*PWD*).
- **Kerberos authentication**: ServiceAdminAuthMethods=kerberos. See "Configuring Kerberos Authentication" on page 303 for details.
- **Integrated NT authentication**: ServiceAdminAuthMethods=integrated_nt.

For SequeLink services on Windows, you can configure multiple authentication mechanisms for administrative activities. For example, if you configure ServiceAdminAuthMethods=integrated_nt and ServiceAdminAuthMethods=OSLogon(*UID*,*PWD*), the SequeLink Server uses the Integrated NT authentication mechanism and does not require a SequeLink Manager running on a Windows workstation to provide user ID and password information. SequeLink Clients on Linux/UNIX still must provide a valid Windows user ID and password.

Configuring Single Sign-On Security

The combination of security features provided by SequeLink and the security provisions offered by the DBMS and the Linux, UNIX, and Windows operating systems allows you to configure a single sign-on environment for ODBC, JDBC, and ADO applications. Users log on to the security system on the network and can connect to the DBMS using their operating system credentials (if allowed by the DBMS security configuration).

Kerberos authentication provides single sign-on for ODBC, ADO, and JDBC applications on Linux, UNIX, and Windows.
NT Integrated authentication provides single-sign on for ODBC and ADO applications on Windows only.

To configure single sign-on security:

1. Configure your DBMS security as described in "Configuring SequeLink® Security on Linux, UNIX, and Windows" on page 300.

   For instructions on how to configure your DBMS security, refer to your database documentation.

2. Grant the required database access rights to users using the appropriate data store provided by your DBMS. For instructions on how to grant database access rights, refer to your database documentation.

3. Choose one of the following authentication methods:

   - Kerberos. See "Configuring Kerberos Authentication" on page 303 for details on configuring Kerberos authentication.

   - NT Integrated authentication. Set the ServiceAuthMethods or ServiceAdminAuthMethods service attribute to a value of integrated_nt.

4. Restart the SequeLink Server.

NOTE: If SequeLink Client *for* ADO.NET clients will be connecting to this service, add another ServiceAuthMethods attribute, and specify `ServiceAuthMethods=OSLogon(HUID,HPWD)` or `ServiceAuthMethods=OSLogon(UID,PWD)`. These SequeLink Clients must provide a valid server user ID and password.

Configuring Kerberos Authentication

Verify that your environment meets the requirements listed in Table 13-1 before you configure SequeLink for Kerberos authentication.

Table 13-1. Kerberos Authentication Requirements

Component	Requirements
Database server	The database server must be running one of the following databases: **DB2**: ■ DB2 v8.1 or higher for Linux/UNIX/Windows **Oracle**: ■ Oracle 10*g* R2 ■ Oracle 9i R2 **Microsoft SQL Server**: ■ Microsoft SQL Server 2005 ■ Microsoft SQL Server 2000 In addition, your database security must be configured to accept operating system user IDs. See "Configuring Database Security" on page 304 for more information.
Kerberos server	The Kerberos server is the machine where the user IDs for authentication are administered. The Kerberos server is also the location of the Kerberos KDC.

Configuring Database Security

To use single-sign on with Kerberos authentication, your database security must be configured to accept operating system user IDs:

- For Microsoft SQL Server, configure the DBMS security for integrated security.
- For Oracle, configure the DBMS security to allow external authentication (formerly OPS$functionality).
- For DB2, configure your database with authentication type SERVER. Credentials are validated on the database server.

For instructions on how to configure your DBMS security, refer to your database documentation.

Kerberos and Operating System User ID Mapping

A Kerberos user principal has the form:

name@*REALM*

To allow single-sign on for Kerberos in a SequeLink environment, the user principal name must be mapped to the operating system user ID. For example:

- On Linux and UNIX, the "name" part of the principal must match the operating system user ID.
- On Windows, the "name" part and the "REALM" part of the principal matches the operating system user ID using the format *REALM**name* where *REALM* is the Windows domain name.

As a result, on Linux and UNIX, the database administrator must be aware that principals from different Kerberos realms, but with the same name, are mapped to the same operating system user ID. For example, the following principals are mapped to the same user ID:

```
user_A/@UVW.COM
```

and

```
user_A/@XYZ.COM
```

Configuring Kerberos Authentication for SequeLink® Server

1. In your SequeLink data access service or your SequeLink Agent service, set ServiceAuthMethods attribute to a value of kerberos.
2. In the server data source, set the DataSourceLogonMethod attribute to a value of OSIntegrated. This allows the SequeLink service to authenticate the client using the operating system user ID and password.

If you installed the SequeLink Server on Windows, skip to Step 4.

3. If you installed the SequeLink Server on Linux and UNIX, modify the krb5.conf file to contain your Kerberos realm name and the KDC name for that Kerberos realm. Modify the krb5.conf file by editing the file with a text editor.

   NOTE: In Windows Active Directory, the Kerberos realm name is the Windows domain name and the KDC name is the Windows domain controller name.

   For example, if your Kerberos realm name is XYZ.COM and your KDC name is kdc1, your krb5.conf file would look like this:

```
[libdefaults]
        default_realm = XYZ.COM
```

```
        default_keytab_name = FILE:/install_dir/lib/krb5/krb5.keytab
        default_tkt_enctypes = des3-cbc-sha1 des-cbc-md5 des-cbc-crc
        default_tgs_enctypes = des3-cbc-sha1 des-cbc-md5 des-cbc-crc

[realms]
        XYZ.COM = {
             kdc = kdc1:88
             admin_server = server1.eu.xyz.com:749
             default_domain = kdc1
        };

[domain_realm]
        .xyz.com = XYZ.COM
        .eu.xyz.com = XYZ.COM

[logging]
        kdc = FILE:/var/krb5/log/krb5kdc.log
        admin_server = FILE:/var/krb5/log/kadmin.log
        default = FILE:/var/krb5/log/krb5lib.log
```

4 (Optional) If you want to use a service principal name other than the default name to identify the SequeLink service, specify a value for the ServiceKerberosPrincipalName attribute for data access services or the ServiceAdminKerberosPrincipalName attribute for SequeLink Agent services. See "ServiceKerberosPrincipalName" on page 579 and "ServiceAdminKerberosPrincipalName" on page 561 for details.

Configuring Kerberos Authentication for the ODBC Client and ADO Client

If the SequeLink Client is installed on Windows, the Kerberos realm and KDC name for that realm are automatically detected by the Windows operating system.

NOTE: In Windows Active Directory, the Kerberos realm name is the Windows domain name and the KDC name is the Windows domain controller name.

UNIX

If the SequeLink Client is installed on UNIX or Linux, you must modify a Kerberos configuration file named krb5.conf to specify values for the Kerberos realm and the KDC name for that realm. Typically, this file is installed by the Kerberos implementation provided by your UNIX or Linux operating system and is located in the /etc subdirectory of your UNIX or Linux machine.

For example, if your Kerberos realm name is XYZ.COM and your KDC name is kdc1, your krb5.conf file would look like this:

```
[libdefaults]
        default_realm = XYZ.COM
        default_keytab_name = FILE:/etc/krb5/krb5.keytab
        default_tkt_enctypes = des3-cbc-sha1 des-cbc-md5 des-cbc-crc
        default_tgs_enctypes = des3-cbc-sha1 des-cbc-md5 des-cbc-crc

[realms]
        XYZ.COM = {
             kdc = kdc1:88
             admin_server = server1.eu.xyz.com:749
             default_domain = kdc1
        };

[domain_realm]
        .xyz.com = XYZ.COM
        .eu.xyz.com = XYZ.COM

[logging]
        kdc = FILE:/var/krb5/log/krb5kdc.log
        admin_server = FILE:/var/krb5/log/kadmin.log
        default = FILE:/var/krb5/log/krb5lib.log
```

Configuring Kerberos Authentication for the JDBC Client

During installation, the JDBC Client installs the following files required for Kerberos authentication in the /lib subdirectory of the SequeLink Client installation directory:

- JDBCDriverLogin.conf file is a configuration file that specifies which Java Authentication and Authorization Service (JAAS) login module to use for Kerberos authentication. This file is configured to load automatically unless the java.security.auth.login.config system property is set to load another configuration file.

  NOTE: Do not modify this file.

- krb5.conf is a Kerberos configuration file containing values for the Kerberos realm and the KDC name for that realm. The JDBC Client installs a generic file that you must modify for your environment.

  If the krb5.conf file does not contain a valid Kerberos realm name and KDC name, the following exception is thrown:

  ```
  Message:[DataDirect][SequeLink JDBC Driver]Could not
  establish a connection using integrated security: No
  valid credentials provided

  [DataDirect][SequeLink JDBC Driver]A username was not
  specified and the driver was unable to establish a
  connection using integrated security.
  ```

The krb5.conf file installed by the JDBC Client is configured to load automatically unless the java.security.krb5.conf system property is set to point to another Kerberos configuration file.

Obtaining a Kerberos Ticket Granting Ticket

To use Kerberos authentication, the application user first must obtain a Kerberos Ticket Granting Ticket (TGT) from the Kerberos server. The Kerberos server verifies the identity of the user and controls access to services using the credentials contained in the TGT.

If the application uses Kerberos authentication from a Windows client, the application user is not required to log onto the Kerberos server and explicitly obtain a TGT. Windows Active Directory automatically obtains a TGT for the user.

UNIX

If an application uses Kerberos authentication from a UNIX or Linux client, the user must log onto the Kerberos server using the kinit command to obtain a TGT. For example, the following command requests a TGT from the server with a lifetime of 10 hours, which is renewable for 5 days:

```
kinit -l 10h -r 5d user
```

where *user* is the application user.

Refer to your Kerberos documentation for more information about using the kinit command and obtaining TGTs for users.

Configuring Authorization for a SequeLink® Data Access Service

To configure user authorization for a SequeLink data access service, set the ServiceUser attribute, or, if you want to configure user authorization for user groups defined on Windows or UNIX, set the ServiceUserGroup attribute.

Configuring ServiceUser

To configure user authorization, set the ServiceUser attribute:

- ServiceUser=*user_ID*, where *user_ID* is the user ID of each user who is allowed to access data using this service. To configure authorization for multiple users, you must set this attribute multiple times, one instance for each user. For example:

  ```
  ServiceUser=RSMITH
  ServiceUser=DJONES
  ServiceUser=TCONRAD
  ```

  NOTE: On Windows servers, you must prefix the user ID with the Windows server name or the Windows domain name, for example, SALES\DJONES. When connecting, the user must also prefix the user ID with the Windows server name, if connecting to a local server, or the Windows domain name.

- ServiceUser=authenticated. Any user who can provide a valid host user ID and password or who uses the Integrated NT authentication process will be allowed to send data access requests to the data access service.

- ServiceUser=everyone. All connections will receive the user authorization level, regardless of how they are authenticated. If ServiceAuthMethods=anonymous, ServiceUser=everyone *must* be specified.

Configuring ServiceUserGroup

To configure user authorization for user groups defined on a Linux, UNIX, and Windows server, set the ServiceUserGroup attribute. Specify ServiceUserGroup=*user_group*, where *user_group* is a valid user group defined on Linux, UNIX, and Windows. To configure user authorization for multiple user groups, you must set this attribute multiple times, one instance for each user. For example:

```
ServiceUserGroup=SLUSERG1
ServiceUserGroup=SLUSERG2
ServiceUserGroup=SLUSERG3
```

NOTE: On Windows servers, you must prefix the user group ID with the Windows server name or the Windows domain name where the user group is defined, for example, SALES\SLUSERG1. When connecting, the user must also prefix the administrator ID with the Windows server name, if connecting to a local server, or the Windows domain name.

Configuring Authorization for a SequeLink® Agent Service

To configure administration authorization, set the ServiceAdministrator service attribute, or, if you want to configure authorization for user groups defined on Linux, UNIX, and Windows, set the ServiceAdministratorGroup service attribute.

NOTE: On Windows, each user who is allowed to make SequeLink Manager requests must have administrator rights.

Configuring ServiceAdministrator

To configure administration authorization, set the ServiceAdministrator attribute:

- ServiceAdministrator=*user_ID*, where *user_ID* is the user ID of each user who is allowed to make SequeLink Manager requests. To configure administration authorization for multiple users, you must set this attribute multiple times, one instance for each user. For example:

  ```
  ServiceAdministrator=RSMITH
  ServiceAdministrator=DJONES
  ServiceAdministrator=TCONRAD
  ```

  NOTE: On Windows servers, you must prefix the user ID with the Windows server name or the Windows domain name, for example, SALES\DJONES. When connecting, the user must also prefix the user ID with the Windows server name, if connecting to a local server, or the Windows domain name.

- ServiceAdministrator=authenticated. Any user who can provide a valid host user ID and password or who uses the Integrated NT authentication process will receive the same administration authorization.

- ServiceAdministrator=everyone. All connections will receive the same administration authorization, regardless of how they are authenticated. If ServiceAdminAuthMethods=anonymous, you must specify ServiceAdministrator=everyone.

Configuring ServiceAdministratorGroup

To configure authorization for user groups defined on Windows and UNIX, set the ServiceAdministratorGroup attribute. Specify ServiceAdministratorGroup=*user_group*, where *user_group* is a valid user group defined on Windows or UNIX. To configure administration authorization for multiple user groups, you must

set this attribute multiple times, one instance for each user. For example:

```
ServiceAdministratorGroup=SLUSERG1
ServiceAdministratorGroup=SLUSERG2
ServiceAdministratorGroup=SLUSERG3
```

NOTE: On Windows servers, you must prefix the user group ID with the Windows server name or the Windows domain name where the group is defined, for example, SALES\SLUSERG1. When connecting, the user must also prefix the user group with the Windows server name, if connecting to a local server, or the Windows domain name.

Configuring Data Privacy

To ensure privacy of data, SequeLink provides data scrambling (all SequeLink environments) and "real" encryption through the use of SSL.

Configuring Data Scrambling

To configure SequeLink to use DES, 3DES, or byteswap, you must set the ServiceEncryptionAlgorithm service attribute, for example, ServiceEncryptionAlgorithm=DES. Data scrambling is not enabled by default, which means that if you do not configure data scrambling, cleartext messages are transmitted between the client and server over the network.

Configuring SSL Encryption

SequeLink supports SSL for the following types of data transfers:

- Between a SequeLink Client and a SequeLink Server. SequeLink uses SSL for data encryption.
- Between a JDBC Client and the SequeLink Proxy Server. SequeLink uses SSL for data encryption and authentication.

The configuration process is different for each data transfer type.

Data Transfers Between SequeLink Client and SequeLink Server

By default, when SSL encryption is enabled, SSL uses the TLS_DH_anon_WITH_AES_128_CBC_SHA cipher suite for data encryption. If you want to change the cipher suite used for SSL, set a value for the ServiceSSLCipherSuites service attribute. See "ServiceSSLCipherSuites" on page 584 for a description of valid values.

Because creating an SSL session requires CPU-intensive computations, SSL session reuse results in a relatively large performance gain over setting up completely new security sessions for each connection. Cached sessions allow the client to reuse a session in a subsequent connection. The ServiceSSLSessionCacheSize service attribute controls how many entries can be cached. Set this attribute to a number less than or equal to the maximum connections setting for the server. The ServiceSSLSessionCacheTimeout value specifies how long cached session IDs remain valid.

To configure SSL data encryption:

1. Enable SSL encryption on the SequeLink Server. Set the ServiceSSLEnabled service attribute to true.
2. Enable SSL encryption on the SequeLink Client:
   - JDBC Client: Using a connection URL, specify a value of 1 for the Encrypted connection property. For example:

     ```
     jdbc:sequelink://localhost:16007;DatabaseName=odbc;
     Encrypted=1
     ```

   - ODBC Client: Using the ODBC Administrator, on the General tab, select the **Encrypted (SSL)** check box. Alternatively, you can use the Encrypted connection attribute in the connection string. For example:

     ```
     DSN=sqlnk600_sqlsrv;SDSN=Default;HST=localhost;PRT=1
     6007;DB=odbc;UID=odbc;PWD=sqlnk001;Encrypted=1
     ```

   - ADO Client: On the DataDirect SequeLink for ADO 6.0 Provider dialog box, on the General tab, select the **Encrypted (SSL)** check box. Alternatively, you can use the Encrypted connection attribute in the provider string. For example:

     ```
     DSN=sqlnk600_sqlsrv;SDSN=Default;HST=localhost;PRT=1
     6007;DB=odbc;UID=odbc;PWD=sqlnk001;Encrypted=1
     ```

Alternatively, if you are using LDAP in your SequeLink environment, you can configure SSL encryption as shown in the following examples. The LDAP attribute used to configure encryption depends on whether you are configuring an LDAP entry for the JDBC, ODBC, or ADO client.

Example 1: LDAP Entry for the JDBC Client

Configure SSL encryption using the JavaReferenceAddress attribute (`javaReferenceAddress: #5#encrypted#1`).

```
dn: cn=DB2V8 on server1,ou=JVS,ou=Arne,ou=USERS,o=Development
cn: DB2V8 on server1
objectClass: top
objectClass: javaContainer
objectClass: javaObject
objectClass: javaNamingReference
javaReferenceAddress: #0#description#
javaReferenceAddress: #1#portNumber#6007
javaReferenceAddress: #2#serverName#server1
javaReferenceAddress: #3#databaseName#jvs
javaReferenceAddress: #4#serverDataSource#Default
javaReferenceAddress: #5#encrypted#1
javaFactory: com.ddtek.jdbcx.sequeLink.SequeLinkDataSourceFactory
javaClassName: com.ddtek.jdbcx.sequelink.SequeLinkDataSource
```

Example 2: LDAP Entry for the ODBC Client or the ADO Client

Configure SSL encryption using the SequelinkEncrypted attribute (`SequeLinkEncrypted: 1`).

```
# usercn.ldiff template
# USER     Arne
# DATABASE SequeLink
# HOST     belg-keidis
# PORT     6007
#
dn: cn=DB2V8 on belg-keidis,ou=OVS,ou=Arne,ou=USERS,o=Development
cn: DB2V8 on belg-keidis
objectClass: top
objectClass: SequeLinkODBCDSN
SequeLinkPort: 6007
SequeLinkHost: belg-keidis
SequeLinkServerDatasource: Default
SequeLinkEncrypted: 1
```

Data Transfers Between the JDBC Client and the SequeLink Proxy Server

For details on configuring SSL encryption for data transfers between the JDBC Client and the SequeLink Proxy Server, see Chapter 14 "Configuring the SequeLink® Proxy Server" on page 351.

Configuring Data Store Logon

A client application establishing a connection to the database must provide a valid DBMS user ID and password when DataSourceLogonMethod=DBMSLogon(*UID*,*PWD*) or DataSourceLogonMethod=DBMSLogon(*DBUID*,*DBPWD*).

Depending on how the SequeLink service is configured, the SequeLink Server may require the SequeLink Client to provide two user IDs and passwords. SequeLink Clients typically provide user ID and password information using the UID and PWD attributes in a connection string (ODBC and ADO) or a connection URL (JDBC).

To avoid possible conflict with a standard keyword pair (UID,PWD) with two sets of values, make sure that you set non-conflicting values for the ServiceAuthMethods and DataSourceLogonMethod attributes. For example, when ServiceAuthMethods=OSLogon(*UID*,*PWD*) and DataSourceLogonMethod=DBMSLogon(*DBUID*,*DBPWD*), the SequeLink Client must provide the operating system user and password using the keywords UID and PWD and the database user and password must be specified using the DBUID and DBPWD keywords.

To allow the DBMS to inherit the operating system (or network) user ID and password for data store authorization, specify DataSourceLogonMethod=OSIntegrated. For example, you may want to specify DataSourceLogonMethod=OSIntegrated to

allow Kerberos to use the operating system user Id and password for authentication.

NOTE: Do not use this method when ServiceAuthMethods=anonymous.

Some databases allow anonymous access, that is, access without any user credentials. To allow anonymous access to the database, specify DataSourceLogonMethod=anonymous.

Storing Client Credentials

The credentials for the ODBC Client can be stored in the registry on Windows or in the odbc.ini on Linux/UNIX. This allows the user to logon without specifying credentials.

Storing Credentials in the Windows Registry

The following keys are added:

- For a user data source:

```
HKEY_CURRENT_USER\Software\ODBC\ODBC.INI\<dsname>\LogonId
HKEY_CURRENT_USER\Software\ODBC\ODBC.INI\<dsname>\Password
```

- For a system data source:

```
HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBC.INI\<dsname>\LogonId
HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBC.INI\<dsname>\Password
```

The slencpwd.exe utility is installed in the Client's installation directory. This utility allows the customer to enter an encrypted password in the registry and optionally add a LogonId. The LogonId is stored a plain text, the Password is encrypted. Note that the credentials are always sent over the wire in encrypted form.

```
C:\>slencpwd
SequeLink Client for ODBC (tm) Utility to store an encrypted password in a
Data Source
Choose System DSN [S] or User DSN [U].
Choice: [U] ==> U
Enter Data Source Name: sqlnk600_mydsn

Choose which password to encrypt and store.
[1]=Password (PWD)
[2]=DBPassword (DBPWD)
[3]=HPassword (HPWD)
[4]=NewPassword (NPWD)
Choice: [1] ==> 1

Enter Password (PWD) (keystrokes are not shown):
DSN=sqlnk600_mydsn : Password (PWD) stored encrypted.

Would you like to store a LogonID (UID) as well? [Y] [N]
Choice: ==> Y
Enter LogonID (UID):odbc

DSN=sqlnk600_mydsn : LogonID (UID) stored.
```

UNIX *Storing Client Credentials in the odbc.ini file*

The slencpwd utility is installed in the tools directory of the driver's installation directory. This utility allows you to enter an encrypted password in the odbc.ini and, optionally, to add a LogonId.

The following keys are added to the odbc.ini:

```
LogonId=my_logonid
Password=my_password_in_encrypted_format
```

The LogonId is stored plain text while the Password is encrypted. Note that the credentials are always sent over the wire in encrypted form.

```
$ ./slencpwd
SequeLink Client for ODBC (tm) Utility to store an encrypted password in a
Data Source
Enter Data Source Name: sqlnk600_mydsn
Choose which password to encrypt and store.
[1]=Password (PWD)
[2]=DBPassword (DBPWD)
[3]=HPassword (HPWD)
[4]=NewPassword (NPWD)
Choice: [1] ==> 1
Enter Password (PWD) (keystrokes are not shown):
DSN=sqlnk600_mydsn : Password (PWD) stored encrypted.
Would you like to store a LogonID (UID) as well? [Y] [N]
Choice: ==> Y
Enter LogonID (UID):my_logonid
DSN=sqlnk600_mydsn : LogonID (UID) stored.
```

SequeLink® Security Attribute Defaults for Linux, UNIX, and Windows

This section lists the installation defaults for SequeLink's security attributes for Linux, UNIX, and Windows and describes the effect each combination of settings has on security.

Defaults for a SequeLink® Agent Service

`ServiceAdminAuthMethods=OSlogon(`*`UID`*`,`*`PWD`*`)`
`ServiceAdministrator=`*`SequeLink_administrator`*

The combination of defaults for these attributes means that only the person who logs on using the user ID that was entered when the SequeLink Server software was installed is allowed to manage the SequeLink environment. The SequeLink Server installer prompts for a user ID for the SequeLink administrator

when you install the SequeLink Server. On Windows, the SequeLink administrator must have administrator rights.

`ServiceEncryptionAlgorithm=none`

The default for this attribute means that cleartext messages are transmitted between the client and server. Note that user IDs and passwords are *never* sent as cleartext.

Defaults for a Data Access Service

```
ServiceAuthMethods=anonymous
ServiceUser=everyone
DataSourceLogonMethod=DBMSLogon(UID,PWD)
DataSourceReadOnly=No
```

The combination of defaults for these attributes means that anyone who can provide a valid DBMS user name and password will be allowed to access the database using this service. The database connection accepts all types of SQL statements. Once connected to the DBMS, the database security system will guarantee that the user can only perform actions that are allowed by the database administrator.

`ServiceEncryptionAlgorithm=none`

The default for this attribute means that cleartext messages are transmitted between the client and server. Note that user IDs and passwords are never sent as cleartext.

Configuring SequeLink® Security for z/OS

z/OS This section describes how to configure SequeLink authentication, authorization, and data store logon, and how to activate terminal security for z/OS. The security of the SequeLink Server for z/OS is integrated with the z/OS security system using the SAF interface.

Configuring Authentication

On z/OS, SequeLink supports Kerberos, in addition to the anonymous and user ID and password authentication mechanisms:

- For the user ID and password authentication mechanism, the client application must provide a valid user ID and password for the platform on which the SequeLink Server is running. The server verifies the user ID and password with the security package installed on z/OS. If verified, the server accepts the user ID as the identity of the client. When a password change is required (for example, when a password expires), the client application must also provide a new password.
- For anonymous authentication, a UID map with a generic entry is needed. See "Using UID Mapping" on page 329 for more information.
- For Kerberos authentication, the SequeLink service principal name is retrieved from the Kerberos segment of the SequeLink Server user ID defined in your security manager, for example, RACF.

Configuring Authentication for Data Access Activities on z/OS

To configure authentication for data access activities, set the ServiceAuthMethods attribute to anonymous or OSLogon:

```
ServiceAuthMethods=anonymous
```

or

```
ServiceAuthMethods=OSLogon(UID,PWD,NPWD)
```

Configuring Authentication for Administrative Activities on z/OS

To configure authentication of administrative activities, set the ServiceAdminAuthMethods attribute to the following values:

```
ServiceAdminAuthMethods=OSLogon(UID,PWD)
```

Configuring Authorization

To configure resource-based authorization for z/OS, enable the authorization setting for your service. You can also specify a security class and a security resource within this security class.

How you configure authorization depends on whether you are configuring it for data access activities, administrative activities, or enabling RACF resource security for a server data source.

Configuring Authorization for Data Access Activities at the Service or Data Source Level

To configure resource-based authorization for z/OS, set either MVSServiceAuthorizationEnable or MVSDataSourceAuthorizationEnable to True. The SequeLink

Server validates the SequeLink Client's identity using the client's user ID and password and the client's authority to access the service or data source. If one of these options is used, you may specify a security resource and a security class. Also, any user that requires access to this SequeLink service or data source must be granted READ access to the specified resource defined in this class.

To configure resource-based authorization for data access activities, set the following attributes to the following values:

- MVSServiceAuthorizationEnable=True for authorization on the service level or MVSDataSourceAuthorizationEnable=True for authorization on the data source level
- MVSServiceAuthorizationClass=*sec_class_name*, where *sec_class_name* is the name of the security class where the MVSServiceAuthorizationResource is defined. The default value is FACILITY.
- MVSServiceAuthorizationResource=*sec_resource_name*, where *sec_resource_name* is the name of the security resource where access is defined for your users. The default value is the name of the SequeLink data access service.

  For data source authorization, the data source name is appended to the name of the security resource, for example, *sec_resource_name.data_source_name*.

Configuring Authorization for Administrative Activities

To configure resource-based authorization for administrative activities on z/OS, set MVSServiceAdminAuthorizationEnable to True. The SequeLink Server validates the SequeLink Client's identity using the client's user ID and password, and the client's authority to access the service. If this option is used, you also may specify a security resource and security class. Also, any user that

requires access to this SequeLink service must be granted READ access to the specified resource defined in this class.

To configure resource-based authorization for administrative activities, set the following attributes to the following values:

- Set MVSServiceAdminAuthorizationEnable to TRUE.
- MVSServiceAdminAuthorizationResource=*server_name*, where *server_name* is the name of the SequeLink Server, or, on z/OS, the service name of the SequeLink agent.
- MVSServiceAdminAuthorizationClass=*sec_class_name*, where *sec_class_name* is the name of the security class where the MVSServiceAdminAuthorizationResource is defined. The default value is FACILITY.

Configuring Data Store Logon

To allow the DBMS to inherit the operating system (or network) user identification to use for database authorization, set the DataSourceLogonMethod attribute to the following value:

```
DataSourceLogonMethod=OSIntegrated
```

Activating Terminal Security

When terminal security is enabled and the RACF TERMINAL security class is activated, SequeLink Server verifies that the TCP/IP address (terminal ID) requesting the connection has permission to connect to the SequeLink Server. Even when the TERMINAL security class is not activated, a RACF (or equivalent) message informs you of the TCP/IP address of each user or application requesting a connection to the SequeLink Server.

Each TCP/IP address originating the client connection to SequeLink Server has an equivalent terminal ID used by the SAF security interface. SequeLink encodes each part of a TCP/IP

address into its hexadecimal equivalent, and then, concatenates the four groups of two hexadecimal digits. For example, the TCP/IP address 10.131.40.59 would use the terminal ID 0A83283B as shown in the following example:

```
RDEFINE TERMINAL(0A83283B) UACC(NONE)
PERMIT ZORGR CLASS TERMINAL ID(0A83283B) ACCESS(READ)
```

An asterisk (*) can be used as a wildcard in a terminal ID to specify a range of TCP/IP addresses associated with a specific network. For example:

```
RDEFINE TERMINAL(0A8328*)  UACC(NONE)
PERMIT GROP1 CLASS TERMINAL ID( 0A8328*)  ACCESS(READ)
```

allows all users from group GROP1 to use the 10.131.40.*xx* network when connecting to the SequeLink Server. Any other users connecting from this network are denied access.

Before activating the RACF TERMINAL class:

- Review the information on terminal security in your Security Manager documentation.
- Remember that other applications, such as TSO, use the RACF TERMINAL class.
- Define several terminals that have access to TSO to avoid creating a situation in which no one can log on.

The following examples show how to configure terminal security for SequeLink. Example A shows how to restrict specific TCP/IP addresses to specific users; example B shows how to restrict groups of users to specific TCP/IP addresses.

Prerequisite tasks for the following examples include:

1. Create the RACF GROUP named ZORGR with option NOTERMUACC.
2. Create the RACF userid ZORRO, making sure that the group ZORGR is the default group.

3 Connect user ZORRO to ZORGR.

Example A: Restricting TCP/IP Addresses to Specific Users

1 Grant read access to all terminals by setting `UACC(READ)`. This ensures that no one can log on to TSO when you activate the TERMINAL class. For example:

```
SETROPTS  TERMINAL(READ)
```

2 Set the terminal address so that it is not generally accessible. For example, to limit access to Terminal 0A83283B = TCPIP address 10.131.40.59:

```
RDEFINE TERMINAL(0A83283B) UACC(NONE)
PERMIT ZORGR CLASS TERMINAL ID(0A83283B) ACCESS(READ)
```

Group ZORGR users are allowed to access System from Terminal 0A83283B = TCPIP address 10.131.40.59.

3 Activate the TERMINAL class and load it in storage.

```
SETROPTS CLASSACT(TERMINAL)  RACLIST(TERMINAL)
```

Example B: Restricting Groups of Users to Specific TCP/IP Addresses

This example shows how to restrict groups of users to specific TCP/IP addresses. For example, you may want to make sure that a user ID associated with an application running on an application server, such as IBM WebSphere, can only log on the SequeLink Server from a specific TCP/IP address.

1 Ensure that the NOTERMUACC option, which enforces `UACC(NONE)` on the TERMINAL class, is in effect for the Group WEBSPHR, even if you specified SETROPS TERMINAL(READ). Ensure that this is the default and only group.

```
ALTGROUP WEBSPHR  NOTERMUACC
```

2 Grant UACC(READ) for all terminals to avoid a situation in which no one can log on to TSO once you activate the TERMINAL class. This will not apply to group WEBSPHR because of the NOTERMUACC option.

```
SETROPTS  TERMINAL(READ)
RDEFINE TERMINAL(0A83283B) UACC(NONE)
```

3 Define `Terminal 0A83283B = TCPIP address 10.131.40.59` as not accessible.

```
PERMIT WEBSHPR  CLASS TERMINAL ID(0A83283B) ACCESS(READ)
```

Group WEBSHPR users are allowed to access the system from Terminal 0A83283B = TCPIP address 10.131.40.59, in combination with the NOTERMUACC option on GROUP WEBSPHR. This effectively restricts users of this Group to Terminal 10.131.40.59.

4 Activate class TERMINAL and load it in storage.

```
SETROPTS  CLASSACT(TERMINAL)  RACLIST(TERMINAL)
```

SequeLink® Manager Security Attribute Defaults for z/OS

This section lists the installation defaults for SequeLink's security attributes for z/OS and describes the effect each combination of settings has on security.

Authentication Defaults for Administrative Activities

`ServiceAdminAuthMethods=OSLogon(UID,PWD)`

The default for this attribute means that everyone who can provide a valid host user name and password will be allowed to administer and monitor the SequeLink Server.

Authentication Defaults for Data Access Activities

```
ServiceAuthMethods=OSLogon(UID,PWD)
DataSourceLogonMethod=OSIntegrated
DataSourceReadOnly=No
```

The combination of defaults for these attributes means that everyone who can provide a valid host user name and password will be allowed to access the database using this service. The database connection accepts all types of SQL statements. Once connected to the database, the database security system will guarantee that the user can only perform actions that are allowed by the database administrator.

```
ServiceEncryptionAlgorithm=none
```

The default for this attribute means that cleartext messages are transmitted between the client and server. Note that user IDs and passwords are never sent as cleartext.

Authorization Defaults

By default, authorization is disabled for administrative activities and data access.

Using UID Mapping

UID mapping is the mapping of user IDs to alternate user IDs using a UID map. You can use UID mapping to prevent users from updating DB2 tables using commonly available tools, such as QMF or SPUFI, while preserving their ability to update DB2

tables using SequeLink. For example, suppose a user, SMITH, has privileges defined in a UID map as shown:

User ID	DB2 Table Privilege	SequeLink Plan Privilege	Application
SMITH	UPDATE	EXECUTE	SequeLink service
SMITH	UPDATE	EXECUTE	SPUFI

In this example, SMITH can update DB2 tables using SPUFI and the SequeLink service.

To prevent SMITH from updating DB2 tables using SPUFI, you can map the logon ID to an alternate user ID (for example, SMITH=SMITHB). Once the logon ID SMITH has been mapped to the alternate user ID SMITHB, you can specify DB2 table privileges as shown:

User ID	DB2 Table Privilege	SequeLink Plan Privilege	Application
SMITHB	UPDATE	EXECUTE	SequeLink service
SMITH	SELECT	EXECUTE	SPUFI

The UPDATE privilege set for SMITHB allows SMITH to update DB2 tables using a SequeLink service. The SELECT privilege set for SMITH allows read-only access to the DB2 tables using SPUFI.

NOTE: Alternate UIDs are used internally for UID mapping only. If a SequeLink Client attempts to log on with an alternate UID, the logon will be rejected. You can also map an RACF group to a single alternate user, simplifying the administrative task of managing multiple users.

You can define multiple UID maps for the SequeLink Server. These maps can be referenced by the SequeLink server data sources; only one map can be referenced by the SequeLink service. The UID map defined at the SequeLink server data source takes precedence over the UID map defined at the SequeLink service level.

When a UID map is specified for a service or data source, the SequeLink Client's user ID is mapped to an alternate UID as specified in the UID map. If a UID map has been specified for the SequeLink service or data source, the alternate UID in the UID map will be used as the DB2 authorization ID when logging on to DB2. If an alternate UID cannot be found in the UID map, the SequeLink Client's logon ID will be used as the DB2 authorization ID when MVSUIDDefaultAccess=PERMIT. All status displays will continue to show the SequeLink Client's logon ID.

To configure UID mapping for a SequeLink service or data source, set the following attributes:

MVSDataSourceUIDMap	Specify the name of the UID map you want the data source to use.
MVSServiceUIDMap	Specify the name of the UID map you want the service to use.
MVSUIDDefaultAccess	Controls UID mapping behavior for a service or data source. Valid values include: ■ PERMIT=If user ID mapping is set for the service or data source and the user ID cannot be found in the UID map, the connection is accepted. ■ DENY=If user ID mapping is set for the service or data source and the user ID cannot be found in the UID map, the connection is refused.

To configure UID map entries, set the following attribute:

MVSUID

Specify an entry in the UID map using the format *user=mapped_user* or **=mapped_user*, where:

- *user* is a valid user or user group for the z/OS security system.
- * is a wildcard for any user.
- *mapped_user* is a valid DB2 authorization ID.

**=mapped_user* is required when the service attributes ServiceAuthMethods=Anonymous and MVSServiceSecurity=SAFNONE.

For example, suppose you wanted to configure UID maps for a service and several data sources. First, you define the UID maps that the SequeLink service and the server data sources will use, as shown in Table 13-2:

Table 13-2. Defining UID Maps

UIDMap	UID	UID Map Definitions
UIDMap1	MVSUID	SMITH=SMITHA
	MVSUID	ERICK=ERICKA
	MVSUIDDefaultAccess	DENY
UIDMap2	MVSUID	SMITH=APPDBA
	MVSUID	EDWARD=APPDBA
	MVSUID	*=APPDBU
	MVSUIDDefaultAccess	Permit

Then, you configure service SLDB2 with one of the UID maps defined in Table 13-2, specifying MVSServiceUIDMap=UIDMap1.

Finally, you configure data sources, both with and without UID maps, as shown in Table 13-3:

Table 13-3. SequeLink Data Source Definitions

Data Source	UID Service Settings
Accounting	MVSDataSourceUIDMap=UIDMap2
Shipping	No UID map was specified for this SequeLink server data source.

Using this example, the following scenarios could occur:

User ID	Service	Map ID	Action	Explanation
ALBERT	SLDB2A	n/a	Denied	Connection denied because ALBERT was not in UIDMAP1
ALBERT	SLDB2C	n/a	Permit	Connection permitted to SLDB2C as ALBERT
SMITH	SLDB2B	APPDB2B	Connect	Connection to SLDB2B as APPDB2B
EDWARDS	SLDB2B	APPDB2B	Connect	Connection to SLDB2B as APPDB2B
ERICK	SLDB2B	n/a	Denied	Connection denied because ERICK was not in UIDMAP2
ERICK	SLDB2C	ERICKC	Connect	Connection to SLDB2C as ERICKC
SMITH	SLDB2D	n/a	Connect	No UID mapping for SLDB2D

User ID	Service	Map ID	Action	Explanation
Anonymous	SLDB2B	n/a	Denied	Connection denied because no *=mapped_user entry in UID map
Anonymous	SLDB2E	APPDDB2Z	Connect	Connection to SLDB2E as APPDB2Z

Table 13-4 show several sample scenarios that could occur using the mappings in the preceding tables. In each case, the data source UIDMap, when present, takes precedence over the UIDMap of the service.

Table 13-4. Sample UID Mapping Scenarios

User ID	Data Source	Map ID	Action and Explanation
ALBERT	Accounting	UIDMap2	Connection is permitted because ALBERT falls under the default access (PERMIT); the user ID is mapped to APPDB2U.
ALBERT	Shipping	UIDMap1	Connection to Shipping is refused because ALBERT is not named in UIDMap1, and the default access is set to DENY.
ERICK	Accounting	UIDMap2	Connection is permitted because ERICK falls under the default access (PERMIT); the user ID is mapped to APPDB2U.
ERICK	Shipping	UIDMap1	Connection is permitted because ERICK is an entry in UIDMap1; the user ID is mapped to ERICKA.

Table 13-4. Sample UID Mapping Scenarios *(cont.)*

User ID	Data Source	Map ID	Action and Explanation
EDWARD	Accounting	UIDMap2	Connection is permitted because EDWARD is an entry in UIDMap2; the user ID is mapped to APPDBA.
EDWARD	Shipping	UIDMap1	Connection to Shipping is refused because EDWARD is not named in UIDMap1, and the default access is set to DENY.

Using Application IDs to Restrict User Access

Application IDs are alphanumeric strings passed by a SequeLink Client that identify the client application to a SequeLink service that has been configured to accept connections only from specific application IDs.

Application IDs add another layer of security for the connection to the data store beyond that provided by the Data Store Logon security mechanism. Data Store Logon allows all users of client applications to access the data store if they meet the qualifications set by Data Store Logon. Using application IDs, you can restrict connections to the data store to only those client applications that identify themselves to the SequeLink Server through an application ID.

On Windows platforms, application IDs can be specified explicitly by the client application or they can be automatically generated by the ODBC Client or the ADO Client. The advantage of using application IDs generated by the ODBC Client or ADO Client is the application itself does not need to contain the application ID; however, you must specify in the client application that you

want to turn on the automatic generation of application IDs. The application ID is generated using the sha-1 hashing algorithm, resulting in a 160-bit hash value.

Specifying Application IDs Using ODBC Client Applications

This section describes how to specify application IDs explicitly using ODBC client applications and by turning on the automatic generation of application IDs.

Specifying Application IDs Explicitly

ODBC client applications can identify themselves explicitly to the SequeLink service in any of the following ways:

- **Specifying the application ID in the ODBC connection string that is passed to SQLDriverConnect**. For example:

  `....;APPID=MyAppID;`

  or

  `....;ApplicationID=MyAppID;`

  where *MyAppID* is the application ID.

- **Specifying the application ID using SQLSetConnectAttr**. Immediately after each call to SQLConnect or SQLDriverConnect connecting to the ODBC Client, call SQLSetConnectAttr as shown:

  `SQLSetConnectAttr(hdbc, 1053, "myAppId", SQL_NTS)`

  where *myAppId* is the application ID.

  The SQLSetConnectAttr is defined in sql.h. If an incorrect application ID is specified, the SQLSetConnectAttr fails and all subsequent SQL statements fail.

Generating Application IDs Automatically

ODBC client applications can turn on automatic application ID generation in any of the following ways:

- **Specifying the automatic application ID method in the ODBC connection string that is passed to SQLDriverConnect**. For example:

  ```
  ....;AutomaticApplicationID=x;
  ```

  where *x* is set to one of the following values:

  - When set to 1, the full path of the application executable is used as input for the hash function.
  - When set to 2, the executable binary file is used as input for the hash function.
  - When set to 3, both the full path of the application executable and the executable binary file are used as input for the hash function.
  - When set to 4, the full directory name of the application executable is used as input for the hash function.

- **Specifying SQLSetConnectAttr**. Immediately after each call to SQLConnect or SQLDriverConnect connecting to the ODBC Client, call SQLSetConnectAttr as shown:

  ```
  SQLSetConnectAttr(hdbc, 1054, x, SQL_IS_INTEGER)
  ```

  where *x* is one of the following values:

  - When set to 1, the full path of the application executable is used as input for the hash function.
  - When set to 2, the executable binary file is used as input for the hash function.

- When set to 3, both the full path of the application executable and the executable binary file are used as input for the hash function.
- When set to 4, the full directory name of the application executable is used as input for the hash function.

Specifying Application IDs Using ADO Client Applications

This section describes how to specify application IDs explicitly using ADO client applications and by turning on the automatic generation of application IDs.

Specifying Application IDs Explicitly

Using the ADO Client, the client application specifies the following *key-value* pair in the DBPROP_INIT_PROVIDERSTRING property of the DBPROPSET_DBINITALL property set:

```
ApplicationID=MyAppID;
```

where *myAppID* is the application ID.

Generating Application IDs Automatically

Using the ADO Client, the client application specifies the following *key-value* pairs in the DBPROP_INIT_PROVIDERSTRING property of the DBPROPSET_DBINITALL property set:

```
Automatic Application ID=x
```

where:

- When Automatic Application ID is set to 1, the full path of the application executable is used as input for the hash function.
- When Automatic Application ID is set to 2, the executable binary file is used as input for the hash function.
- When Automatic Application ID is set to 3, both the full path of the application executable and the executable binary file are used as input for the hash function.
- When Automatic Application ID is set to 4, the full directory name of the application executable is used as input for the hash function.

Specifying Application IDs Using JDBC Client Applications

After establishing a connection with the JDBC driver, immediately invoke setApplicationId. The setApplicationId method is defined on the interface com.ddtek.jdbc.extensions.SlExtensionInterface, and uses the following method prototype:

```
public void setApplicationId(String s) throws SQLException
```

You can set the application ID as shown in the following example:

```
import java.sql.*;
import com.ddtek.jdbc.extensions.SlExtensionInterface;
```

```
...
Connection con = DriverManager.getConnection(...);

String appId = "myAppID";
if (con instanceof SlExtensionInterface)
   {
   SlExtensionInterface slCon = (SlExtensionInterface)con;
   slCon.setApplicationId(myAppID);
   }
```

where *myAppID* is the application ID.

Configuring the List of Authorized Application IDs

How you configure the list of authorized application IDs depends on whether the SequeLink Client specifies the application ID explicitly or allows the ODBC Client or ADO Client to automatically generate an application ID:

- When the application explicitly specifies an application ID, set the DataSourceApplID service attribute to the application ID string.
- When the application generates an automatic application ID (ODBC Clients and ADO Clients only), set the DataSourceAutoApplID service attribute to the value of the automatically generated application ID. Optionally, you can add a description of the attribute, for example:

  ```
  DataSourceAutoAppId=InventoryControl=
  aaf7798c8c66e6b3a6b7be6946
  ```

  Continue to "Obtaining the Value of Automatically Generated Application IDs" on page 341 for instructions on obtaining the value of automatically generated application IDs.

NOTE: The DataSourceApplID and DataSourceAutoApplID service attributes are not, by default, included in a data access service template; therefore, you must explicitly add them. See "Adding a Service Attribute" on page 73 for instructions on adding service attributes.

Obtaining the Value of Automatically Generated Application IDs

1. Turn on the debug log level for the SequeLink service the client application will be using. For example, set ServiceDebugLogLevel=4 (Debug).
2. Connect to the SequeLink service with your ODBC or ADO application using the values 1, 2, 3, or 4 to turn on automatic application ID generation:
   - If 1 is specified, the full path of the application executable is used as input for the hash function.
   - If 2 is specified, the executable binary file is used as input for the hash function.
   - If 3 is specified, both the full path of the application executable and the executable binary file are used as input for the hash function.
   - If 4 is specified, the full directory name of the application executable is used as input for the hash function.

   The connection request will fail and the following message will be generated:

   ```
   [DataDirect][ODBC SequeLink driver][SequeLink
   Server]The application specified an invalid application
   identifier
   ```

3 Open the log file and look for the following entry:

```
CHAIN
   PROVIDE
     refNum       :0
     refNumType   :connect
       direction  :set
         000) Id  :kSSP PID CLOSEDID
           Type   :binary
           Info   :0xGAppID
```

where the set of 40 characters following `Info :0x` is the generated application ID.

4 Configure the SequeLink service to accept the generated application ID by setting the DataSourceAutoApplID service attribute to the list of IDs you generated in Step 3.

NOTE: Remember to turn off the debug log level for the SequeLink service. (Set ServiceDebugLogLevel=3 (Error))

Configuring TCP/IP Location Filters

TCP/IP location filters allow you to control which clients have access to a SequeLink service based on the network address of the client originating the request. To configure TCP/IP location filters for:

- A SequeLink data access service, set the ServiceAuthorizedClient service attribute.
- A SequeLink agent service, set the ServiceAuthorizedAdminClient service attribute.

You can specify one or multiple location filters in either address or name formats as shown in the following examples:

Client TCP/IP host name	`burner.ddtek.com`
Client TCP/IP domain names (using a wildcard)	`192.16.2.*` or `*.ddtek.com`
Client TCP/IP address	`127.0.0.1`
Client TCP/IP address range (using a wildcard)	`192.16.*.*`

NOTE: When using host names, the ServiceResolveHostNames service attribute must be set to TRUE, and only primary domain names can be used.

To configure multiple TCP/IP location filters, you must set the ServiceAuthorizedClient or ServiceAuthorizedAdminClient attribute multiple times, one instance for each location filter. For example:

```
ServiceAuthorizedClient=192.16.*.*
ServiceAuthorizedClient=192.17.*.*
ServiceAuthorizedClient=192.18.*.*
```

Required Permissions for the Java 2 Platform

Using the SequeLink Client *for* JDBC on a Java 2 Platform with the standard Security Manager enabled requires certain permissions to be set in the security policy file of the Java 2 Platform. This security policy file can be found in the jre/lib/security subdirectory of the Java 2 Platform installation directory. The required permissions can be found in the JDBCDriver.policy file in the *install_dir*/driver/lib directory.

NOTE: Web browser applets running in the Java 2 plug-in are always running in a Java Virtual Machine with the standard Security Manager enabled.

To run an application on a Java 2 Platform with the standard Security Manager, use the following command:

"`java -Djava.security.manager` *`application_class_name`*"

where *`application_class_name`* is the class name of the application.

Refer to your Java 2 Platform documentation for more information about setting permissions in the security policy file.

Permissions for Establishing Connections

To establish a connection to the database server, the JDBC Client must be granted the permissions as shown in the following example:

```
grant codeBase "file:/install_dir/lib/-" {
   permission java.net.SocketPermission "*", "connect";
};
```

where *install_dir* is the JDBC Client installation directory.

Granting Access to Temporary Files

Access to the temporary directory specified by the Java Virtual Machine configuration must be granted in the security policy file of the Java 2 Platform to use insensitive scrollable cursors or to perform client-side sorting of DatabaseMetaData result sets. The following example shows permissions that have been granted for the C:\TEMP directory:

```
grant codeBase "file:/install_dir/lib/-" {
// Permission to create and delete temporary files.
// Adjust the temporary directory for your environment.
   permission java.io.FilePermission "C:\\TEMP\\-", "read,write,delete";
};
```

where *install_dir* is the JDBC Client installation directory.

Permissions for Kerberos Authentication

To use Kerberos authentication with the JDBC Client running under a security manager, the application and driver code bases must be granted security permissions in the security policy file of the Java 2 Platform as shown in the following example.

```
grant codeBase "file:/install_dir/lib/-" {
permission javax.security.auth.AuthPermission
"createLoginContext.JDBC_DRIVER_01";
permission javax.security.auth.AuthPermission "doAs";
permission javax.security.auth.kerberos.ServicePermission
"krbtgt/your_realm@your_realm", "initiate";
permission javax.security.auth.kerberos.ServicePermission
"principal_name/hostname@your_realm", "initiate";
};
```

where:

install_dir is the JDBC Client installation directory.

principal_name is the service principal name registered with the Kerberos Key Distribution Center (KDC) that identifies the SequeLink service.

your_realm is the Kerberos realm (or Windows Domain) to which the SequeLink Server host machine belongs.

hostname is the host name of the machine running the SequeLink service.

Service Attributes that Affect Security

Table 13-5 briefly describes the service attribute that SequeLink Server uses to set security features. For more detailed information, see Appendix D "SequeLink® Service Attributes" on page 491.

Table 13-5. Attributes Used to Set Security

Service Attribute	Description
"DataSourceApplID" on page 503	Specifies a list of application IDs for the service attribute.
"DataSourceAutoApplId" on page 504	Specifies an application ID that is automatically generated by the ODBC Client to identify the client application to the SequeLink service.
"DataSourceLogonMethod" on page 521	Specifies the method to be used to log on to the data store.
"MVSDataSourceUIDMap" on page 548	Specifies the name of a user ID (UID) map. UID maps can be referenced at both the service and data source level. The valid value is a defined UIDMap name.
"MVSServiceAdminAuthorizationClass" on page 553	Specifies a general resource class name to be used by the server.
"MVSServiceAdminAuthorizationEnable" on page 554	Enables authorization for the SequeLink agent.
"MVSServiceAuthorizationClass" on page 555	Specifies a general resource class name used by the SequeLink Server when authorization is enabled.
"MVSServiceAuthorizationEnable" on page 555	Enables authorization for data access to the service.
"MVSServiceAuthorizationResource" on page 555	If MVSServiceAuthorization is enabled, the resource name is used to validate a connection request against the z/OS security system. If the attribute's value is blank or an empty string, the server uses the service name as the resource name to be checked.

Table 13-5. Attributes Used to Set Security (cont.)

Service Attribute	Description
"MVSServiceDataSourceAuthorization" on page 556	Enables authorization for data access to the server data source.
"ServiceAdminAuthMethods" on page 559	Specifies authentication mechanisms that the SequeLink Manager can use to authenticate itself to the server.
"ServiceAdministrator" on page 559	Sets authorization for users who are allowed to manage SequeLink services using the SequeLink Manager.
"ServiceAdministratorGroup" on page 561	Sets authorization for defined Linux, UNIX, and Windows user groups who are allowed to manage SequeLink services using the SequeLink Manager.
"ServiceAdminKerberosPrincipalName" on page 561	Specifies a service principal name other than the default service principal name to be used for Kerberos authentication for the SequeLink Agent service.
"ServiceAuthMethods" on page 562	Specifies one or multiple authentication mechanisms the service accepts.
"ServiceAuthorizedAdminClient" on page 563	Specifies one or multiple client TCP/IP network identifiers that are allowed to access the service using an administrator client.
"ServiceAuthorizedClient" on page 564	Specifies client TCP/IP network identifiers to limit the number of clients that are allowed data access connections to data access services.
"ServiceEncryptionAlgorithm" on page 572	Specifies the data scrambling algorithm used when sending requests or replies across the network between client and server.
"ServiceKerberosPrincipalName" on page 579	Specifies the case-sensitive service principal name to be used for Kerberos authentication to the SequeLink data access service on Linux/UNIX/Windows.
"ServiceSSLCipherSuites" on page 584	Specifies the supported cipher suites. The default is TLS_DH_anon_WITH_AES_128_CBC_SHA.

Table 13-5. Attributes Used to Set Security *(cont.)*

Service Attribute	Description
"ServiceSSLEnabled" on page 584	Specifies whether SSL or TLS is enabled for the SequeLink service. The default is false.
"ServiceSSLSessionCacheSize" on page 584	Specifies the maximum number of SSL or TLS session identifiers that are cached on the SequeLink Server before the sessions are flushed from the cache.
"ServiceSSLSessionCacheTimeout" on page 585	Specifies the maximum time in seconds to keep a SSL or TLS session in the SequeLink Server session cache. The default is 300 seconds.
"ServiceSSLVersions" on page 585	Specifies the version of the SSL standard that is used for encryption. The default is TLS 1.0.
"ServiceUser" on page 587	Sets authorization for users who are allowed to access the service for data access. The setting can be affected by the value of ServiceAuthMethods.

14 Configuring the SequeLink® Proxy Server

This chapter describes how to configure the SequeLink Proxy Server.

Using the SequeLink® Proxy Server

Untrusted applets cannot open a connection to a machine other than the originating host. Therefore, if any JDBC Client will be used by an untrusted applet, your SequeLink Server software must be installed on the same machine as your Web server software. This is a Java restriction. To circumvent this restriction, SequeLink provides a component written in Java that you can install on your Web server host called the *SequeLink Proxy Server*.

Installing the SequeLink Proxy Server on the Web server from which your JDBC applets are downloaded allows untrusted applets to connect to SequeLink Servers on hosts other than the Web server, as shown in Figure 14-1.

Figure 14-1. SequeLink Proxy Server Installed on a Web Server

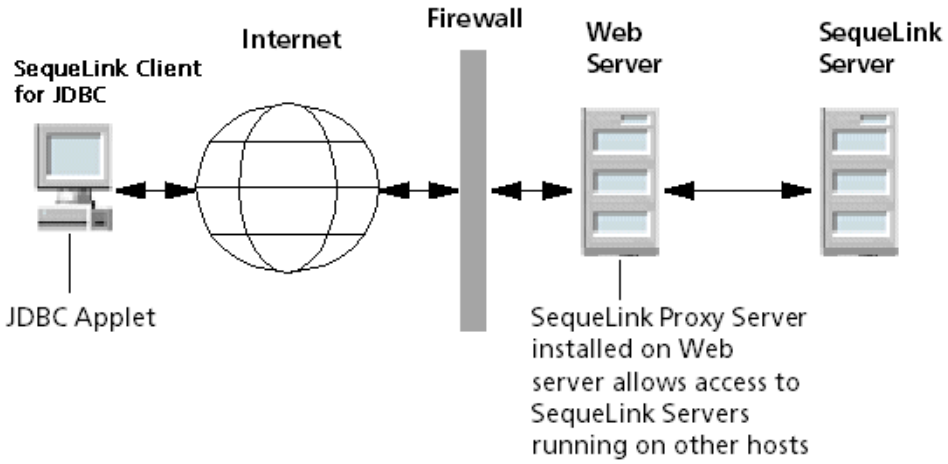

The SequeLink Proxy Server maps incoming TCP/IP connection requests from the JDBC Client to outgoing TCP connections to other hosts. When the SequeLink Proxy Server receives a connection request on a particular TCP/IP port, the SequeLink Proxy Server establishes a TCP/IP connection to a remote host and transfers data packets between the JDBC Client and the remote host.

In addition, you can use SSL encryption with the proxy server to encrypt data between the SequeLink Proxy Server and the JDBC Client. You can also use SSL with a Java application running on your Intranet to secure data over your entire network by installing the SequeLink Proxy Server on the same machine as the SequeLink Server. For example, you may want to use SSL to encrypt the data sent between an application server and the data store serviced by a SequeLink Server on another machine. See "Using SSL Encryption" on page 360 for more information about SSL.

Configuring the SequeLink® Proxy Server

Each SequeLink service serviced by the SequeLink Proxy Server must be described in a configuration file, *service_name*.cfg, where *service_name* is the name of the service. We recommend that the service name be the same as the SequeLink service it is servicing. Configuration files are stored in the proxy server directory and use the following *keyword=value* pairs:

Port	The incoming TCP/IP port. The JDBC applet or application must specify this TCP/IP port (and the IP address of the Proxy Server host) in the JDBC connection string.
ServerPort	The TCP/IP port of the service to which the final connection is made. This port must be the same port defined in the service configuration on the remote host. A default SequeLink service installation uses the port 19996.
Host	The IP address of the remote host or a symbolic host name.
AdminPort	The TCP/IP port on which the SequeLink Proxy Server listens for administration requests (for example, requests to stop the SequeLink Proxy Server). NOTE: If you do not want the SequeLink Proxy Server to listen for administration requests, omit this keyword from the configuration file. For example, if the SequeLink Proxy Server is installed on a Web server that is accessible by the Internet, your firewall may be configured to block requests from the Internet to the proxy server administration port.

You can find a configuration file template (proxyserver.cfg) in the proxy server directory. The configuration file must be located

in the directory from which you start or stop the SequeLink Proxy Server.

Configuration File Example:

```
Port=4000
ServerPort=4003
Host=189.23.5.132
AdminPort=5000
```

NOTES:

- Keywords in the configuration file are case-sensitive.
- Make sure that you use different port numbers for the Port and AdminPort keywords. Also, the port numbers for the Port and AdminPort keywords must be unique (cannot be used by another TCP/IP service).

- **On Windows**, you can use the SequeLink Manager to obtain the TCP/IP ports used by SequeLink services. In addition, you can verify the TCP/IP ports in the system32\etc\drivers\services file.

- **On Linux and UNIX**, you can verify the TCP/IP ports in /etc/services.

Starting and Stopping the SequeLink® Proxy Server

This section provides instructions for starting and stopping the SequeLink Proxy Server.

Starting the SequeLink® Proxy Server

On Windows:

Open a command-line window and change the working directory to the proxy server directory. Start the SequeLink Proxy Server by running the command appropriate for the Java Virtual Machine (JVM) you are using:

```
proxyserver14 -s [-v jview] configfile
```

where *configfile* is the name of the proxy server configuration file without the .CFG extension. By default, this batch file uses the J2SE JVM. If you want to use the Microsoft JVM, specify the optional parameter `-v jview` as shown in the preceding example.

On Linux/UNIX:

Start the SequeLink Proxy Server by running the shell script: appropriate for the JVM you are using:

```
proxyserver14.sh -s [-v jview] configfile
```

where *configfile* is the name of the proxy server configuration file without the .CFG extension. The configuration file must be located in the directory from which you start or stop the SequeLink Proxy Server.

Stopping the SequeLink® Proxy Server

On Windows servers:

Open a command-line window and change the working directory to the proxy server directory. Stop the SequeLink Proxy Server by running the command appropriate for the JVM you are using:

```
proxyserver14 -q [-v jview] configfile
```

where *configfile* is the name of the proxy server configuration file without the .CFG extension. By default, this BAT file uses the J2SE JVM. If you want to use the Microsoft JVM, specify the optional parameter `-v jview` as shown in the preceding example.

On Linux/UNIX:

Stop the SequeLink Proxy Server by running the shell script appropriate for the JVM you are using:

```
proxyserver14.sh -q [-v jview] configfile
```

where *configfile* is the name of the proxy server configuration file without the .CFG extension. The configuration file must be located in the directory from which you start or stop the SequeLink Proxy Server.

SequeLink® Proxy Server Logging

All messages generated by the SequeLink Proxy Server are written to a log file in the *installdir*/proxy/log/ directory, where *installdir* is your installation directory. The log file name has the format:

proxy_server_name.log

where *proxy_server_name* is the name of the SequeLink Proxy Server. Severe errors and information, such as `server started` or `server stopped` messages display on the screen also.

Using the SequeLink® Proxy Server as a Windows Service

Before you install the SequeLink Proxy Server as a Windows service, check the following requirements:

- Make sure that you have administrator rights. Installing and un-installing the SequeLink Proxy Server as a Windows service requires making changes to the HKEY_LOCAL_MACHINE key in the Windows Registry.
- Make sure that the directory your JVM is installed and is specified in the correct sequence in the system definition of the PATH environment variable. Because the SequeLink Proxy Server Windows service is configured to run under the local system account, access to network resources is not available. If the system definition of the PATH environment variable contains a network directory before the directory in which the JVM is installed, you will not be able to start the SequeLink Proxy Server.

  If you cannot start the SequeLink Proxy Server, either:

  - Redefine the system definition of the PATH environment variable so that the network directory appears in the system definition after the directory in which the JVM is installed. Then, reboot to make your changes effective for the local system account.
  - Change the definition of the SequeLink Proxy Server Windows service to run under an account that has access to the specific network drive.
- Make sure the CLASSPATH environment variable is defined correctly for your JVM and that the SequeLink Proxy Server .jar files are added to the CLASSPATH.

Installing the SequeLink Proxy Server as a Windows Service

1. Create a proxy server configuration file.
2. Open a Windows command window and change the working directory to the proxy subdirectory of the SequeLink Client *for* JDBC directory.
3. Issue the following command:

   ```
   cmdsrvc -s service_name -c -r [-v jview]
   ```

   where *service_name* is the name of the proxy server configuration file. This command creates a Windows service for the SequeLink Proxy Server. Use the Windows Event Viewer to verify that the service was created successfully (in the Application log for the source cmdsrvc). By default, the JDK JVM is used. If you want to use the Microsoft JVM, specify the optional parameter `-v jview` as shown in the preceding example.

   The Windows service you created should have the following attributes:

   - Automatic startup
   - Log on as System Account
   - Allow service to interact with the desktop

   In addition, a Windows Event Viewer source is defined with the name of the SequeLink Proxy Server. The SequeLink Proxy Server logs start and stop messages to this source.
4. Start the Windows service using the Windows Services control panel. Because the service is configured for automatic startup, it will also start when the Windows machine is initialized.

NOTE: Make sure that the following files located in the proxy/lib directory are added to the CLASSPATH definition of your JVM:

For a SequeLink Proxy Server running in...	Add these files to the CLASSPATH of your JVM...
J2SE Platform JVM without SSL	slproxy.jar
J2SE Platform JVM with SSL or data scrambling enabled	slproxy.jar, slssl14.jar, and iaik_jce_full.jar

Un-Installing the SequeLink® Proxy Server as a Window Service

Before you un-install the SequeLink Proxy Server as a Windows service, make sure that you have administrator rights.

To un-install the SequeLink Proxy Server:

1. Stop the SequeLink Proxy Server Windows service using the Windows Services control panel.
2. Open a Windows command-line window.
3. Change the working directory to the proxy server subdirectory in the SequeLink Client *for* JDBC directory.
4. Issue the following command:

   `cmdsrvc -s service_name -d`

Using SSL Encryption

If your SequeLink environment requires greater data privacy than that provided by fixed-key DES, fixed-key 3DES, or byteswap, you can use SSL to encrypt data exchanged between the JDBC Client and the SequeLink Server. This assumes that the communication between the SequeLink Client machine (for applets, the Web server from which the applets are downloaded) and the SequeLink Server machine is secure, meaning that:

- Only authorized persons can obtain login access to the Web server machine.
- Only authorized persons can eavesdrop on (or monitor) the communication (physical communication lines and any intermediate routers) between the Web server host and the database server host. Because the data on your Intranet is not encrypted, you also must ensure that only authorized access to internal communication lines and internal routers is permitted.

NOTE: SequeLink data scrambling (fixed-key DES, fixed-key 3DES, and byteswap) can work with SSL, resulting in a completely secure combination between the SequeLink Client *for* JDBC and the SequeLink Server.

Using SSL with a Java application running on your Intranet, you can secure data over your entire network by installing the SequeLink Proxy Server on the same machine as the SequeLink Server (as shown in Figure 14-2) and specifying `localhost` as the host name of the SequeLink Server in the proxy server configuration file. The cleartext messages that are sent between the SequeLink Proxy Server and the SequeLink Server do not leave the machine.

Figure 14-2. Using SSL with the SequeLink Proxy Server Installed on the SequeLink Server

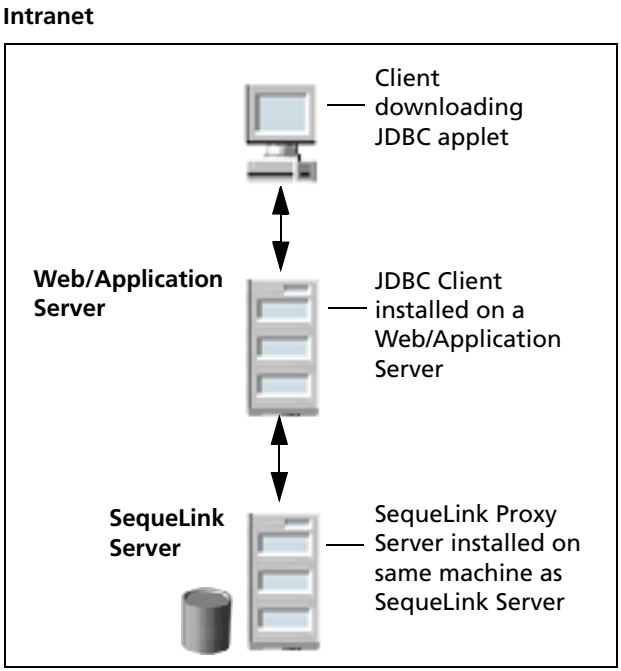

NOTE: SequeLink Proxy Server uses the IETF TLS (Transport Layer Security) 1.0 standard, the successor to the SSL 3.0 protocol.

SSL Cipher Suites

SSL cipher suite definitions have the format:

SSL_KeyExchangeMethod_WITH_DataTranserCipher_DigestFunction

The following cryptographic strong SSL cipher suites are supported by SequeLink when using SSL with the SequeLink Proxy Serve.

- SSL_DH_anon_WITH_RC4_128_MD5
- SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
- SSL_DH_anon_WITH_DES_CBC_SHA
- SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_DSS_WITH_DES_CBC_SHA
- SSL_DHE_DSS_WITH_RC4_128_SHA
- SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_DHE_RSA_WITH_DES_CBC_SHA
- SSL_RSA_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_WITH_DES_CBC_SHA
- SSL_RSA_WITH_RC4_128_MD5
- SSL_RSA_WITH_RC4_128_SHA

Cryptographic Characteristics of Key Exchange Algorithms

Table 14-1 lists the cryptographic characteristics of SSL key exchange algorithms, including a description, the key-size limit, and the type of situation for which specific algorithms are most appropriate.

Table 14-1. Cryptographic Characteristics of Key Exchange Algorithms

Key Exchange Algorithm	Description	When to Use
DH_anon	The Diffie-Hellman parameters are generated during session establishment.	When there is no risk of man-in-the-middle attacks.
DHE_DSS	The Diffie-Hellman parameters are generated during session establishment. They are signed by the DSS certificate.	When the DSS certificate of the server is used for signing only and not used for key exchange.
DHE_RSA	The Diffie-Hellman parameters are generated during session establishment. They are signed by the RSA certificate.	When the RSA certificate of the server is used for signing only and not used for key exchange.
RSA	The public key from the RSA certificate is used for key exchange.	When the server uses an RSA certificate.

Cryptographic Characteristics of Data Transfer Ciphers

Table 14-2 lists the cryptographic characteristics of data transfer ciphers, including the algorithm used and the effective key size.

Table 14-2. Cryptographic Characteristics of Data Transfer Ciphers

Data Transfer Cipher	Algorithm	Effective Key size
DES_CBC	DES in cipher block chaining mode	56
3DES_EDE_CBC	Triple DES in cipher block chaining mode	168
RC4	RC4 from RSA	128

Configuring SSL Encryption for the SequeLink® Proxy Server

You configure SSL encryption in the proxy server configuration file by adding the *keyword=value* pairs:

```
Network=ssl
CipherSuites=value
```

NOTES:

1. The Network and CipherSuites keywords in the proxy server configuration file are case-sensitive.
2. If you do not want to use SSL, specify `Network=socket` in the proxy server configuration file or omit the Network keyword from the configuration file.
3. The value of the CipherSuites keyword is a list of cipher suites to use, in order of preference. The listed cipher suites are separated by commas with no blank spaces allowed. You must specify cipher suites that use the same type of certificate. For example, you cannot specify a combination of RSA cipher suites and DSS cipher suites. See "SSL Cipher Suites" on page 362 for a list of supported cipher suites.
4. For cipher suites that require a DSS or RSA certificate, you must specify the X.509 certificate (with the public key) and the corresponding private key in the proxy server configuration file. See Table 14-3 for a list of the *keyword=value* pairs you can specify in the proxy server configuration file for each key exchange algorithm.
5. When the JDBC Client and the SequeLink Proxy Server agree on a cipher suite that requires a certificate, the JDBC Client must specify the certificate checker class that will be used to verify the certificate chain the SequeLink Proxy Server sends to the JDBC Client. See "Verifying the SequeLink® Proxy Server Certificate" on page 371 for more information on certificate checker classes.

Table 14-3 lists the key exchange algorithms you can use and the *keyword=value* pairs you can specify in the proxy server configuration file when using a particular key exchange algorithm.

Table 14-3. Key Exchange Algorithms and Keyword/Value Pairs for the SequeLink Proxy Server

Key Exchange Algorithm	Keyword	Value
DHE_DSS	DSSCertificate	Name of the file with the DSS certificate in DER format or a PKCS #7 certificate chain.
	DSSPrivateKey	Name of the file with the DSS private key in PKCS #8 encrypted format.
	PassPhrase	Pass phrase with which the private key file is encrypted. If this keyword is unspecified, the Proxy Server will prompt for the pass phrase.
	UsePassPhraseDialog	To be prompted for the pass phrase using the standard input/output instead of a dialog box, set this keyword to `No`. Remember that the pass phrase will be shown on the screen as you type.
DHE_RSA	RSACertificate	Name of the file with the RSA certificate in DER format or a PKCS #7 certificate chain.
	RSAPrivateKey	Name of the file with the RSA private key in PKCS #8 encrypted format.
	PassPhrase	Pass phrase with which the private key file is encrypted. If this keyword is unspecified, the Proxy Server will prompt for the pass phrase.
	UsePassPhraseDialog	To be prompted for the pass phrase using the standard input/output instead of a dialog box, set this keyword to `No`. Remember that the pass phrase will be shown on the screen as you type.
RSA	RSACertificate	Name of the file with the RSA certificate in DER format or a PKCS #7 certificate chain.
	RSAPrivateKey	Name of the file with the RSA private key in PKCS #8 encrypted format.

Table 14-3. Key Exchange Algorithms and Keyword/Value Pairs for the SequeLink Proxy Server (cont.)

Key Exchange Algorithm	Keyword	Value
	PassPhrase	Pass phrase with which the private key file is encrypted. If this keyword is unspecified, the Proxy Server will prompt for the pass phrase.
	UsePassPhraseDialog	To be prompted for the pass phrase using the standard input/output instead of a dialog box, set this keyword to `No`. Remember that the pass phrase will be shown on the screen as you type.

Using Private Keys with the SequeLink® Proxy Server

The SSL cipher suites that use server authentication require a valid server certificate and associated private key. The SequeLink Proxy Server must access the private key from a private key file. Because it is not safe to store the private key as cleartext in a file, the SequeLink Proxy Server expects the private key to be stored in PKCS #8 format, which is a standard method of storing encrypted private keys when the encryption key is derived from a pass phrase.

Providing the Pass Phrase for the SequeLink® Proxy Server

The SequeLink Proxy Server requires the pass phrase to start. The private key can be retrieved in either of the following ways:

- **When the SequeLink Proxy Server starts, it prompts for the private key**. In graphical user interface (GUI) environments, a dialog box may appear. For example:

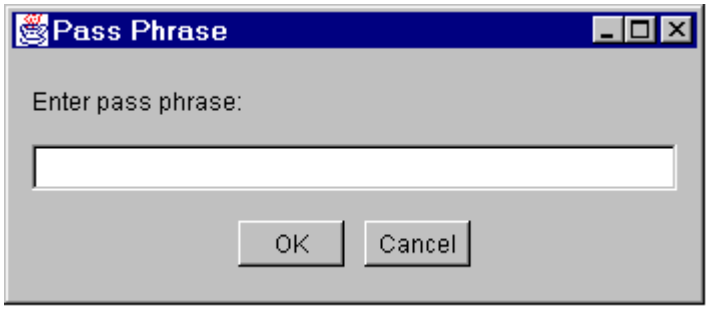

  Type the pass phrase in the appropriate field of the dialog box and click **OK**.

  In situations without a GUI, such as when the SequeLink Proxy Server is running in a terminal session on a UNIX machine, specify `UsePassPhraseDialog=No` in the proxy server configuration file. The SequeLink Proxy Server will use the standard input/output of your environment to prompt for the private key. When you type the pass phrase and press ENTER, the pass phrase displays on your standard output. When you are finished, make sure to scroll the output window so that unauthorized persons cannot see the pass phrase on your screen.

- **Code the pass phrase in the proxy server configuration file**. Add the *keyword=value* pair:

  `PassPhrase=`*`pass_phrase`*

  where *`pass_phrase`* is the pass phrase required to access the private key. Leading and trailing blanks are stripped from the

value when the pass phrase is retrieved from the configuration file; therefore, the pass phrase cannot have leading or trailing blanks in the configuration file. Make sure that only trusted accounts have access to the configuration file.

NOTE: If the SequeLink Proxy Server will be started as a Windows service, you must specify the pass phrase in the configuration file because the SequeLink Proxy Server cannot prompt for the pass phrase.

Storing the Private Key in PKCS #8 Format

If your private key is in cleartext format, you can use the encrypt.bat utility (on Windows) or the encrypt.sh shell script (on Linux/UNIX) to store the key in a file in PKCS #8 format.

The private keys are encrypted with triple DES with a 168-bit key derived from the pass phrase using a one-way hash function (SHA).

To provide sufficient randomness in the generated keys, you must provide sufficient randomness in the pass phrase. The English language has approximately 1.3 bits of randomness for each character; therefore, to provide 168 random bits for the two keys, you must have 130 characters (conservatively) of English text. Using punctuation characters and a mix of upper and lowercase letters, you can construct pass phrases that have more randomness with fewer characters.

Using the Encryption Tool

On Windows servers:

```
encrypt [-v virtual_machine] infile outfile
```

where:

virtual_machine	is the executable name of the JVM that is installed on the machine where you encrypt the key. By default, this BAT file uses the J2SE JVM. If you want to use the Microsoft JVM, specify the optional parameter `-v jview`.
infile	is the name of the cleartext file.
outfile	is the name of the encrypted file.

On UNIX:

```
encrypt.sh infile outfile
```

where:

infile	is the name of the cleartext file.
outfile	is the name of the encrypted file.

You may want to run the encryption tool on a machine other than the one running the SequeLink Proxy Server and transfer the encrypted file to the SequeLink Proxy Server host to avoid writing a copy of the private key in cleartext on the SequeLink Proxy Server host. Make sure that you transfer the complete proxy/lib directory to the machine on which you want to run the encryption tool.

The proxy server installation directory also contains a decryption tool that can be used to decrypt a file that has been encrypted with the encryption tool. The encryption and decryption tools prompt for the pass phrase and show it on the screen as you type, so make sure that you close the terminal session window after

you have encrypted or decrypted the file to prevent unauthorized people from viewing it.

Using the Decryption Tool

On Windows servers:

```
decrypt [-v virtual_machine] infile outfile
```

where:

virtual_machine	is the executable name of the JVM that is installed on the machine where you encrypt the key. By default, this BAT file uses the J2SE JVM. If you want to use the Microsoft JVM, specify the optional parameter `-v jview`.
infile	is the name of the encrypted file.
outfile	is the name of the cleartext file.

On UNIX:

```
decrypt.sh infile outfile
```

where:

infile	is the name of the encrypted file.
outfile	is the name of the cleartext file.

Verifying the SequeLink® Proxy Server Certificate

When you use a cipher suite that specifies server authentication, the SSL handshake protocol ensures that the server knows the private key that corresponds to the public key in the certificate. Subsequently, the client application must verify that the server is indeed the server with which it wants to communicate by

verifying that the received certificate is the certificate that it expects from the server.

The JDBC application or applet provides the JDBC Client with a class that implements the com.ddtek.sequelink.cert.CertificateCheckerInterface interface. If you do not supply a class that implements this interface, the connection will be refused.

This interface is defined as:

```
package com.ddtek.sequelink.cert;
public interface CertificateCheckerInterface
   {
   public void CheckCertificate(byte [][] certChain)
      throws SecurityException;
   }
```

The JDBC driver calls this method and passes the X.509 certificate chain that it received during the SSL handshake to the method. All certificates are DER encoded and the server certificate is the first certificate in the array. The checkCertificate method must verify that the received certificate is trusted and is, for example, signed by a trusted authority. If the certificate is not trusted, the method must throw a Security Exception. You specify the name of the class that implements this interface in the certificateChecker keyword in the JDBC connection URL or the data source.

The driver/examples subdirectory contains the Java source files listed in Table 14-4 as examples of classes that implement CertificateCheckerInterface.

Table 14-4. Java Source Files Implementing CertificateCheckerInterface

Java Source File	Description
CheckAgainstCertificateFromJar.java	Adapt and use for downloaded applets.
CheckAgainstCertificateFromFile.java	Adapt and use for Java applications on a client machine.
KeyStoreCertificateChecker.java	Adapt and use for Java applications that use a Java keystore to verify that the provided certificate chain is trusted.

These classes retrieve the server certificate from a JAR file, or local file, and compare it with the certificate that is passed as the first element of the certChain parameter to the checkCertificate method. You can change these files as appropriate for your environment.

Coding the certificate you want to compare other certificates against in the downloaded applet is safe only if no one tampers with the applet while it is being downloaded from your Web server. You must use signed applets and you must configure your Web browser to explicitly check the signer of downloaded applets. Alternatively, you can use a secure and authenticated SSL connection to the web server when downloading the applet.

Using the Demo Certificates, Certificate Checker, and Private-Key Format Conversion Tool

SequeLink provides some demo applications in the *installdir*/proxy/demos directory, where *installdir* is your installation directory, that allow you to create or convert certificates.

Demo Certificates

The demo certificates that SequeLink provides are intended for testing purposes only and cannot be used to provide secure connections. Table 14-5 lists the private key files and describes the corresponding certificates.

Table 14-5. Demo Certificates

File	Descriptions
demo-DSA-CA.cer	Demo Certificate Authority with a DSS X.509 certificate. This certificate is self signed.
demo-DSA-CA.pk8	Corresponding (PKCS #8 encrypted) private key of the public key provided by the certificate demo-DSA-CA.cer.
demo-DSA-server.p7b	Demo DSS server X.509 certificate signed with the public key provided by the certificate demo-DSA-CA.cer.
demo-DSA-server.pk8	Corresponding (PKCS #8 encrypted) private key of the public key provided by the certificate demo-DSA-server.cer.
demo-RSA-CA.cer	Demo Certificate Authority with an RSA X.509 certificate. This certificate is self signed.

Table 14-5. Demo Certificates *(cont.)*

File	Descriptions
demo-RSA-CA.pk8	Corresponding (PKCS #8 encrypted) private key of the public key provided by the certificate demo-RSA-CA.cer.
demo-RSA-server.p7b	Demo RSA server X.509 certificate signed with the public key provided by the demo-RSA-CA.cer.
demo-RSA-server.pk8	Corresponding (PKCS #8 encrypted) private key of the public key provided by the certificate demo-RSA-server.cer.

NOTES:

- To use the demo certificates, you must add slssl.jar (if you are using J2SE 1.3) or slssl14.jar and iaik_jce_full.jar (if you are using J2SE 1.4 or higher) to your CLASSPATH variable.
- You can re-generate demo certificates by running the following Java program in the *installdir*/proxy/ directory, where *installdir* is your installation directory:

  ```
  java com.ddtek.sequelink.demo.GenerateDemoCertificates
  ```

- You can customize the generation of these demo certificates by editing the demo.properties file in the *installdir*/proxy/demos/com/datadirect/sequelink/demo directory, where *installdir* is your installation directory.

The following examples show how to use the demo certificates:

Example A: Using SSL with an RSA Server Certificate

1 Start the SequeLink Proxy Server with the following configuration:

```
Port=9500
AdminPort=9600
Host=SequeLinkhost
ServerPort=SequeLinkport
Network=ssl
CipherSuites=SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA,SSL_
DHE_RSA_WITH_DES_CBC_SHA,SSL_RSA_WITH_3DES_EDE_CBC_
SHA,SSL_RSA_WITH_DES_CBC_SHA,SSL_RSA_WITH_RC4_128_MD5,
SSL_RSA_WITH_RC4_128_SHA
RSACertificate=cert/demo-RSA-server.p7b
RSAPrivateKey=cert/demo-RSA-server.pk8
PassPhrase=Demo Pass Phrase
```

where *SequeLinkhost* is the TCP/IP host name or address of the SequeLink Server and *SequeLinkport* is the port on which the SequeLink Server is listening for connection requests.

2 Make a connection to the SequeLink Server, for example, using DataDirect Test:

```
jdbc:sequelink:ssl://proxyserverhost:9500;
cipherSuites=SSL_RSA_WITH_RC4_128_MD5;
certificateChecker=com.ddtek.sequelink.cert.
AcceptAllCertificateChecker
```

where *proxyserverhost* is the IP address or symbolic host name of your proxy server host.

If successful, the following message appears:

```
Certificate accepted by
AcceptAllCertificateChecker.
*** ONLY FOR TESTING PURPOSES ***
Certificate chain:
1: O=SequeLink Demo Certificates, OU=Demo RSA
Server Certificate, CN=demo.ddtek.sequelink.com
```

```
2: O=SequeLink Demo Certificates, CN=Demo RSA CA
Certificate
```

Example B: Using SSL with a DSS Server Certificate

1 Start the proxy server with the following configuration:

```
Port=9500
AdminPort=9600
Host=SequeLinkhost
ServerPort=SequeLinkport
Network=ssl
CipherSuites=SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA,SSL_
DHE_DSS_WITH_DES_CBC_SHA,SSL_DHE_DSS_WITH_RC4_128_SHA
DSSCertificate=cert/demo-DSA-server.p7b
DSSPrivateKey=cert/demo-DSA-server.pk8
PassPhrase=Demo Pass Phrase
```

where *SequeLinkhost* is the TCP/IP host name or address of the SequeLink Server and *SequeLinkport* is the port on which the SequeLink Server is listening for connection requests.

2 Make a connection to the SequeLink Server, for example, using DataDirect Test:

```
jdbc:sequelink:ssl://proxyserverhost:9500;
cipherSuites=SSL_DHE_DSS_WITH_DES_CBC_SHA;
certificateChecker=com.ddtek.sequelink.cert.
AcceptAllCertificateChecker
```

where *proxyserverhost* is the IP address or symbolic host name of your proxy server host.

If successful, the following message appears:

```
Certificate accepted by
AcceptAllCertificateChecker.
*** ONLY FOR TESTING PURPOSES ***
Certificate chain:
1: O=SequeLink Demo Certificates, OU=Demo DSA
Server Certificate, CN=demo.sequelink.ddtek.com
2: O=SequeLink Demo Certificates, CN=Demo DSA CA
Certificate
```

Example C: Using SSL with Anonymous Cipher Suites (No Server Authentication)

1 Start the proxy server with the following configuration:

```
Port=9500
AdminPort=9600
Host=sequeLinkhost
ServerPort=sequelinkport
Network=ssl
CipherSuites=SSL_DH_anon_WITH_3DES_EDE_CBC_SHA,SSL_
DH_anon_WITH_DES_CBC_SHA,SSL_DH_anon_WITH_RC4_128_MD5
```

where *sequeLinkhost* is the TCP/IP host name or address of the SequeLink Server and *sequelinkport* is the port on which the SequeLink Server is listening for connection requests.

2 Make a connection to the SequeLink Server, for example, using DataDirect Test:

```
jdbc:sequelink:ssl://proxyserverhost:9500;
cipherSuites=SSL_DH_anon_WITH_DES_CBC_SHA
```

where *proxyserverhost* is the IP address or symbolic host name of your proxy server host.

Demo Certificate Checker

SequeLink provides a demo certificate checker that accepts all server certificates. It displays on the screen a warning and a description of the certificate the client received from the server through the SSL handshake. This certificate checker is implemented by the com.ddtek.sequelink.cert.AcceptAllCertificateChecker class.

Demo Private-Key Format Conversion Tool

SequeLink provides a private-key format conversion tool that can perform the following tasks:

- Export a private key and X.509 certificate from a Java keystore (JKS) to an encrypted PKCS #8 private-key file and DER-encoded certificate file
- Export a private key and X.509 certificate from a PKCS #12 keystore

The private-key format conversion tool is a command-line tool that uses the following syntax:

```
java.com.ddtek.sequelink.demo.KeyTool
[-keystore keystore]
[-alias alias]
-certfile certfile
-keyfile keyfile
[-storetype storetype]
[-storepass storepass]
[-keypass keypass]
```

where:

Parameter	Java Keystore (JKS) Export	PKCS #12 Keystore Export	Description
keystore	X	X	The file name of the JKS or the PKCS #12 keystore.
alias	X		The alias in the JKS. If supplied, it is assumed that the keystore parameter is a JKS.
certfile	X	X	The DER-encoded X.509 certificate file.

Parameter	Java Keystore (JKS) Export	PKCS #12 Keystore Export	Description
keyfile	X	X	The PKCS #8 encoded private key. The private key ends with the same password as the JKS or the PKCS #12 keystore.
storetype	X	X	The type of JKS. The default is jks. This parameter is optional.
storepass	X		The password used to protect the JKS or the PKCS #12 keystore. If omitted, you will be prompted for this password.
keypass	X		The password that protects the JKS key entry. This parameter is required when the password for the key entry is different from the keystore password.

To use the demo private-key format conversion tool, you must add slssl14.jar and iaik_jce_full.jar to your CLASSPATH variable.

15 Configuring SequeLink® Services for Your Database

Many of the SequeLink service templates provide a configuration that can be used without any modification; however, some databases must be configured before the service template can be used.

This chapter describes how to create a SequeLink service and how to configure SequeLink services for specific databases:

- "Configuring SequeLink® Server for ODBC Socket" on page 383
- "Configuring SequeLink® Server for JDBC Socket" on page 392
- "Configuring SequeLink® Server for DB2 on z/OS" on page 408
- "Configuring a SequeLink® Service for Oracle" on page 426

Creating a SequeLink® Service

When you install SequeLink Server, at least one SequeLink data access service, and one or more service templates determined by the SequeLink Server, are installed. You can use the SequeLink Manager to create additional services based on the SequeLink

service templates and to modify the default service attributes defined in the service templates:

- [SequeLink 6.0] Agent service
- z/OS [SequeLink 6.0] Agent service for z/OS
- z/OS [SequeLink 6.0] DB2 service for z/OS
- [SequeLink 6.0] DB2 UDB LUW service
- [SequeLink 6.0] DB2 UDB LUW service (enhanced code page support)
- [SequeLink 6.0] Informix service (32-bit only)
- [SequeLink 6.0] JDBC Socket service (32-bit only)
- [SequeLink 6.0] ODBC Socket service
- [SequeLink 6.0] ODBC Socket service (enhanced code page support)
- [SequeLink 6.0] ODBC Socket service (enhanced code page support - UTF8 encoding)
- [SequeLink 6.0] ODBC Socket service (enhanced code page support - UTF16 encoding)
- [SequeLink 6.0] Oracle 9 service (32-bit only)
- [SequeLink 6.0] Oracle 9 service (enhanced code page support - 32-bit only)
- [SequeLink 6.0] Oracle 10 service
- [SequeLink 6.0] Oracle 10 service (enhanced code page support)
- [SequeLink 6.0] SQL Server service
- [SequeLink 6.0] SQL Server service (enhanced code page support)
- [SequeLink 6.0] Sybase service
- [SequeLink 6.0] Sybase service (enhanced code page support)
- [SequeLink 6.0] Sybase service (enhanced code page support - UTF8 encoding)
- [SequeLink 6.0] Sybase service (enhanced code page support - UTF16 encoding)

In most cases, the SequeLink service templates provide a configuration that can be used without any modification. Not all SequeLink service attributes are defined in the templates. See Appendix D "SequeLink® Service Attributes" on page 491 for a list of all SequeLink service attributes.

Configuring SequeLink® Server for ODBC Socket

Many of the SequeLink service templates provide a configuration that can be used without any modification. During installation of the SequeLink Server for ODBC Socket, you specify a service, a backend driver, and data source. The following sections describe how to add other services and data sources.

Configuring a SequeLink® Service for ODBC Socket on Windows

You can use the same SequeLink Server for ODBC Socket to connect with different ODBC drivers. For each ODBC driver you want to access, create a new ODBC system data source. In the server data sources of the SequeLink service, you can then specify a different DataSourceSOCODBCConnStr service attribute (see "DataSourceSOCODBCConnStr" on page 536).

Setting Up the ODBC Driver Environment

Before setting up the SequeLink Server for ODBC Socket environment, ensure you have set up the backend ODBC driver environment and that you can make a connection to this ODBC driver.

Verifying the ODBC Driver Environment

The ivcheckodbc utility, which is installed in *install_dir*\tools\ivcheckodbc.bat, is used to test the backend ODBC driver.

To verify the environment for the backend ODBC driver:

1. Create a system data source for your ODBC driver. Refer to your ODBC driver documentation for instructions on creating and testing a system data source.
2. To make sure that you have set up the environment for the ODBC driver that you want to use, run the ivcheckodbc utility from a DOS prompt. For example:

   ```
   run ivcheckodbc.bat
   ```

   The ivcheckodbc utility makes a connection, checks the requirements of the backend ODBC driver, and returns the information you need to enter during the configuration of the SequeLink Server for ODBC Socket on Windows.

Configuring the SequeLink® Server for ODBC Socket

You can use the administrative tool of your choice to configure the SequeLink Server for ODBC Socket. The following procedure uses the MMC.

1. Refer to the documentation for the ODBC driver that you want to access to determine whether it is a Unicode driver. Create the appropriate SequeLink service for the ODBC driver that you will be using.
   - For a non-Unicode driver, use SequeLink ODBC Socket Service.
   - For a Unicode driver, select SequeLink ODBC Socket Service (enhanced code page support).

2 Right-click on **SequeLink Services** to create a new service, using the service template that you identified in Step 2.

3 Enter the service name and the TCP/IP port. Then, click **Next**.

4 Register the SequeLink Service and click **Finish**.

5 Select the newly created SequeLink ODBC Socket Service.

6 Select the data source you want to change, or create a new server data source by right-clicking **Data Source Settings**.

7 Right-click the **Advanced** node, and select **New / Attribute**. The Properties window appears.

8 From the drop-down list, select **DataSourceSOCODBCConnStr**.

9 In the Value field, type the connect string of the data source you created in "Verifying the ODBC Driver Environment" on page 384. For example, if you want the Socket Server to use the Oracle9i_on_host data source, type in the string `DSN=Oracle9i_on_host`.

   NOTE: This attribute contains the name of the ODBC system data source.

10 Click **OK**. The attribute is changed.

11 Save the configuration.

Configuring a SequeLink® Service for ODBC Socket on Linux/UNIX

Before setting up the SequeLink Server for ODBC Socket environment:

- Make sure that you have set up the environment for the ODBC driver that you want to use, as described in "Setting Up the ODBC Driver Environment" on page 386.
- Then, configure the SequeLink Server for ODBC Socket, as described in "Configuring the SequeLink® Server for ODBC Socket" on page 384.

The configuration examples are for a 32-bit environment. The examples are also valid for SequeLink Server for ODBC Socket in a 64-bit environment.

Setting Up the ODBC Driver Environment

Before setting up the SequeLink Server for ODBC Socket environment, ensure you have set up the backend ODBC driver environment and that you can make a connection to this ODBC driver.

Verifying the ODBC Driver Environment

The SequeLink Server for ODBC Socket uses the DataDirect ODBC Driver Manager. This driver manager loads the backend driver upon request, finding all of the necessary information in the backend ODBC data source configuration file odbc.ini, located in *installdir*/bin/odbc.

The odbc.ini file must contain one line in the [ODBC Data Sources] section and one section for the data source:

- In the [ODBC Data Sources] section, for example:

  ```
  [ODBC Data Sources]
  Socket_Default_DataSource=ODBC Driver
  ```

  This entry indicates that a data source called Socket_Default_DataSource uses the driver ODBC DRIVER.

- In the [Socket_Default_DataSource] section, as shown in this example on 32-bit Solaris:

  ```
  [Socket_Default_DataSource]
  Driver=installdir/bin/odbc/lib/odbcdrvr.so
  Description=ODBC Driver
  ```

  This entry indicates that the ODBC driver `odbcdrvr.so` for the data source ODBC Driver is located in the directory *installdir*/bin/odbc/lib/.

  The file extension of the backend ODBC driver depends on the operating system:
  - .a - AIX
  - .so - Linux, Solaris
  - .sl - HP-UX

The ivcheckodbc utility, located in the *install_dir*/tools subdirectory, is used to test the backend ODBC driver with the DataDirect Driver Manager on Linux/UNIX. The DataDirect Driver Manager on Linux/UNIX is a requirement for SequeLink Server for ODBC Socket. For more information about the ivcheckodbc utility, refer to the *SequeLink Troubleshooting Guide and Reference*.

To verify the environment for the backend ODBC driver:

1. Create a data source for your backend ODBC driver in the DataDirect Driver Manager configuration file, odbc.ini. Update the file to reflect the appropriate data source configuration information, based on the documentation of your backend driver.

   The following example uses the definitions for the DataDirect Progress ODBC driver.

   - In the [ODBC Data Source] section, add:

     ```
     [ODBC Data Sources]
     dsn_prog_10=DataDirect 5.3 Progress10
     ```

     This entry identifies a new data source called dsn_prog_10 that uses the DataDirect 5.3 Progress10 ODBC driver. The Progress10 ODBC driver finds all the necessary configuration information in section [dsn_prog_10] of the odbc.ini file.

   - Add a [dsn_prog_10] section, as shown in this example on Solaris:

     ```
     [dsn_prog_10]
     Driver=/usr/DataDirect/ODBC/0530/lib/ivpro1020.so
     Description=DataDirect 5.3 Progress OpenEdge
     DatabaseName=your_db
     HostName=your_host
     LogonID=your_uid
     password=your_pwd
     PortNumber=your_server_port
     ```

     This entry indicates that a Progress10 data source called dsn_prog_10 has been created. The ODBC driver for this datasource is located in the directory /usr/DataDirect/ODBC/0530/lib/.

2. To make sure that you have set up the environment for the ODBC driver that you want to use, run the ivcheckodbc utility, for example:

   ```
   $ /usr/slserver60/tools/ivcheckodbc.sh
   ```

The ivcheckodbc utility makes a connection, checks the requirements of the backend ODBC driver, and returns the information you need to enter during the configuration of the SequeLink Server for ODBC Socket on Linux/UNIX.

Configuring the SequeLink® Server for ODBC Socket

1 Refer to the documentation for the ODBC driver that you want to access to determine whether it is a Unicode driver. Create the appropriate SequeLink service for the ODBC driver that you will be using.

- For a non-Unicode driver, use SequeLink ODBC Socket Service.
- For a Unicode driver in UTF-8 encoding, select SequeLink ODBC Socket Service (enhanced code page support for UTF8 drivers).
- For a Unicode driver in UTF-16 encoding, select SequeLink ODBC Socket Service (enhanced code page support for UTF16 drivers).

2 Use the administration tool of your choice to create a new service. This example uses the MMC. Right-click on **SequeLink Services** to create a new service, using the service template that you identified in "Verifying the ODBC Driver Environment" on page 386.

   a When prompted, perform the following actions:

      - In the Service Name field, type the service name you want to use for the new SequeLink service. The service name must be unique (not used by another service).
      - In the TCP port field, type the number of the TCP/IP port on which the new SequeLink service will be listening. The port must be an available port.
      - Click **Next**.

   b The wizard prompts you to register the SequeLink service on the host machine. The default is to register the service, which makes the service information available to the operating system.

   c Click **Finish** to create the new SequeLink service.

3 The Driver Manager finds its configuration file using the environment variable ODBCINI. This environment variable is automatically set for the SequeLink Server for ODBC Socket environment. If you want to put the odbc.ini file elsewhere, modify the ODBCINI environment variable.

   See "ServiceEnvironmentVariable" on page 573 for more information.

   ```
   ServiceEnvironmentVariable:
   ODBCINI=installdir/bin/odbc/odbc.ini
   ```

   a Select the ODBC Socket Service that you created in Step 2.

   b Select the **Configuration** node.

   c Select the **Service Settings** node.

d Right-click the Environment node and select **New/Attribute**.

e Enter the setting for the environment variable. For example, on Solaris, enter:

```
ServiceEnvironmentVariable: ODBCINI=newdir/odbc.ini
```

4 Specify where the SequeLink Server for ODBC Socket can find the shared library of the driver manager and the backend ODBC driver (LIBPATH for AIX, SHLIB_PATH, for HP-UX, or LD_LIBRARY_PATH for Solaris and Linux).

a Select the ODBC Socket Service that you created in Step 2.

b Select the **Configuration** node.

c Select the **Service Settings** node.

d Right-click the Environment node and select **New/Attribute**.

e Enter the setting for the environment variable. For example, on Solaris, enter:

```
ServiceEnvironmentVariable:LD_LIBRARY_PATH=/usr/Data
Direct/ODBC/0530/lib/ivpro1020.so
```

5 If the ODBC driver that you are using has required specific environment variables, add them as service attributes to the SequeLink Server for ODBC Socket configuration.

a Select the ODBC Socket Service that you created in Step 2.

b Select the **Configuration** node.

c Select the **Service Settings** node.

d Right-click the Environment node and select **New/Attribute**.

e Enter the setting for the environment variable and its value. For example, on Solaris,

```
ServiceEnvironmentVariable: ORACLE_HOME=/db/oracle10
```

6 Set the connection string the SequeLink Server for ODBC Socket service will use to connect:

   a Select **Data Source Settings**; then select a node.

   b Right-click the **Advanced** node and select **New / Attribute**. The Properties window appears.

   c From the drop-down list, select **DataSourceSOCODBCConnStr**.

   d In the Value field, type the connect string of the data source you want SequeLink Server for ODBC Socket to use. For example, if you want SequeLink Server for ODBC Socket to use the data source that you created in "Verifying the ODBC Driver Environment" on page 386, type in the string `DSN=dsn_prog_10`.

7 Click **OK**. The attribute is changed.

8 Save the configuration.

Configuring SequeLink® Server for JDBC Socket

The SequeLink Server for JDBC Socket must be configured before it can be used.

- "Configuring a SequeLink® Service for JDBC Socket on Windows" on page 393
- "Configuring a SequeLink® Service for JDBC Socket on Linux/UNIX" on page 396
- "Tuning the SequeLink® Server for JDBC Socket Service on z/OS USS" on page 405

Configuring a SequeLink® Service for JDBC Socket on Windows

You can use one SequeLink Server for JDBC Socket to connect with different JDBC drivers. For each JDBC driver you want to access, make a new data source and add the JDBC driver to the CLASSPATH. Changing the CLASSPATH requires a restart of the SequeLink Service for JDBC Socket.

Setting Up the JDBC Driver Environment

Before setting up the SequeLink Server for JDBC Socket environment, ensure you have set up the backend JDBC driver environment and that you can make a connection to this JDBC driver.

Verifying the JDBC Driver Environment

The ivcheckjdbcdriver utility is included in the installation package, and is located in the directory where the installer resides.

Run the ivcheckjdbcdriver utility to make a JDBC connection to your backend JDBC driver. This utility will make a connection, check a number of requirements of the backend JDBC driver and return you the information you need to enter during the configuration of a SequeLink Service for JDBC Socket.

Use this utility before configuring the SequeLink Server for JDBC Socket. You must enter the following configuration information:

- The JRE installation directory
- The JDBC Driver classpath
- The driver classname
- The connection URL of your backend JDBC driver

When the ivcheckjdbcdriver script has verified these settings and successfully made a connection, use these settings for the configuration of your SequeLink Server for JDBC Socket. Read the note included in the information that is returned for specific driver issues.

For more information about the ivcheckjdbcdriver utility, refer to the *SequeLink Troubleshooting Guide and Reference*.

Configuring the SequeLink® Server for JDBC Socket

1. Create a new SequeLink JDBC Socket Service using one of the SequeLink administration tools. Enter the service name and the TCP/IP port and register the SequeLink service.

   During the creation of the SequeLink JDBC Socket Service, the SequeLink JDBC Socket ServiceTemplate (configured during the installation) is copied. In the following steps, check these values or correct these values for your specific service.

2. The SequeLink Server for JDBC Socket starts up the JVM and therefore it needs the path to be set in the SequeLink configuration. This is done via a SequeLink Service environment variable.

   ```
   ServiceEnvironmentVariable:
   PATH= java_install_dir\jre\bin;
   java_install_dir\jre\bin\server
   ```

   This path must contain the jvm.dll and all its dependencies.

3. The SequeLink Service for JDBC Socket attempts to load the JDBC drivers, defined in the server side data sources. All JDBC drivers that need to be accessed must reside in the CLASSPATH. This environment variable needs to be defined in the SequeLink service configuration. Add the CLASSPATH and the classes used by the SequeLink service for JDBC Socket. These classes are installed in *install_dir*\bin\classes.

For example, to configure SequeLink with the Apache Derby 10 driver:

```
CLASSPATH=C:\Program Files\Apache\Derby\derby;
C:\Program Files\DataDirect\slserver60\bin\classes
```

4 (Optional) If you want to specify specific settings for the JVM, set the SequeLink Service environment variable SL_JAVA_OPTIONS.

For example, to set the maximum heap size for the JVM:

```
SL_JAVA_OPTIONS=-Xmx32M
```

For Apache Derby 10, you may have to specify the database directory:

```
SL_JAVA_OPTIONS= -Dderby.system.home=C:\Data\Derby
```

5 Create a new server side data source or use the Default server side data source.

6 Specify the required server-side data source settings for the JDBC driver.

- DataSourceSOCJDBCConnectionURL

  For example, for Apache Derby 10:

  ```
  DataSourceSOCJDBCConnectionURL =
  jdbc:dd-derby10:salesdb
  ```

  Where `database` is the database name you want to access.

- DataSourceSOCJDBCDriverClassName

  For example, for Apache Derby 10:

  ```
  DataSourceSOCJDBCDriverClassName =
  com.ddtek.jdbc.derby10.DerbyDriver
  ```

See Appendix D “SequeLink® Service Attributes” on page 491 for information about these attributes.

7 Specify specific, server-side data source settings, for your JDBC driver, for example, `DataSourceSOCJDBCDbPropertiesName=derby`.

   See Appendix D "SequeLink® Service Attributes" on page 491 for information about these attributes.

8 When the SequeLink Service for JDBC Socket and its server-side data sources are set up, you can save your SequeLink configuration and (re)start the SequeLink Service for JDBC Socket.

9 Run your ODBC, JDBC, .NET or ADO application and access the databases through your JDBC backend driver using the SequeLink Service for JDBC Socket. Refer to the SequeLink Developer's Guide for additional information.

NOTE: If you want to run an ADO application, refer to the *SequeLink Installation Guide*, for information on completing the JDBC Socket Installation for SequeLink Client for ADO.

Configuring a SequeLink® Service for JDBC Socket on Linux/UNIX

You can use one SequeLink Server for JDBC Socket to connect with different JDBC drivers. For each JDBC driver you want to access, make a new data source and add the JDBC driver to the CLASSPATH. Changing the CLASSPATH requires a restart of the SequeLink Service for JDBC Socket.

Setting Up the JDBC Driver Environment

Before setting up the SequeLink Server for JDBC Socket environment, ensure you have set up the backend JDBC driver environment and that you can make a connection to this JDBC driver.

Verifying the JDBC Driver Environment

The ivcheckjdbcdriver utility is packaged with SequeLink Server for JDBC Socket for Linux/UNIX, and installed after the untar.

Run the ivcheckjdbcdriver utility to make a JDBC connection to your backend JDBC driver. This utility will make a connection, check a number of requirements of the backend JDBC driver and return the information you need to enter during the configuration of a SequeLink Service for JDBC Socket.

Use the ivcheckjdbcdriver utility before configuring the SequeLink Server for JDBC Socket. You must enter the following configuration information:

- JVM installation directory
- JDBC Driver CLASSPATH
- Driver class name
- Connection URL of the backend JDBC driver

When the ivcheckjdbcdriver script has verified these settings and successfully made a connection, you can use these settings for the configuration of your SequeLink Server for JDBC Socket. Read the note included in the information that is returned for specific driver issues.

For more information about the ivcheckjdbcdriver utility, refer to the *SequeLink Troubleshooting Guide and Reference*.

Configuring the SequeLink® Server for JDBC Socket

1. Create a new SequeLink JDBC Socket Service using one of the SequeLink administration tools. Enter the service name and the TCP/IP port and register the SequeLink service.

   During the creation of the SequeLink JDBC Socket Service, the SequeLink JDBC Socket ServiceTemplate (configured during the installation) is copied. Some of the settings may already be configured during installation. Check these values or correct these values for your specific service when you encounter them in the next steps.

2. The SequeLink Server for JDBC Socket starts up the JVM and therefore it needs the shared library path to be set in the SequeLink configuration. This is done via a SequeLink Service environment variable.

   On Linux/UNIX, use the shared library path, as follows:

   ```
   shared_library_path
   =java_install_dir/shared_library_directory_for_the_jvm
   ```

   where:

   shared_library_path is one of the following values:

Linux	LD_LIBRARY_PATH
HP-UX	SHLIB_PATH
AIX	LIBPATH
Solaris	LD_LIBRARY_PATH
z/OS USS	LIBPATH

shared_library_directory_for_the_jvm for J2SE 1.4 is one of the following values:

Solaris	lib/sparc/server
HP-UX	lib/PA_RISC/server
AIX	bin/classic
Linux	lib/i386/server
z/OS USS	bin/classic

The shared library of the JVM of JRE 1.4 is one of the following values. It must reside in the *shared_library_directory_for_the_jvm*:

Solaris	libjvm.so
HP-UX	libjvm.sl
AIX	libjvm.a
Linux	libjvm.so
z/OS USS	libjvm.so

3 The SequeLink Service for JDBC Socket attempts to load the JDBC drivers, defined in the server side data sources. All JDBC drivers that need to be accessed must reside in the CLASSPATH. This environment variable needs to be defined in the SequeLink service configuration. Add the CLASSPATH and the classes used by the SequeLink service for JDBC Socket. These classes are installed in *install_dir*/bin/odbc2jdbc/classes.

For example, to configure SequeLink with the Apache Derby 10 driver:

```
CLASSPATH=/opt/Apache/Derby/derby;
/usr/slserver60/bin/classes
```

4 Optionally, if you want to specify specific settings for the JVM, set the SequeLink Service environment variable SL_JAVA_OPTIONS.

For example, to set the maximum heap size for the JVM, type:

```
SL_JAVA_OPTIONS=-Xmx32M
```

For example, for Derby, you may have to specify the database directory:

```
SL_JAVA_OPTIONS= -DDerby.system.home=/usr/Data/derby
```

5 Create a new server side data source or use the Default server side data source.

6 Specify the required server-side data source settings for the JDBC driver.

- DataSourceSOCJDBCConnectionURL

  For example, for Apache Derby 10:

  ```
  DataSourceSOCJDBCConnectionURL =
  jdbc:dd-derby10:salesdb
  ```

  Where `salesdb` is the name of the database that you want to access.

- DataSourceSOCJDBCDriverClassName

  For example, for Apache Derby 10:

  ```
  DataSourceSOCJDBCDriverClassName =
  com.ddtek.jdbc.derby10.DerbyDriver
  ```

See Appendix D “SequeLink® Service Attributes” on page 491 for information about these attributes.

7 Specify specific, server-side data source settings, for your JDBC driver, for example, `DataSourceSOCJDBCDbPropertiesName=derby`.

See Appendix D “SequeLink® Service Attributes” on page 491 for information about these attributes.

8 When the SequeLink Service for JDBC Socket and its server-side data sources are set up, you can save your SequeLink configuration and start the SequeLink Service for JDBC Socket.

9 Run your ODBC, JDBC, .NET or ADO application and access the databases through your JDBC backend driver using the SequeLink Service for JDBC Socket. Refer to the SequeLink Developer's Guide for additional information.

Using a SequeLink® Server for JDBC Socket Service on z/OS USS

Although you can start and stop SequeLink Service for JDBC Socket using the SequeLink Manager Snap-in, you might not want the users to log on to the MMC console and manually start the Service. Furthermore, the SequeLink Agent cannot be started this way because SequeLink Manager Snap-in requires the SequeLink Agent Service to be already up and running.

Four scripts are provided in the `/install_dir/admin` subdirectory to help you automate the starting and stopping of SequeLink services:

- startSLAgent
- stopSLAgent
- startSLSocket2JDBC
- stopSLSocket2JDBC

z/OS

Configuring SequeLink® Service for JDBC Socket on z/OS USS

The configuration procedure for SequeLink Service for JDBC Socket on z/OS USS uses the steps described in "Configuring a SequeLink® Service for JDBC Socket on Linux/UNIX" on page 396.

In addition to defining the connection URL and any required environment variables, you can configure the syslog daemon. See "Writing Syslog Messages" on page 404 for more information.

In addition, read the following sections for platform-specific differences:

- "Starting SequeLink® Services" on page 402
- "Stopping SequeLink® Services" on page 404

Starting SequeLink® Services

The SequeLink Agent Service is used to administer and control the SequeLink Service for JDBC Socket. For that reason, we recommend that you first start the SequeLink Agent Service before starting the SequeLink Service for JDBC Socket.

To start the SequeLink Service for JDBC, Socket use the startSLSocket2JDBC script located in the */install_dir*/admin subdirectory.

You can start the scripts either from a shell or with the BPXBATCH utility program, which allows you to start the processes from a batch Job or Started Task. The SLUSS1 sample in the */install_dir*/admin subdirectory shows how you can achieve this.

Use the following guidelines for starting the services:

1. Define the UID associated with SequeLink Server processes.

The OMVS UID associated with both UNIX processes must be the same, and the UID must be the owner of the SequeLink installation files or be superuser.

During installation, the owner of the SequeLink libraries was set to the UID of the person who installed SequeLink. If you need to change the ownership of the installation files, you can do so with the following command:

```
chown -R owner-uid installation_target_directory
```

where *owner-uid* is the user ID of the new owner and *installation_target_directory* is the SequeLink installation directory.

2 Define which UID is associated with SequeLink Server process.

Typically, the UID is inherited from the UNIX parent process starting the script. However, the method you used to start the service determines which UID is used:

- If started manually from a shell, the UID of the logged in user is inherited.
- If started from a batch Job, the UID of the user submitting the batch Job is inherited.
- If started from a started Task, the UID of the user associated with the started Task is inherited.

3 Optionally, you can change the UID to the value of the Environment variable _BPX_USERID.

- Set and export the _BPX_USERID environment variable in both the startSLAgent and startSLSocket2JDBC scripts to override the standard behavior.
- Ensure the MVS user associated with the parent process has UPDATE access on the BPX.DAEMON profile in the FACILITY class.

4 Define the JobName associated with SequeLink Server process.

   Typically, the user of the parent process is used and appended with a number to form the JobName. You can change this default behavior by setting the Environment variable _BPX_JOBNAME:

   - Set and export the _BPX_JOBNAME environment variable in the scripts. The startSLSocker2JDBC script provides an example.
   - Ensure the MVS user associated with the child process (not the parent process) has READ access to the BPX.JOBNAME profile in the FACILITY class.

Stopping SequeLink® Services

You can start the scripts that stop the SequeLink services from a shell, a batch Job, or a Started Task. Ensure the UID connected to either process has authority to access the SequeLink installation directories.

- To stop the SequeLink Service for JDBC Socket, use the stopSLSocket2JDBC script located in the admin directory.
- To stop the SequeLink Agent Service, use the stopSLAgent script.

Writing Syslog Messages

In addition to the normal SequeLink logging that is written to the logging directory, SequeLink Server for JDBC Socket on z/OS USS writes messages to the z/OS USS syslog. The output destination of these messages is controlled by the syslog daemon.

The syslog daemon is shipped with IBM Communication Server, but it is not configured automatically. To configure the syslog daemon, refer to *Communication Server - IP Configuration*, "Configuring the Syslog Daemon (syslogd)".

By default, SequeLink writes syslog messages to the 'user' facility. You can configure the facility that SequeLink uses by setting the ServiceUnixSysLogFacility service attribute (see "ServiceUnixSyslogFacility" on page 586).

For example, when added to the /etc/syslog.conf file, the following statement sends SequeLink syslog messages to the /var/log/user_sqlnk file:

```
# Write ALL messages from applications, which specify
# facility "user" such as SequeLink_USS log
#
user.* /var/log/user_sqlnk
#
```

Tuning the SequeLink® Server for JDBC Socket Service on z/OS USS

Tuning can make a significant difference in the performance of your application. During the performance and scalability tests of this product, tuning improved throughput and response time in some cases by two to three times.

The SequeLink Server for JDBC Socket Service uses the Java Virtual Machine. You can tune this environment using the ServiceEnvironmentVariable SL_JAVA_OPTIONS. For example, to specify options used by the backend JDBC Driver:

```
ServiceEnvironmentVariable=SL_JAVA_OPTIONS=-Xms32m -Xms64m -Xcomp -verbose
```

In the case of the SequeLink Server for JDBC Socket Service running on z/OS USS, the service runs with the IBM Language Environment (LE) runtime library under UNIX System Services on z/OS. The JDBC driver runs in the Java Virtual Machine (JVM), which also runs with the same IBM LE runtime library.

Tuning of this complex environment involves:

- "Tuning of UNIX System Services" on page 406
- "Tuning of the IBM Language Environment" on page 406
- "Tuning of the Java Virtual Machine (JVM)" on page 407

Tuning of UNIX System Services

Because both the SequeLink Server and the JVM use USS for z/OS for base operating system functions, USS must be correctly installed and tuned to get optimum performance.

These topics are discussed in IBM's document, *UNIX System Services: Planning (GA22-7800-02)*, in the chapter "Tuning Performance." See also the IBM USS performance Web page: http://www-1.ibm.com/servers/eserver/zseries/zos/unix/bpxa1tun.html for detailed tuning suggestions.

Also, using a performance monitor that supports USS is helpful.

Tuning of the IBM Language Environment

The settings of heap and stack sizes and the use of heap pools can positively influence the performance of your SequeLink Server for JDBC Socket on z/OS USS.

Heap and Stack Sizes

All of these storage settings need to be determined for your environment. A Language Environment (LE) storage report generated by the RPTSTG(ON) option shows you how much stack and heap storage is being used, the total number of segments allocated, and the recommended values for the various storage runtime options.

However, RPTSTG(ON) and the STORAGE runtime option can have a negative affect on the performance of your application, because as the application runs, statistics are kept on storage requests. Therefore, always use the IBM-supplied default setting

RPTSTG(OFF) when running production jobs. Use RPTSTG(ON) and STORAGE only when debugging or tuning applications.

Refer to IBM's *Language Environment Programming Guide (SA22-7561-02) and Language Environment Programming Reference (SA22-7562-02)* for details.

Performance and tuning tips for LE can also be obtained from the IBM Web page:
http://www-1.ibm.com/servers/eserver/zseries/zos/le/

LE settings can be set in the SequeLink Server for JDBC Socket service using the environment variable _CEE_RUNOPTS.

For example:

```
_CEE_RUNOPTS=HEAP(128M,8M,ANYWHERE,KEEP,8K,4K),
ANYHEAP(8M,512K,ANYWHERE,FREE),STORAGE(NONE,NONE,NONE,1K),
HEAPP(ON,8,1,32,1,128,2,256,1,1024,4,2048,2)
```

Heap Pools

Tuning the heap pools algorithm for an application is a three-step process that is described in IBM's *Language Environment Programming Guide (SA22-7561-02)*.

Tuning of the Java Virtual Machine (JVM)

For hints and tips, consult the IBM JVM Web page:
http://www-1.ibm.com/servers/eserver/zseries/software/java/javafaq.html

Specific start-up settings for the JVM can be configured in the SequeLink for JDBC Socket service through the SequeLink Service environment variable.

For example, following are the JVM options used for the performance and scalability tests:

```
IBM_JAVA_OPTIONS=-Xms32m -Xmx64m -Xcomp
```

Environment variables and command-line parameters for the JVM can be found in IBM® Developer Kit and Runtime Environment, Java™ 2 Technology Edition, Version 1.4.1, *Service Refresh 1 Diagnostics Guide* (SC34-6309-02).

Configuring SequeLink® Server for DB2 on z/OS

z/OS NOTE: SequeLink DB2 services for z/OS must be created, started, stopped, and deleted locally using the SequeLink Manager for z/OS. See Chapter 7 "Configuring SequeLink® Services Using the SequeLink® Manager for z/OS" on page 129 for more information about creating and managing SequeLink services on z/OS.

Configuring the SequeLink® Service for DB2 with the RRSAF Attachment

NOTE: We highly recommend that you use the ThreadPool connection model, which requires the RRSAF attachment. This setting allows you to have a large number of connections.

To configure a server using the DB2 RRSAF attachment:

1. Type `S` beside the Global Settings node for the server; then, press ENTER.

2 In the AttributeList panel, type `C` next to the MVSGlobalDB2Attachment service attribute. The service attribute panel appears.

```
                DataDirect SequeLink Manager for z/OS

Command ===>


MVSGlobalDB2Attachment
Allowed values are listed below:

Use S or / to select a value.

  Allowed values
-----------------------------------------------------------------
_ CAF

/ RRSAF
```

3 Type `/` next to RRSAF to select it; then, press ENTER.

Configuring the DB2 environment for SequeLink Server involves:

- "Binding the SequeLink® Package" on page 409
- "Granting Access to the SequeLink® Server Package" on page 410
- "Granting Cancel Thread Authorization to the SequeLink® Server" on page 411

Binding the SequeLink® Package

The server-specific CNTL library contains generated BIND jobs to bind the SequeLink DBRMs into packages.

Because SequeLink Server for DB2 can change isolation level dynamically, each DBRM is bound into a package using a different isolation level for each package. The collection ID name of each package is created by concatenating the DB2

collection prefix and the DB2 collection suffix as shown in Table 15-1.

Table 15-1. Isolation Levels of SequeLink DBRMs

Data Source Transaction Isolation	Data Source DB2 Collection Prefix	Data Source DB2 Collection Suffix	Collection ID	Isolation
Uncommitted	SLD600	U	SLD600_U	UR
Committed	SLD600	S	SLD600_S	CS
RepeatableRead	SLD600	T	SLD600_T	RS
Serializable	SLD600	R	SLD600_R	RR

NOTE: The collection *prefix* is set when you create the SequeLink data source (SequeLink service attribute DataSourceDB2CollectionPrefix). The collection *suffix* is set when you configure the transaction isolation level of the data source (SequeLink service attribute DataSourceTransactionIsolation). See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes.

Granting Access to the SequeLink® Server Package

Grant access to the SequeLink Server package for all users that will be connecting to the SequeLink Server. When granting access to the SequeLink package, remember to include the user ID associated with the SequeLink Server for each DB2 subsystem to which the SequeLink Server will connect. To grant access to the SequeLink Server package for all users, you can use the following SQL statement:

```
GRANT EXECUTE ON PACKAGE collection-id.* TO PUBLIC
```

where `collection-id.*` is a collection-id used in the SequeLink bind job. See Table 15-1 for the names of possible collection-ids.

Granting Cancel Thread Authorization to the SequeLink® Server

SequeLink automatically correlates output from DB2 with information from Connect ADO. This means that the SequeLink administrator has more control over the SequeLink sessions, and can take action without spending time to investigate why a killed session does not end.

To request the DB2 database to cancel a DB2 Thread, use the DB2 IFI call interface to route a -CANCEL THREAD(token) command to DB2. Connect ADO only requests canceling a DB2 Thread which is owned by it. DB2 threads owned by any other application are never canceled.

NOTE: This enhancement is not supported for the CAF attachment.

1. To enable this feature, define one of the following DB2 authorizations for the USERID that represents the SequeLink Server address space:

   - SYSOPR
   - SYSCTRL
   - SYSADM

   Refer to your DB2 documentation for more information about the DB2 "COMMAND REFERENCE" -CANCEL THREAD(DB2) command.

2. Regenerate the JCL and run the bind job.

Configuring the SequeLink® Service for DB2 with the CAF Attachment

If you cannot use the DB2 RRSAF attachment, you must change additional configuration settings.

To configure a server using the DB2 CAF attachment:

1. Type `S` beside the Global Settings node for the server; then, press ENTER.
2. In the AttributeList panel, type `C` next to the MVSGlobalDB2Attachment service attribute. The service attribute panel appears.

```
                DataDirect SequeLink Manager for z/OS

 Command ===>


 MVSGlobalDB2Attachment
 Allowed values are listed below:

 Use S or / to select a value.

   Allowed values
 ------------------------------------------------------------------
 / CAF

 _ RRSAF
```

3. Type `/` next to CAF to select it; then, press ENTER.

   NOTE: Selecting CAF as the value for MVSGlobalDB2Attachment implies that a DB2 plan (MVSDataSourceDB2Plan) is required. The MVSDataSourceDB2Plan service attribute is set to SLD600PL by default, but the value can be changed.

4 In the management tree of the SequeLink Service, expand the nodes. Type / next to the Advanced node; then, press ENTER. The AttributeList panel is displayed.

```
          DataDirect SequeLink Manager for z/OS - AttributeList
                                                          Row 1 to 2 of 2
COMMAND ===>                                               SCROLL > PAGE
                                                               MORE >>>

Service ACCT1
Enter the 'ADD'-command to add an attribute or
perform one of the actions below on a specific attribute
o (S)elect    o (C)hange    o (D)elete    o (?)Help

  Name                              Value
-------------------------------------------------------------------------
 _  ServiceConnectionModel          ThreadPool
 _  ServiceMaxSessions              2000
 _  ServiceMaxThreads               64
 _  ServiceMinThread                32
 **************************** Bottom of data ****************************
```

5 Type S beside the attribute you want to change; then, press ENTER. The Attribute Display window appears with the cursor positioned at the Value field of the attribute.

6 Type the new values of the following attributes; then, press ENTER:

- Change ServiceConnectionModel to Thread/Connection
- Decrease ServiceMaxSessions to 250
- Increase ServiceMaxThreads to 256

You are returned to the AttributeList panel, and the attribute values, if valid, are changed.

7 Press F3 to return to the server management tree.

Configuring the DB2 environment for SequeLink Server involves:

- "Binding the SequeLink® Package" on page 414
- "Granting Access to the SequeLink® Server Package" on page 415
- "Binding the SequeLink® Plan" on page 415
- "Granting Access to the SequeLink® Server Plan" on page 416

Binding the SequeLink® Package

The server-specific CNTL library contains generated BIND jobs to bind the SequeLink DBRMs into packages.

Because SequeLink Server for DB2 can change isolation level dynamically, each DBRM is bound into a package using a different isolation level for each package. The collection ID name of each package is created by concatenating the DB2 collection prefix and the DB2 collection suffix as shown in Table 15-2.

Table 15-2. Isolation Levels of SequeLink DBRMs

Data Source Transaction Isolation	Data Source DB2 Collection Prefix	Data Source DB2 Collection Suffix	Collection ID	Isolation
Uncommitted	SLD600	U	SLD600_U	UR
Committed	SLD600	S	SLD600_S	CS
RepeatableRead	SLD600	T	SLD600_T	RS
Serializable	SLD600	R	SLD600_R	RR

NOTE: The collection *prefix* is set when you create the SequeLink data source (SequeLink service attribute DataSourceDB2CollectionPrefix). The collection *suffix* is set when you configure the transaction isolation level of the data source (SequeLink service attribute DataSourceTransactionIsolation). See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes.

Granting Access to the SequeLink® Server Package

Grant access to the SequeLink Server package for all users that will be connecting to the SequeLink Server. When granting access to the SequeLink package, remember to include the user ID associated with the SequeLink Server for each DB2 subsystem to which the SequeLink Server will connect. To grant access to the SequeLink Server package for all users, you can use the following SQL statement:

```
GRANT EXECUTE ON PACKAGE collection-id.* TO PUBLIC
```

where `collection-id.*` is a collection-id used in the SequeLink bind job. See Table 15-2 for the names of possible collection-ids.

Binding the SequeLink® Plan

The server-specific CNTL library contains a generated BIND job to bind the SequeLink Server packages into plans.

The SequeLink plan name is set when you create the SequeLink data source (SequeLink service attribute MVSDataSourceDB2Plan). For a list of SequeLink service attributes, refer to the *SequeLink Installation Guide*.

Granting Access to the SequeLink® Server Plan

When you are setting up your server to use the CAF attachment, you must grant access to the SequeLink plan for all users who will connect to the SequeLink Server. When granting access to the SequeLink plan, remember to include the user ID associated with the SequeLink Server for each DB2 subsystem to which the SequeLink server will connect.

To grant access to the SequeLink plan for all users, take the following steps:

1. Use the following SQL statement:

   ```
   GRANT EXECUTE ON PLAN planname TO PUBLIC
   ```

   where `planname` is the name of the SequeLink plan.

   Refer to your DB2 documentation for more information about the DB2 "COMMAND REFERENCE" -CANCEL THREAD(DB2) command.

2. Regenerate the JCL and run the bind job.

Using z/OS Workload Manager (WLM)

Workload Manager (WLM) manages *workloads* and optimizes throughput based on policies defined by the WLM administrator. A workload is a collection of service classes that are tracked and reported as a unit. WLM distributes the available computer resources to achieve the goals defined in the policy. You can define different performance goals and processing rules for different workloads. For example, you can assign priorities so that the workload for your Customer Orders department has a higher priority than the workloads for the Human Resources and Marketing departments, which do not require the same level of service.

The SequeLink Server for z/OS supports WLM by using *enclaves*. An enclave is a z/OS construct that serves as a unit of priority and

accounting for a transaction. You define performance criteria for an enclave using service policies, which are a collection of performance goals and processing capacity rules. You can use service policies to define groups of service classes that can be used for workloads with similar performance requirements. For example, you can define a service class that requires an average response time, sets a high business importance, and uses a resource group that guarantees a minimum amount of CPU is granted even if the enclave does not meet the response time goal. For more information about defining service classes, refer to your IBM documentation.

When SequeLink is enabled for WLM and a client connection is made to the Server, an independent enclave is created and classification information is passed to WLM. See Appendix A "z/OS Workload Manager (WLM) Classification" on page 437 for more information.

This classification information allows the WLM administrator to define rules so that WLM can determine the "Service class" to be used for the enclave(s) created for this connection. Each enclave is managed separately according to its performance goal or performance group.

When the connection to the SequeLink Server ends, the CPU consumption of all of the enclaves for this connection is reported in the SMF record, if activated (see "Configuring SMF Accounting" on page 420).

Configuring SequeLink® Server for Workload Management

The default setting MVSGlobalWLMEnclaves =NONE disables WLM support in the SequeLink Server for DB2 for z/OS. To enable WLM for a SequeLink Server, you must use one of the following settings for the MVSGlobalWLMEnclaves attribute:

- MVSGlobalWLMEnclaves =CONNECTION. SequeLink creates a single WLM enclave for each new incoming client connection. This enclave exists for the duration of the connection. Because the enclave includes client think-time and network wait-time, it is not appropriate to define WLM velocity goals.
- MVSGlobalWLMEnclaves=RPC. SequeLink creates a new WLM enclave for each SequeLink RPC (each network access to the SequeLink Server). This setting allows the WLM administrator to specify response time goals without regard to network wait time or client application processing time.

When MVSGlobalWLMEnclaves is set to either CONNECTION or RPC, SequeLink Server for DB2 for z/OS connects during startup to WLM as a Workmanager of type VAI. The instance name used on the initial connection to WLM is the value of the MVSGlobalSubSysID attribute (see "MVSGlobalSubSysID" on page 553).

Additionally you must define classification rules in WLM. Appendix A "z/OS Workload Manager (WLM) Classification" on page 437 gives an overview of ALL information passed to WLM by SequeLink. You can select the criteria you need to classify one or more connections to a Service Class.

SequeLink® Cluster in a Sysplex

To achieve high availability, ensure a consistent response time, and provide a future path of growth for your applications, IBM introduced the Sysplex cluster technology. Beginning with OS/390 V2R7, the TCP/IP stack changed to fully exploit this new technology, and now supports TCP/IP connection distribution in a Sysplex. Three different implementation models are available:

- **Sysplex distributor**
- **Domain Name System (DNS)/Workload Manager (WLM)**
- **Network dispatcher**

All three implementations use multiple Server images distributed over your Sysplex cluster, which appear to your clients as a single server application. New incoming TCP/IP client connections are then distributed over the server images based on workload and availability information.

Not all solutions require support from Sequelink Server for DB2 for z/OS. The Sysplex distributor and network dispatcher models can be implemented without any specific interaction with Sequelink. The DNS/WLM solution requires Sequelink to register to WLM during startup of the Server.

The MVSGlobalClustername attribute has been added to support the DNS/WLM implementation model (see "MVSGlobalClustername" on page 549 for more information). All servers on the different clusters must use the same clustername to register to WLM and thus appear as a single application to your clients. DNS regularly queries WLM and based on this information, routes new TCP/IP connection requests for this clustername to the appropriate server image.

The IBM Redbook, *TCP/IP in a Sysplex*, document number SG24-5235-02 explains in detail the different implementation models. If you are considering the use of TCP/IP in a Sysplex environment, this book is required reading.

Configuring SMF Accounting

When SMF accounting is enabled, records are written to the active SMF data set using the SMF record ID at the end of each client connection. SMF accounting is enabled by default and the SMF record ID is configured using the GlobalSMFRecordType attribute, which is a Global Settings attribute of the SequeLink Server. If you do not want to use the SMF accounting facility, you can turn it off by deleting the GlobalSMFRecordType attribute.

SMF records contain statistical data that can be used to track information for charge-back systems. The following data is provided for each database session:

- Job/STC name of the server
- Client logon ID
- Internal Thread ID assigned by SequeLink Server
- Mapped user ID
- Session start/end times
- Number of packets and bytes received/sent
- Total CPU consumption
- Service name
- Application loadmodule name
- Network node name of the client
- DataSource name
- Thread abend/abort information
- Abend code
- Number of SQL calls, checkpoint calls, and DB2 OPEN calls
- Number of records switched for ReadTimeout
- Number of records switched for Max RPC
- Accumulated CPU consumption and elapsed time for SQL calls
- WLM ServiceClass

The following example shows a typical SMF record. A sample job (RUNSMF) to print out the SMF records is provided in the SequeLink_HLQ.CNTL library.

```
*********************************************************************
*                  SMF RECORD HEADER
*********************************************************************
SMFRECORD   DS     0F
SMFRECLEN   DS     H        RECORD LENGTH
SMFDESC     DS     H        DESCRIPTOR (MUST BE ZERO)
SMFSYS      DS     X        SYSTEM INDICATOR (SET TO ZERO)
SMFTYPE     DS     X        Record type (MVSGlobalSMFRecordType)
SMFTIME     DS     AL4      CURRENT TIME
SMFDATE     DS     AL4      CURRENT DATE
SMFSID      DS     CL4      SYSTEM ID
SMFSUBS     DS     CL4      SERVER SUBSYSTEMID OR BLANKS
SMFHDRLEN EQU   *-SMFRECORD HEADERLENGTH
SMFUSER     DS     0X       START OF SUBTYPE FIELD(S)
*
*  VARIABLE USER FIELDS START HERE
*
********************************************************************
* SUBTYPE X'0001' - SESSION ACCOUNTING RECORD                      *
********************************************************************
@ACCTREC  DSCET
@ACCTLEN    DS    XL2       Accounting record length
@ACCTTYP    DS    XL2       Accounting record subtype (X'0001')
*
@ACCTJNM  DS   CL8      Serv STC/JOB/Name
@ACCTUID    DS    CL8       RACF Userid connected to thread
@ACCTTID    DS    CL8       Thread ID
@ACCTAPL    DS    CL8       Application name
@ACCTSVC    DS    CL8       Service name
*
@ACCTNOD    DS    CL16      Client node name (TCP/IP address)
@ACCTCNT    DS    F         Message count
@ACCTBIN    DS    F         Total Input packet byte count
@ACCTBOT    DS    F         Total Output packet byte count
*
* Session Start - Time and date (SMF header type)
*
```

```
@ACCTBGT    DS    F      Time since Midnight in 100th of second
@ACCTBGD    DS    PL4    Date in 0CYYDDDF format -  C is 1 if year is 20YY,
                         is 0 if 19YY
*
*  Session Ending - Time and date  ( SMF header type )
*
@ACCTENT    DS    F      Time since Midnight in 100th of second
@ACCTEND    DS    PL4    Date in 0CYYDDDF format
*
@ACCTCPU    DS    F      CPU TIME used - measured in 100th of a second
*
@ACCTERR    DS    X
ACCTABND    EQU   1      Thread abended - abend code in @ACCTABN
ACCTABIT    EQU   2      Idle Timeout
ACCTABOP    EQU   4      Aborted by Operator
ACCTMXCP    EQU   8      Aborted for MaxCPU
*
@ACCTABN    DS    XL3    Left  12 bits - Systemabend code in Hex
*                        Right 12 bit  - Userabend   code in Hex
*
@ACCTDBC    DS    F      Database Calls - count
@ACCTDBO    DS    F      Database Opens - count
@ACCTCKP    DS    F      Database Checkpoints - count
@ACCTSQC    DS    F      Accum SQL CPU time in 100th of seconds
@ACCTSQE    DS    F      Accum SQL Elapsed time in 100th of seconds
@ACCTWSC    DS    CL8    WLM ServiceClass or blanks

@ACCTUID2   DS    CL8    Mapped Userid - Primary authid DB2
@ACCTDSRC   DS    CL64   DataSource Name
@ACCTSWTO   DS    F      Count - switched for ReadTimeout
@ACCTSWMR   DS    F      Count - switched for Max RPC

@ACCTLTH  EQU  *-@ACCTREC  SUBTYPE LENGTH
```

If you have your own version of an SMF processing program you can check the accounting record length if you need to differentiate your SMF records between the previous version layout and the new SMF record layout. For example:

```
@ACCTLEN DS    XL2           Accounting record length
```

Adding a UID Map to the SequeLink® Service

UID mapping is the optional mapping of user IDs to alternate user IDs using a UID map. You can use UID mapping to prevent users from updating DB2 tables using commonly available tools, such as QMF or SPUFI, while preserving their ability to update DB2 tables using SequeLink. A UID map is an independent entity in the SequeLink configuration and can be specified at the level of a SequeLink service or/and at the level of a SequeLink Server data source. (See "MVSServiceUIDMap" on page 557 for more information.)

UID mapping is required when ServiceAuthMethods=Anonymous.

You can specify a UID map for a SequeLink service and another UID map for a SequeLink Server data source. See "Using UID Mapping" on page 329 for more information about using UID maps.

To add a UID map:

1. From the server management tree, type `A` beside the UID Maps node to add a UID map to the SequeLink Server; then, press ENTER. The Add UID Map panel appears.

```
      DataDirect SequeLink Manager for z/OS - Add UID Map

Complete the following fields for the new mapping table:

UID Map Name . . . .
Default access . . .                (PERMIT / DENY)

Description

Command ===>
```

2 Provide the following information; then, press ENTER.

UID Map Name: Type the name of the UID map to add to the SequeLink Server. The corresponding service attribute is MVSUIDMap.

Default access: Choose one of the following options for the default behavior of the UID map you are adding by typing:

- `PERMIT`. If the user ID cannot be found in the UID map, the connection request is accepted. The userid is passed to DB2 as the primary authid.
- `DENY`. If the user ID cannot be found in the UID map, the connection request is refused.

The corresponding service attribute is MVSUIDDefaultAccess.

Description: Type a description of the UID map.

The corresponding service attribute is MVSUIDMapDescription.

3 If the UID map was added successfully, a message appears to confirm it.

4 To add mapping entries to the UID map, type `S` beside the UID map; then, press ENTER. The AttributeList panel appears.

```
        DataDirect SequeLink Manager for z/OS - AttributeList
                                                        Row 1 to 2 of 2
                                                               MORE >>>

  UID Map UIDACCT
  Enter the 'ADD'-command to add an attribute or
  perform one of the actions below on a specific attribute
  o (S)elect    o (C)hange    o (D)elete    o (?)Help

    Name                            Value
 ---------------------------------------------------------------------------
    MVSUIDDefaultAccess             PERMIT
    MVSUIDMapDescription            UID map for ACCT1
  **************************** Bottom of data ****************************

  COMMAND ===>                                            SCROLL > PAGE
```

5 Type `Add` at the command prompt; then, press ENTER. A panel appears allowing you to specify a UID mapping entry for the MVSUID attribute.

```
                    Attribute Value

Press F1 for help, F3 to leave.

MVSUID
    ===>
```

6 Specify a value for the MVSUID attribute using the format user=mapped_user or *=mapped_user where:

- user is a valid user or user group for the z/OS security system.
- * represents any user. This value is required when ServiceAuthMethods=Anonymous.
- mapped_user is a valid DB2 authorization ID.

Then, press ENTER.

7 You are returned to the AttributeList panel. Add another MVSUID entry or press F3 to return to the server management tree.

8 To configure a SequeLink service to use the UID map, set the MVSServiceUIDMap service attribute and/or MVSDataSourceUIDap attribute.

NOTE: You must restart the SequeLink Server before this change will take effect.

See "Using UID Mapping" on page 329 for more information about using UID maps.

Configuring a SequeLink® Service for Oracle

The default for SequeLink Server for Oracle is to use the Oracle Bequeath protocol to communicate to the Oracle database engine. However, we strongly recommend that during installation, you select the Oracle Net Service instead. If you accept the default driver during installation and later want to use the TCP or IPC protocol through Oracle Net Service, you can change the configuration.

To use the Oracle Net Service:

1. Configure the Oracle listener for Net Service (refer to your Oracle documentation for instructions).
2. Configure the SequeLink Oracle service to use the Oracle Net Service. Using the SequeLink Manager, add the SequeLink service attribute ServiceEnvironmentVariable to your SequeLink Oracle service:
   - On Windows, set the value to ServiceEnvironmentVariable=LOCAL=NetServiceName where NetServiceName is the Oracle Net service name you configured in Step 1.
   - On UNIX, set the value to ServiceEnvironmentVariable=TWO_TASK=NetServiceName where NetServiceName is the Oracle Net service name you configured in Step 1.
3. If the SequeLink service attribute ServiceEnvironmentVariable=ORACLE_SID=OracleSid exists for this SequeLink service, delete it.
4. Restart the SequeLink Oracle service.

16 Using LDAP with the SequeLink® Clients

This chapter explains how SequeLink Clients use LDAP directories to retrieve connection information and describes how to create and update LDAP entries for SequeLink services.

What is LDAP?

SequeLink ODBC, JDBC, and ADO Clients can connect directly to a SequeLink Server or retrieve connection information from a Lightweight Directory Access Protocol (LDAP) directory. LDAP is a standard protocol for accessing and updating common directory information. Storing connection information centrally in an LDAP directory provides flexibility to make environment changes and reduces the time it takes to reconfigure your infrastructure when a change takes place.

For example, if a database must be moved to a different server, you do not have to reconfigure the user applications or the client data sources that must now access the new server. Because the connection information is stored in an LDAP directory, you need only update the LDAP directory entries so that the SequeLink Clients can connect to the new server.

SequeLink supports LDAP V3.

Retrieving Connection Information from LDAP Directories

When a SequeLink Client retrieves connection information from an LDAP directory, the connection to the SequeLink Server takes place as a two-step process as shown in Figure 16-1.

1 The SequeLink Client connects to the LDAP server and retrieves connection information from an LDAP entry.

2 Using the connection information retrieved from the LDAP entry, the SequeLink Client connects to the SequeLink Server.

Figure 16-1. Retrieving Connection Information from an LDAP Directory

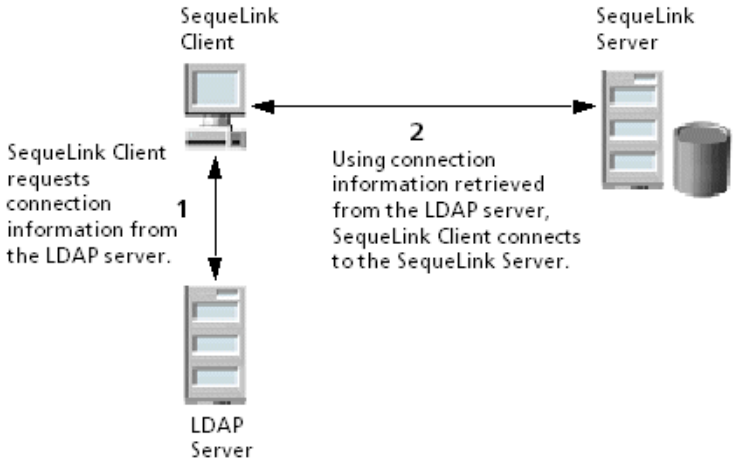

To retrieve connection information from an LDAP directory, you configure the following information at the SequeLink Client:

LDAP Server Host	is the TCP/IP host name of the LDAP server.
LDAP Server Port	is the TCP/IP port that the LDAP server is listening on for incoming connection requests. If unspecified, the SequeLink Client will use the default LDAP port 389.
Distinguished Name (DN)	is an identifier that uniquely identifies the LDAP entry where connection information is stored. See "Creating LDAP Entries for SequeLink® Services" on page 430 for information about what you must include in an LDAP entry.

For information about configuring ODBC Clients and ADO Clients to connect to an LDAP server to retrieve connection information:

- See Chapter 8 "Configuring the ODBC Client" on page 169
- See Chapter 9 "Configuring the ADO Client" on page 211

For JDBC Clients, JDBC data sources store connection instructions in a JNDI (Java Naming and Distribution Interface) infrastructure, which can support LDAP. See "Using an LDAP Server as Your JNDI Service Provider" on page 433 for more information.

See "OpenLDAP Support" on page 431 for information about the SequeLink LDAP schema definition for OpenLDAP.

Creating LDAP Entries for SequeLink® Services

If you are using LDAP directories to store connection information for SequeLink services, you must create a single LDAP directory entry for each SequeLink service and identify that LDAP directory entry by a Distinguished Name (DN). The LDAP entry must contain the TCP/IP host name of the SequeLink Server and the TCP/IP port on which the SequeLink Server is listening. Optionally, the LDAP entry can contain a server data source to use for the connection to the SequeLink Server. If a server data source is not specified, the default data source is used by the client.

Table 16-1 shows an example of an LDAP directory entry and attributes you may want to store in an LDAP entry for a SequeLink service.

Table 16-1. Example: LDAP Directory Entry for a SequeLink Service

Attribute	Syntax	Description	Example	Required?
SequeLinkHost	cis	The TCP/IP host name or the TCP/IP address on which the SequeLink Server is running.	123.4.5.6 dino.yourcompany	Yes
SequeLinkPort	cis	The TCP/IP port on which the SequeLink Server is listening.	5004	Yes
ServerDataSource	cis	A server data source to use for the connection.	Sales	No

How you create LDAP entries depends on the LDAP product you are using. For instructions on creating LDAP entries, refer to your LDAP product documentation.

Updating LDAP Entries for SequeLink® Services

You must update LDAP directory entries for SequeLink services when a:

- New SequeLink service is installed or added
- SequeLink service is assigned a new IP address or port
- SequeLink data source is created or deleted
- SequeLink data source is renamed

For instructions on updating LDAP directory entries, refer to your LDAP product documentation.

OpenLDAP Support

The ODBC Client and ADO Client use the Open LDAP v3 library. This library is developed and maintained by The OpenLDAP Project (http://www.openldap.org/).

SequeLink LDAP Schema Definitions

NOTE: The SequeLink Server for DB2 on z/OS cannot be configured to support SSL. Do not use the SequeLinkEncrypted encryption attribute in a DB2 for z/OS configuration.

The following LDAP schema definition is applicable to both the ODBC Client and the ADO Client.

```
# SequeLink ODBC schema
# ADO Client uses the same schema
#
#
```

```
objectIdentifier ourenterpriseOID 1.1.2.3.4.5.6
objectIdentifier sequelinkOID ourenterpriseOID:8

attributetype (sequelinkOID:200001
    NAME 'SequeLinkHost'
    DESC 'SequeLink Host'
    EQUALITY caseIgnoreMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype (sequelinkOID:200002
    NAME 'SequeLinkPort'
    DESC 'SequeLink Port'
    EQUALITY caseIgnoreMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.15)

attributetype (sequelinkOID:200003
    NAME 'SequeLinkServerdatasource'
    DESC 'SequeLink Server DataSource'
    EQUALITY caseIgnoreMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.15)

attributetype (sequelinkOID:200004
    NAME 'SequeLinkEncrypted'
    DESC 'SequeLink Encrypted'
    EQUALITY caseIgnoreMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.15)

objectclass (sequelinkOID:100001
    NAME 'SequeLinkODBCDSN'
    SUP top STRUCTURAL
    MUST (SequeLinkHost $ SequeLinkPort )
    MAY ( cn $ SequeLinkServerdatasource $
SequeLinkEncrypted ) )
```

Example: LDAP Entry for ODBC

```
# usercn.ldiff template
# create OVS datasource
```

```
# USER     Arne
# DATABASE SequeLink
# HOST     belg-keidis
# PORT     6007
#
dn: cn=DB2V8 on belg-keidis,ou=OVS,ou=Arne,ou=USERS,o=
Development
cn: DB2V8 on belg-keidis
objectClass: top
objectClass: SequeLinkODBCDSN
SequeLinkPort: 6007
SequeLinkHost: belg-keidis
SequeLinkServerDatasource: Default
SequeLinkEncrypted: 0
```

Using an LDAP Server as Your JNDI Service Provider

You can use an LDAP server as your JNDI service provider for storing connection information for the JDBC Client. This information has been taken from RFC 2713 Schema for Representing Java Objects in an LDAP Directory, which can be found at http://www.faqs.org/rfcs/rfc2713.html.

NOTE: The SequeLink Server for DB2 on z/OS cannot be configured to support SSL. Do not use the encrypted encryption attribute in a DB2 for z/OS configuration.

An example of an LDAP schema definition for JDBC is shown below:

```
dn: cn=DB2V8 on belg-keidis,ou=JVS,ou=Arne,ou=USERS,o=Development
cn: DB2V8 on belg-keidis
objectClass: top
objectClass: javaContainer
objectClass: javaObject
objectClass: javaNamingReference
```

```
javaReferenceAddress: #0#description#
javaReferenceAddress: #1#portNumber#6007
javaReferenceAddress: #2#serverName#belg-keidis
javaReferenceAddress: #3#databaseName#jvs
javaReferenceAddress: #4#serverDataSource#Default
javaReferenceAddress: #5#encrypted#1
javaFactory: com.ddtek.jdbcx.sequeLink.SequeLinkDataSourceFactory
javaClassName: com.ddtek.jdbcx.sequelink.SequeLinkDataSource
```

Part 4: Reference

This part contains the following appendixes:

- Appendix A "z/OS Workload Manager (WLM) Classification" on page 437 describes the information used by SequeLink Server to classify WLM enclaves.
- Appendix B "SequeLink® Manager Commands" on page 439 lists all available SequeLink Manager commands.
- Appendix C "Operator Interface Commands for z/OS" on page 479 lists all available Operator Interface commands by category.
- Appendix D "SequeLink® Service Attributes" on page 491 lists the SequeLink Manager attributes you can use to configure and manage your SequeLink environment.
- Appendix E "SequeLink® Events" on page 589 lists and defines the SequeLink events, the attributes associated with events, and explains how to write a filter for an event.
- Appendix F "Internationalization, Localization, and Unicode" on page 597 provides an overview of how internationalization, localization, and Unicode relate to each other.

A z/OS Workload Manager (WLM) Classification

This appendix describes the information used by SequeLink Server to classify WLM enclaves.

IMPORTANT: The IBM WLM ISPF application allows you to customize WLM classification criteria, using a starting position of 1 for the subsystem parameter. Because Table A-1 uses a zero-based offset, you must add 1 to convert from the offset to the starting position.

Table A-1. Workload Manager (WLM) Classification Information

WLM Qualifier Name	Value
Subsystem Type	VAI
Subsystem Instance	SequeLink Subsystem Name (MVSGlobalSubSysID attribute)
Function Name	SQLNK
Collection Name	DB2 Collection Name used when the enclave is created
Connection Type	One of the following client types: ■ ODBC ■ JDBC ■ OLEDB ■ ADMIN ■ .NET
Correlation Information	Connection ID, for example, T*xxxxxxx*
Package Name	DB2 Package Name used when the enclave is created
Plan Name	One of the following plan names: ■ DB2 Plan Name for attachment CAF ■ ?RRSAF for attachment RRSAF

Table A-1. Workload Manager (WLM) Classification Information *(cont.)*

WLM Qualifier Name	Value	
Subsystem Parameter	Information presented in the following pattern:	
	0-7	Space character
	9-23	Client TCP/IP address
	24	Space character
	25-39	Server TCP/IP address
	40	Space character
	41-46	Port number of the SequeLink Server
	47	Space character
	48-55	Original user ID
	56	Space character
	57-64	Mapped UID
	65	Space character
	66-73	SequeLink service name
	74	Space character
	75-138	SequeLink server data source name
Transaction Name / Job Name	SequeLink Server jobname	
Userid	Mapped UID	

B SequeLink® Manager Commands

This appendix lists all available SequeLink Manager commands in alphabetical order. See "Invoking the SequeLink® Manager Command-Line Tool" on page 103 for information about starting the SequeLink Manager Command-Line Tool.

IMPORTANT:

- SequeLink Manager command names are not case-sensitive; however, the command parameter *service_name* is case-sensitive.
- If the value of a command parameter contains spaces, the value must be enclosed within single quotes (') or double quotes (").
- If the value of a command parameter contains single or double quotes, use single quotes to quote double quotes and double quotes to quote single quotes.
- The pound sign (#) is a comment character. All text that follows the pound sign on the same line is ignored.
- When issuing commands using the direct or batch method, you must specify all of a command's required and optional parameters in the correct position (specify the parameters in the order they are documented). See "Direct Command Execution" on page 100 and "Batch Command Execution" on page 101 for more information about direct and batch methods of issuing commands.

About | a

Prints to the console the version of the SequeLink Manager Command-Line Tool and copyright information.

Syntax `{about | a}`

Example `about`

ActivateLocalConfig | alc

Connects to the SequeLink Agent on a local machine.

Syntax {`ActivateLocalConfig | alc`}

Example `alc`

ActivateOfflineConfig | aoc

Opens the local configuration file.

Syntax {`ActivateOfflineConfig | aoc`} *`configuration_file`*

where *`configuration_file`* is the path and name of local configuration file.

Example `aoc "c:\Program Files\DataDirect\slserver60\cfg\swandm.ini"`

NOTE: The path is enclosed in quotes because there is a space in the path name. Otherwise, the path would not have to be quoted, for example, `aoc c:\DataDirect\cfg\swandm.ini`.

ActivateRemoteConfig | arc

Connects to a remote SequeLink Agent.

Syntax

{`ActivateRemoteConfig | arc`} *`agent_connection_info`*

where *`agent_connection_info`* is the information required to connect to the SequeLink Agent—host:port, which is the host name and port of the server on which the SequeLink Agent resides.

Example

```
arc caspar.ddtek.com:19996
```

ActivateSecureRemoteConfig | asrc

On Linux/UNIX/Windows, connects to a remote SequeLink Agent using SSL.

Syntax

{`ActivateSecureRemoteConfig | asrc`} *`agent_connection_info`*

where *`agent_connection_info`* is the information required to connect to the SequeLink Agent—*host*:*port*, which is the host name and port of the server on which the SequeLink Agent resides.

Example

```
asrc caspar.ddtek.com:19996
```

CloseConfig | cc

Closes the activated configuration.

Syntax

{`CloseConfig | cc`}

Example

```
cc
```

DataSourceAttributeAdd | dsaa

Adds an attribute to the specified server data source and the specified SequeLink service.

Syntax

{`DataSourceAttributeAdd | dsaa`} *`service_name data_source_name attribute_name value`*

where:

`service_name` is the name of a data access service. Service names can be obtained using the ServiceList | sl command.

`data_source_name` is the name of the server data source for which you want to add an attribute. This server data source must belong to the data access service you specified.

`attribute_name` is the name of the attribute you want to add.

value is the value for the attribute.

See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes and their valid values.

Example

In the following example, the DataSourceCurrentCatalog attribute is being added to the DS_Employees server data source that belongs to the SLOracle10 data access service. The value for the attribute is employees.

```
dsaa SLOracle10 DS_Employees DataSourceCurrentCatalog
employees
```

DataSourceAttributeDelete | dsad

Deletes an attribute from a server data source.

Syntax

{`DataSourceAttributeDelete | dsad`} *`service_name data_source_name attribute_name`*

where:

`service_name` is the name of a data access service. Service names can be obtained using the ServiceList | sl command.

`data_source_name` is the name of the server data source from which you want to delete an attribute. This server data source must belong to the data access service you specified.

`attribute_name` is the name of the attribute you want to delete.

See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes and their valid values.

Example

The following example deletes the DataSourceCurrentCatalog attribute from the DS_Employees server data source that belongs to the SLOracle10 data access service:

```
dsad SLOracle10 DS_Employees DataSourceCurrentCatalog
```

DataSourceAttributeReplace | dsar

Changes the value of a server data source attribute.

Syntax

{`DataSourceAttributeReplace | dsar`} *`service_name data_source_name attribute_name value`*

where:

`service_name` is the name of a data access service. Service names can be obtained using the ServiceList | sl command.

data_source_name is the name of the server data source you want to modify. This server data source must belong to the data access service you specified.

attribute_name is the attribute for which you want to change the value.

value is the new value of the attribute.

See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes and their valid values.

Example

The following example changes the value of the DataSourceCurrentCatalog attribute in the DS_Employees server data source that belongs to the SLOracle10 data access service. The value is changed to partners.

```
dsar SLOracle10 DS_Employees DataSourceCurrentCatalog
partners
```

DataSourceCreate | dsc

Creates a server data source.

Syntax

{`DataSourceCreate | dsc`} *service_name* *data_source_name*

where:

service_name is the name of the data access service in which to create the new server data source.

data_source_name is the name of the server data source you want to create.

Example

The following example creates a server data source named DS_Employess within the SLOracle10 data access service:

```
dsc SLOracle10 DS_Employees
```

DataSourceDelete | dsd

Deletes a server data source.

Syntax

{`DataSourceDelete | dsd`} *`service_name`* *`data_source_name`*

where:

`service_name` is the name of the data access service from which to delete a server data source. Service names can be obtained using the ServiceList | sl command.

`data_source_name` is the name of the server data source you want to delete.

Example

The following example deletes the DS_Employees server data source from the SLOracle10 data access service:

```
dsd SLOracle10 DS_Employees
```

DataSourceInfo | dsi

Lists all attributes and their values for a server data source.

Syntax

{`DataSourceInfo | dsi`} *`service_name`* *`data_source_name`*

where:

`service_name` is the name of a data access service. Service names can be obtained using the ServiceList | sl command.

`data_source_name` is the name of the server data source for which you want to list attributes and their values. This server data source must belong to the data access service you specified.

See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes and their valid values.

Example

The following example lists attributes and their values for the DS_Employees server data source. This server data source belongs to the SLOracle10 data access service.

```
dsi SLOracle10 DS_Employees
```

DataSourceList | dsl

Lists all available SequeLink server data sources.

Syntax

{`DataSourceList | dsl`}

Example

```
dsl
```

Echo | e

Echoes a user-defined string. This command is useful when you issue commands from a file. For example, you could issue the DataSourceCreate command and use the Echo command to display text that says "A new data source was created in the SLOracle10 data access service."

Syntax

{`Echo | e`} *`string`*

where *`string`* is the text you want displayed to your console. If the string contains spaces, you must surround the string with either double or single quotes.

Example

```
echo "Testing echo command."
```

EventList | el

Lists events from a specified data access service or from a specified event trace file.

Syntax

{`eventlist | el`} *`service_name`* | [remote]file=
`event_trace_file_name`
[details]
[[{service | srvc}] |
[{session | sess}] |
[{statement | stmt}] |
[{transaction | trans}] |
[{network | net}] |
[{error | err}] |
[{other | oth}]]]
[count=[{ + | - }] {all | *`number`*}]
[offset={begin | end} [{ + | - }]*`number`*]
[query='*custom_event_filter_string*']

where:

`service_name` is the name of a data access service. The event trace file for this specified service will be listed. Service names can be obtained using the ServiceList | sl command.

`event_trace_file_name` is the path and name of the event trace file you want to list.

`number` is the event number of events to list when used with the count option. When used with the offset option, number is the number of the event from which to start listing events. For example, if offset=10, SequeLink would list all events starting with event 10. Another example, if you specify count=20 and offset=begin, SequeLink will list the first 20 events. If you specify count=20 and offset=5, SequeLink will list 20 events starting from event 5.

`custom_event_filter_string` is an event filter statement. See "Filtering Events" on page 595 for more information.

Options

Details: If you specify Details, SequeLink will list detail information about the event.

Event Types: You can specify one or more of the following event groups for which to list event information: Service, Session, Statement, Transaction, Error, or Other. When you specify one or more of the event groups, SequeLink Manager lists all the events of the type you specified. For example, if you specify Service, SequeLink lists all service events such as Service Started and Service Stopping (these events start with "Service"). Event names that do not start with Service, Session, Network, Error, Statement, or Transaction are Other events (for example, Cursor Closed).

Examples

Local host or remote configuration examples:

Example A: The following example lists detailed event information for the SLAgent service:

```
el SLAgent details
```

Example B: The following example lists detailed event information for all service events starting with event number 10 in the event trace file associated with the SLOracle data access service:

```
el SLOracle details service count=all offset=10
```

Example C: The following example defines a query for the information it will list for the SLOracle10 data access service. The query returns all SQL statements that do not return a return code of 0.

```
el SLOracle10 stmt query='${ReturnCode} != 0'
```

Example D: The following example lists event information for the first 10 events in the local file named SLoracle10.trc:

```
el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLoracle10.trc" count=10
```

Example E: The following example lists detailed event information for only service events in the remote file named SLoracle10.trc:

```
el "remotefile=C:\Program Files\DataDirect\slserver60\
tracing\SLOracle10.trc" service details
```

Offline configuration examples:

Example A: The following example lists event information from the local file named SLoracle10.trc file:

```
el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc"
```

Example B: The following example lists event information for the first 10 events in the local file named SLoracle10.trc:

```
el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" count=10
```

Example C: The following example lists event information from the local file named SLoracle10.trc, starting from the end of the file and listing all events:

```
el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" count=-all offset=end
```

Example D: The following example lists only session and service event information from the local file named SLoracle10.trc, starting with the fifth event:

```
el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" offset=5 service session
```

Example E: The following example lists detailed event information for service events only from the local file named SLoracle10.trc:

```
el "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" service details
```

EventExport | ee

Exports events from a specified service or specified event trace file to a text file or XML file.

Syntax

{eventexport | ee} *service_name* | [remote]file=
event_trace_file_name export_format export_file_name |
[remote]file=*event_trace_file_name*|
[[[{service | srvc}] |
[{session | sess}] |
[{statement | stmt}] |
[{transaction | trans}] |
[{network | net}] |
[{error | err}] |
[{other | oth}]]]
[count=[{ + | - }] {all | *number*}]
[offset={begin | end} [{ + | - }]*number*]
[query='*custom_event_filter_string*']

where:

service_name is the name of a data access service. The event trace file from the specified service is listed. Service names can be obtained using the ServiceList | sl command.

event_trace_file_name is the path and name of the event trace file.

export_format is the format of the file to which you want to export the events, where:

- {delimited_txt | deltxt} is a text file with commas separating each event you export
- {raw_xml | rxml} is a well-formed XML file.
- {validated_xml | vxml} is a valid XML file.

export_file_name is the path and name of the file to which you are exporting the events.

number is the event number of events to export when used with the count option. When used with the offset option, number is the number of the event from which to start exporting events. For example, if offset=10, SequeLink would export all events starting with event 10. Another example, if you specify count=20 and offset=begin, SequeLink will export the first 20 events. If you specify count=20 and offset=5, SequeLink will export 20 events starting from event 5.

custom_event_filter_string is an event filter statement. See "Filtering Events" on page 595.

Options

Event Types: You can specify one or more of the following event groups for which to export event information: Service, Session, Statement, Transaction, Error, or Other. When you specify one or more of the event groups, SequeLink Manager exports all the events of the type you specified. For example, if you specify Service, SequeLink exports all service events such as Service Started and Service Stopping (these events start with "Service"). Event names that do not start with Service, Session, Network, Error, Statement, or Transaction are Other events (for example, Cursor Closed).

Examples

Local host or remote configuration examples:

Example A: The following example exports event information in the local file named SLoracle10.trc to a text file named export.txt:

```
ee "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" deltxt export.txt
```

Example B: The following example exports event information for only service events in the remote file named SLoracle10.trc to a valid XML file named export.xml:

```
ee "remotefile=C:\Program Files\DataDirect\slserver60\
tracing\SLOracle10.trc" vxml export.xml service
```

Example C: The following example exports event information for all service events starting with event number 10 in the event trace file associated with the SLOracle data access service to a well-formed XML file named export.xml:

```
ee SLOracle rxml export.xml service count=all offset=10
```

Example D: The following example exports session event information for the SLAgent service to a text file named export.txt:

```
ee SLAgent deltxt export.txt
```

Example E: The following example defines a query for the information it will export for the SLOracle10 data access service to a valid XML file named export.xml. The query returns all SQL statements that do not return a return code of 0:

```
ee SLOracle10 vxml export.xml stmt
query='${ReturnCode} != 0'
```

Offline configuration examples:

Example A: The following example exports event information from the local file named SLOracle10.trc to a text file named export.txt:

```
ee "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" deltxt export.txt
```

Example B: The following example exports event information for the first 10 events in the local file named SLoracle10.trc to a text file named export.txt:

```
ee "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" deltxt export.txt count=10
```

Example C: The following example exports event information from the local file named SLoracle10.trc, starting from the end of the file and listing all events, to a well-formed XML file named export.xml:

```
ee "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" rxml export.xml count=-all offset=end
```

Example D: The following example exports only session and service event information from the local file named SLoracle10.trc, starting with the fifth event, to a well-formed XML file named export.xml:

```
ee "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" rxml export.xml offset=5 service session
```

Example E: The following example exports event information for service events only from the local file named SLoracle10.trc to a valid XML file named export.xml:

```
ee "file=C:\Program Files\DataDirect\slserver60\tracing\
SLOracle10.trc" vxml export.xml service
```

Exit | e

Quits the SequeLink Manager Command-Line Tool.

Syntax {Exit | e}

Example

```
exit
```

Help | ?

Displays help about the syntax for invoking the command-line tool or displays help about individual commands.

Syntax

{`Help | ?`} [{*`short_command_name`* | *`long_command_name`*}]

where:

`short_command_name` is the short version of the command name. For example, e is the short version of the Echo command.

`long_command_name` is the long version of the command name, for example, Echo.

Example

The following example displays help for the Echo command:

```
? Echo
```

MVSDB2InterfaceAttributeAdd | mdiaa

Adds an attribute to a DB2 interface.

Syntax

{MVSDB2InterfaceAttributeAdd | mdiaa} *DB2_interface_ID attribute_name value*

where:

DB2_interface_ID identifies the DB2 interface.

attribute_name is the attribute you want to add.

value is the value of the attribute.

Example

The following example adds the MVSDB2RootDescription attribute to DB2_Interface. The value is set to DB2v9.1.

```
mdiaa DB2_interface MVSDB2RootDescription DB2v9.1
```

MVSDB2InterfaceAttributeDelete | mdiad

Deletes an attribute for a DB2 interface.

Syntax

{MVSDB2InterfaceAttributeDelete | mdiad} *DB2_interface_ID attribute_name*

where:

DB2_interface_ID identifies the DB2 interface from which you want to delete an attribute.

attribute_name is the name of the attribute you want to delete.

Example

```
mdiad DB2_Interface MVSDB2RootDescription
```

MVSDB2InterfaceAttributeReplace | mdiar

Changes the value of a DB2 interface attribute.

Syntax

{MVSDB2InterfaceAttributeReplace | mdiar} *DB2_interface_ID attribute_name value*

where:

DB2_interface_ID identifies the DB2 interface.

attribute_name is the name of the attribute you want to change.

value is the new value of the attribute.

Example

The following example changes the value of the MVSDB2TraceTableSize attribute to 256:

```
mdiar DB2_Interface MVSDB2TraceTableSize 256
```

MVSDB2InterfaceInfo | mdii

Lists all attributes and their values for a DB2 interface.

Syntax

{MVSDB2InterfaceInfo | mdii} *DB2_interface_ID*

where *DB2_interface_ID* identifies the DB2 interface for which you want to list attributes and their values.

Example

```
mdii DB2_Interface
```

MVSDB2InterfaceList | mdil

Lists all available external DB2 interfaces.

Syntax

{MVSDB2InterfaceList | mdil}

Example

```
mdil
```

MVSGlobalAttributeDelete | mgad

Deletes an attribute in the z/OS global settings.

Syntax

{MVSGlobalAttributeDelete | mgad} *attribute_name*

where *attribute_name* is the name of the attribute you want to delete.

Example

```
mgad MVSGlobalCompTrace
```

MVSGlobalAttributeReplace | mgar

Changes the value of the specified attribute in the z/OS global settings.

Syntax

{MVSGlobalAttributeReplace | mgar} *attribute_name value*

where:

attribute_name is the name of the attribute you want to change.

value is the new value of the attribute.

Example

```
mgar MVSGlobalCompTrace D
```

MVSGlobalInfo | mgi

Lists all attributes and their values for the z/OS global settings.

Syntax

{MVSGlobalInfo | mgi}

Example

```
mgi
```

MVSUserIDMapAttributeAdd | muimaa

Adds an attribute to a user ID map.

Syntax

{MVSUserIDMapAttributeAdd | muimaa} *user_ID_map attribute_name value*

where:

user_ID_map is the user ID map to which you want to add an attribute.

attribute_name is the name of the attribute you want to add.

value is the value of the attribute.

Example

The following example adds the MVSUIDDefaultAccess attribute to the UserID1 user ID map. The value is set to PERMIT.

```
miumaa UserID1 MVSUIDDefaultAccess PERMIT
```

MVSUserIDMapAttributeDelete | muimad

Deletes an attribute for a user ID map.

Syntax

{MVSUserIDMapAttributeDelete | muimad} *user_ID_map attribute_name*

where:

user_ID_map is the user ID map from which you want to delete an attribute.

attribute_name is the name of the attribute you want to delete.

Example

```
miumad UserID1 MVSUIDDefaultAccess
```

MVSUserIDMapAttributeReplace | muimar

Changes the value for a user ID map attribute.

Syntax

{MVSUserIDMapAttributeReplace | muimar} *user_ID_map attribute_name value*

where:

user_ID_map identifies the user ID map.

attribute_name is the name of the attribute you want to change.

value is the value of the attribute.

Example

```
miumar UserID1 MVSUIDDefaultAccess DENY
```

MVSUserIDMapInfo | muimi

Lists all attributes and their values for a user ID map.

Syntax

{MVSUserIDMapInfo | muimi} *user_ID_map*

where *user_ID_map* is the user ID map for which you want to list attributes and their values.

Example

```
muimi UserID1
```

MVSUserIDMapList | muiml

Lists all available user ID maps.

Syntax {MVSUserIDMapList | muiml}

Example `miuml`

NoOperation | noop

Performs no operation.

Syntax {`NoOperation | noop`}

Example `noop`

ProfileEventTraceCreate | petc

Creates an event trace profile for a SequeLink service.

Syntax {`ProfileEventTraceCreate | petc`} *`service_name`*

where *`service_name`* is the service for which you want to create the event trace profile.

Example `petc SLOracle10`

ProfileEventTraceDelete | petd

Deletes an event trace profile for a SequeLink service.

Syntax

{`ProfileEventTraceDelete | petd`} *`service_name`*

where *`service_name`* is the service from which you want to delete an event trace profile. Service names can be obtained using the ServiceList | sl command.

Example

```
petd SLOracle10
```

ProfileEventTraceEdit | pete

Changes an event trace profile for a SequeLink service.

Syntax

{`ProfileEventTraceEdit | pete`} *`service_name`* {enable | disable | *`event_group`*} [{enable | disable} | {*`event_name`* | *`event_id`*} [{on | off}]] [*`custom_event_filter`*]

where:

`service_name` is the service for which to change an event trace profile. This is a requirement value. Service names can be obtained using the ServiceList | sl command.

`event_group` is the type of event within the service to change. Valid values are:

- {Service | srvc}
- {Session | sess}
- {Statement | stmt}
- {Transaction | trans}
- {Network | net}
- {Error | err}
- {Other | oth}

event_name is the name of the event within in the specified event group to change. You can use the ProfileEventTraceInfo command to list available event names.

event_id is the numeric identifier of the event within the specified event group to change. You can specify either *event_name* or *event_id*, not both. You can use the ProfileEventTraceInfo command to list available event identifiers.

custom_event_filter is an optional event filter that can be specified for all events being enabled. See "Filtering Events" on page 595 for more information.

Examples

Example A: The following example disables all events for the SLOracle10 data access service, meaning that no events are written to the event trace file:

```
pete SLOracle10 disable
```

Example B: The following example enables the event specified by the custom event filter:

```
pete SLOracle10 Session enable '${ClientInfo}="127.0.0.1"'
```

Example C: The following example enables all service events for the `SLOracle10` data access service:

```
pete SLOracle10 service enable
```

Example D: The following example disables all error events for the SLOracle10 data access service:

```
pete SLOracle10 err disable
```

Example E: The following example changes the state of the transaction event named Transaction Rollback to off:

```
pete SLOracle10 trans "Transaction Rollback" off
```

Example F: The following example changes the state of the session event identified by 2 to "on" *if* the event meets the custom event query filter:

```
pete SLOracle10 sess 2 on '${DbmsUser} = "scott"'
```

ProfileEventTraceInfo | peti

Lists event trace profile information for a SequeLink service.

Syntax

{`ProfileEventTraceInfo | peti`} *`service_name`*

where *`service_name`* is the service for which to list event trace profile information. Service names can be obtained using the ServiceList | sl command.

Example

```
peti SLOracle10
```

ProfileMonitorCreate | pmc

Creates a monitoring profile for a SequeLink service.

Syntax

{`ProfileMonitorCreate | pmc`} *`service_name`* {yes | no}

where *`service_name`* is the service for which you want to create the monitoring profile.

You must specify either yes or no. Yes indicates that shared counters will be enabled. No indicates that shared counters will not be enabled. Enabling shared counters is for Windows only.

Example

```
pmc SLOracle10
```

ProfileMonitorDelete | pmd

Deletes a monitoring profile for a SequeLink service.

Syntax

{`ProfileMonitorDelete | pmd`} *`service_name`*

where *`service_name`* is the service from which you want to delete the monitoring profile. Service names can be obtained using the ServiceList | sl command.

Example

```
pmd SLOracle10
```

ProfileMonitorEdit | pme

Changes a monitoring profile for a SequeLink service.

Syntax

{`ProfileMonitorEdit | pme`} *`service_name`* {enable | disable | *`event_group`*} [{enable | disable} | {*event_name* | *event_id*} [{on | off}]]

where:

`service_name` is the service for which to change the monitoring profile. Service names can be obtained using the ServiceList | sl command.

`event_group` is the type of event to change for the monitoring profile. Valid values are:

- {Service | srvc}
- {Session | sess}
- {Statement | stmt}

event_name is the name of the event within in the specified event group to change. You can use the ProfileMonitorInfo command to list available event names.

event_id is the numeric identifier of the event within the specified event group to change. You can specify either *event_name* or *event_id*, not both. You can use the ProfileMonitorInfo command to list available event identifiers.

Examples

Example A: The following example enables all events for the SLOracle10 data access service, meaning that all events are monitored:

```
pme SLOracle10 enable
```

Example B: The following example disables all events for the SLOracle10 data access service, meaning no events are monitored:

```
pme SLOracle10 disable
```

Example C: The following example enables all service events for the SLOracle10 data access service:

```
pme SLOracle10 service enable
```

Example D: The following example disables all session events for the SLOracle10 data access service:

```
pme SLOracle10 session disable
```

Example E: The following example changes the state of the session event named database user to off:

```
pme SLOracle10 session "database user" off
```

Example F: The following example changes the state of the session event identified by 10 to off:

```
pme SLOracle10 session 10 off
```

Example G: The following example changes the state of the statement event named sql to on:

```
pme SLOracle10 stmt sql on
```

ProfileMonitorInfo | pmi

Lists profile monitoring information for a SequeLink service.

Syntax

{`ProfileMonitorInfo | pmi`} *`service_name`*

where *`service_name`* is the service for which you want to list profile monitoring information. Service names can be obtained using the ServiceList | sl command.

Example

```
pmi SLOracle10
```

Quit | q

Quits the SequeLink Manager Command-Line Tool.

Syntax

{`Quit | q`}

Example

```
q
```

SaveConfig | save

Saves the current configuration. This command is available only when AutoSave setting is set to off. For more information about setting AutoSave, see the Set | s command.

Syntax

{`SaveConfig | save`}

Example

```
save
```

ServiceActiveDebugLogLevel | sadll

Displays or changes the debug log level of an active SequeLink service. When no debug log level values are provided, the current settings are listed.

Syntax

```
{ServiceActiveDebugLogLevel | sadll} service_name
[ [{dis | disable}]
[{enall |  enableall}]
[{en | enable}]
[{ferr | fatalerror}={off | on}] |
[{err | errors=off | on}] |
[{war | warnings}={off | on}] |
[{info | informationals}={off | on}] |
[{debug | debugging}={off | on}] |
[{sspp | ssppackets}={off | on}] |
[{sspr | ssprequests}={off | on}] | ... ]
```

where *service_name* is the active service for which you want to display or change debug log levels. Service names can be obtained using the `ServiceList | sl` command.

Example

The following example turns on debug messages and turns off error messages in the debug log for the SLOracle10 data access service:

```
sadll SLOracle10 debug=on err=off
```

ServiceActiveInfo | sai

Lists specific information about an active SequeLink service.

Syntax

{`ServiceActiveInfo | sai`} *`service_name`*

where *`service_name`* is the active service for which you want to list information. Service names can be obtained using the ServiceList | sl command.

Example

```
sai SLOracle10
```

ServiceAttributeAdd | saa

Adds an attribute to a SequeLink service.

Syntax

{`ServiceAttributeAdd | saa`} *`service_name attribute_name value`*

where:

`service_name` is the service to which you want to add an attribute. Service names can be obtained using the ServiceList | sl command.

`attribute_name` is the name of the attribute you want to add.

`value` is the value of the attribute.

See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes and their valid values.

Example

```
saa SLOracle10 ServiceUser sqlnk
```

ServiceAttributeDelete | sad

Deletes an attribute from a SequeLink service.

Syntax

{`ServiceAttributeDelete | sad`} *`service_name attribute_name`*

where:

`service_name` is the service from which you want to delete an attribute. Service names can be obtained using the ServiceList | sl command.

`attribute_name` is the name of the attribute you want to delete.

See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes and their valid values.

Examples

```
sad SLOracle10 ServiceCodePage
```

```
sad SLOracle10 ServiceUser[2]
```

ServiceAttributeReplace | sar

Changes the value for a SequeLink service attribute.

Syntax

{ServiceAttributeReplace | sar} *`service_name attribute_name`* *value*

where:

`service_name` is the service for which you want to change the value of an attribute. Service names can be obtained using the ServiceList | sl command.

`attribute_name` is the name of the attribute you want to change.

`value` is the new value of the attribute.

See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes and their valid values.

Example

```
sar SLOracle10 ServiceUser[2] devuser
```

ServiceCreate | sc

NOTE: This command is not applicable to z/OS.

Creates a SequeLink service based on a service template ID.

Syntax

{`ServiceCreate | sc`} *`service_name service_ID tcp_port`*

where:

`service_name` is the name of the service to create.

`service_ID` identifies the ID of the service template to use to create the new service. To get a listing of available templates and their IDs, use the ServiceTemplateList | stl command.

`tcp_port` is the TCP/IP port on which the service is listening.

Example

```
sc SLOracle10 SL6_Oracle10 19996
```

ServiceDebugLogLevel | sdll

Displays or changes the debug log level of a SequeLink service. When no debug log level values are provided, the current settings are listed.

Syntax

```
{ServiceDebugLogLevel | sdll} service_name
[ [{dis | disable}]
[{enall |  enableall}]
[{en | enable}]
[{ferr | fatalerror}={off | on}] |
[{err | errors=off | on}] |
[{war | warnings}={off | on}] |
[{info | informationals}={off | on}] |
[{debug | debugging}={off | on}] |
[{sspp | ssppackets}={off | on}] |
[{sspr | ssprequests}={off | on}] | ... ]
```

where *service_name* is the service to which the session belongs. Service names can be obtained using the `ServiceList | sl` command.

Example

The following example turns on debug messages and turns off error messages in the debug log file for the SLOracle10 data access service:

```
sdll SLOracle10 debug=on err=off
```

ServiceDelete | sd

NOTE: This command is not applicable to z/OS.

Deletes a SequeLink service.

Syntax

{`ServiceDelete | sd`} *`service_name`*

where *`service_name`* is the name of the service you want to delete. Service names can be obtained using the ServiceList | sl command.

Example

```
sd SLOracle10
```

ServiceInfo | si

Lists all attributes and their values for a SequeLink service.

Syntax

{`ServiceInfo | si`} *`service_name`*

where *`service_name`* is the service for which you want to list attributes and their values. Service names can be obtained using the ServiceList | sl command.

See Appendix D "SequeLink® Service Attributes" on page 491 for a list of SequeLink service attributes and their valid values.

Example

```
si SLOracle10
```

ServiceList | sl

Lists all available SequeLink services.

Syntax

{`ServiceList | sl`}

Example

```
sl
```

ServiceRegister | sr

NOTE: This command is not applicable to z/OS.

Registers a SequeLink service.

Syntax

{`ServiceRegister | sr`} *`service_name`*

where *`service_name`* is the service you want to register. Service names can be obtained using the ServiceList | sl command.

Example

```
sr SLOracle10
```

ServiceStart | ss

NOTE: This command is not applicable to z/OS.

Starts a SequeLink service.

Syntax

{`ServiceStart | ss`} *`service_name`*

where *`service_name`* is the service you want to start. Service names can be obtained using the ServiceList | sl command.

Example

```
ss SLOracle10
```

ServiceStop | sst

NOTE: This command is not applicable to z/OS.

Stops a SequeLink service.

Syntax

{`ServiceStop | sst`} *`service_name`*

where *`service_name`* is the service you want to stop. Service names can be obtained using the ServiceList | sl command.

Example

```
sst SLOracle10
```

ServiceTemplateList | stl

Lists all available SequeLink service templates and their IDs.

Syntax

{`ServiceTemplateList | stl`}

Example

```
stl
```

ServiceTemplateInfo | sti

Lists all attributes and their values for a SequeLink service template.

Syntax

{`ServiceTemplateInfo | sti`} *`service_template_ID`*

where *`service_template_ID`* identifies the service template. To get a listing of available templates and their IDs, use the ServiceTemplateList | stl command.

Example

```
sti SL6_Oracle10
```

ServiceUnregister | su

NOTE: This command is not applicable to z/OS.

Unregisters the specified SequeLink service.

Syntax

{`ServiceUnregister | su`} *`service_name`*

where *`service_name`* is the service you want to unregister. Service names can be obtained using the ServiceList | sl command.

Example

```
su SLOracle10
```

SessionDebugLogLevel | sesdll

NOTE: This command is not applicable on Linux/UNIX or z/OS.

Displays or changes the debug log level of the specified session. When no debug log level values are provided, the current settings are listed.

Syntax

```
{SessionDebugLogLevel | sesdll} service_name
[ [{dis | disable}]
[{enall | enableall}]
[{en | enable}]
[{ferr | fatalerror}={off | on}] |
[{err | errors}={off | on}] |
[{war | warnings}={off | on}] |
[{info | informationals}={off | on}] |
[{debug | debugging}={off | on}] |
[{sspp | ssppackets}={off | on}] |
[{sspr | ssprequests}={off | on}] | ... ]
```

where *`service_name`* is the service to which the session belongs. Service names can be obtained using the `ServiceList | sl` command.

Examples

Example A: The following example displays the current settings of the debug log for session 5 of the SLOracle10 data access service:

```
sesdll SLOracle10 5
```

Example B: The following example turns on debug messages and turns off error messages in the debug log for session 5 of the SLOracle10 data access service:

```
sesdll SLOracle10 5 debug=on err=off
```

SessionInfo | sesi

Lists specific information about a session and its associated statements (only for data access sessions) for SequeLink service.

Syntax

{`SessionInfo | sesi`} *`service_name session_ID`*

where:

`service_name` is the service to which the session belongs. Service names can be obtained using the ServiceList | sl command.

`session_ID` identifies the session for which you want to display information. This session must belong to the specified service. Session IDs can be obtained using the ServiceActiveInfo | sai command.

Example

```
sesi SLOracle10 5
```

SessionStop | sess

Stops a session for a SequeLink service.

Syntax

{`SessionStop | sess`} *`service_name session_ID`*

where:

`service_name` is the service to which the session belongs. Service names can be obtained using the ServiceList | sl command.

`session_ID` identifies the session you want to stop. This session must belong to the specified service. Session IDs can be obtained using the ServiceActiveInfo | sai command.

Example

```
sess SLOracle10 5
```

Set | s

Sets the following configuration for the command-line tool settings. Valid values are:

- AutoSave
- Echo
- IgnoreErrors

AutoSave specifies whether to save the configuration when a change is made. When set to on, the configuration is saved when a change is made. When set to off, the configuration is not automatically saved, and you must issue a SaveConfig command to save the configuration.

Echo specifies whether all commands entered on the command line are printed to output. When set to on, all commands entered on the command line are printed to output. When set to off, this setting is ignored. The default is off.

IgnoreErrors specifies whether the SequeLink Manager Command-Line Tool stops when an error occurs. When set to on, the SequeLink Manager Command-Line Tool stops when an error occurs. This setting is useful when you want the processing of batch files to stop when an error occurs. When set to off, this setting is ignored. The default is off.

Syntax

{`Set | s`} [*option1*={on | off} [*option2*={on | off} ...]]

Examples

```
set autosave=off
```

```
set echo=on
```

```
set echo=on ignoreerrors=on
```

C Operator Interface Commands for z/OS

This appendix describes the Operator Interface commands by category:

See "Using the SequeLink® Manager for z/OS Operator Interface" on page 120 for information about using the Operator Interface.

Syntax of Operator Interface Commands

The syntax for Operator Interface commands is:

compid command parms

where:

compid is the ID of the server component to which the command is directed. Valid component IDs are:

- CNTL - (Server Controller Tasks)
- LOGR - (Messaging Component)
- DB2 - (DB2 Component)
- RRS - (RRS Component)
- THPL - (Threadpool Component)
- XA - (XA Component)

command is the command name.

parms are the parameters of the command.

For example:

`CNTL HALT` shuts down the Server immediately.

NOTE: When using the Operator Interface, you can repeat the last command you entered by pressing F24 or Shift+F12.

Server Controller Task (CNTL) Commands

CLOSE

Shuts down the SequeLink Server system in Normal or Immediate mode.

Parameters

TYPE=NORM

Shuts down the SequeLink Server. All connections are allowed to end normally. New connections are refused. If you start shutdown normally, you can still override it to IMMEDIATE mode.

TYPE=IMMED

Shuts down the SequeLink Server immediately. All client connections are terminated in a consistent manner.

Example

`CNTL CLOSE TYPE=NORM` shuts down the SequeLink Server. All connections end normally and new connections are refused. You can still override shutdown to IMMEDIATE mode.

`CNTL CLOSE TYPE=IMMED` shuts down the SequeLink Server immediately. All client connections are terminated in a consistent manner.

HALT

Shuts down the SequeLink Server system.

Parameters

None

Example

`CNTL HALT` shuts down the SequeLink Server immediately.

STATUS

Displays the general status of the SequeLink Server system.

Parameters

SHOW=ALL

Lists all active tasks (server core tasks and service tasks) and shows their status. All services known to the server are listed.

SHOW=SERVICES

Lists all active application service tasks known to the server.

NOTE: The SHOW parameter is optional.

Example

`CNTL STATUS` displays the status of each attached component. The output displays on the operator's console.

`CNTL STATUS SHOW=ALL` displays all known active tasks and services. The output displays on the operator's console.

`CNTL STATUS SHOW=SERVICES` lists all active services. The output displays on the operator's console.

LIST

Lists free storage, expressed in kilobytes (K), above and below the 16 MB line.

Parameters

Type=Free

Example

`CNTL LIST TYPE=FREE` lists the free storage above and below the 16 MB line.

Logging Controller (LOGR) Commands

ALTER

Starts or stops logging of VAILOG trace messages.

Parameters

TRACE=ON | OFF

Turns on or off logging of trace messages sent by components.

COMPNT=ALL | *component_ID*

Specifies which components will have messages logged. This can be set to ALL (messages from all system components) or to a component ID (messages for that component only). Valid component IDs are:

C Common or shared components (operator interface, for example)

D DB2 or RRS component

L Generic log component

S Server controller component

T Threadpool component

Examples

`LOGR ALTER TRACE=ON COMPNT=ALL` sets tracing on for all components.

`LOGR ALTER TRACE=OFF COMPNT=T` sets tracing off for the threadpool component.

`LOGR ALTER TRACE=OFF COMPNT=ALL` sets tracing off for all components.

PRINT

Prints a SequeLink Server system log.

Parameters

LOG=*log_name*

Prints the primary log, VAILOGP, or the secondary log, VAILOGS. The default is the inactive log.

CLASS=*class_name*

Indicates which SYSOUT class to which to spool the printout. The default is A.

HOLD=Y | N

Indicates whether the output is to be held on the output queue. Valid values are Y (Yes) or N (No). The default is N.

NOTE: All parameters are optional.

Example

`LOGR PRINT CLASS=L` spools the inactive log to the JES2 output class L; the output is not held in the output queue.

STATUS

Displays the status of the message logging component (LOGR).

Parameters

None

Example

`LOGR STATUS` returns a single line, providing the general status of the logging component.

SWITCH

Changes the active log. This command switches to the alternate log, making the current log available for printing or archiving.

Parameters None

Example `LOGR SWITCH` changes the active log. For example, if the primary log (VAILOGP) is active when this command is issued, the secondary log (VAILOGS) becomes the active log. Issuing the command again makes VAILOGP the active log again.

DB2 Component (DB2) Commands

SHOW

Show the active DB2 interface.

Parameters None

Example `DB2 SHOW` shows the active DB2 interface.

STATUS

Displays the status of SequeLink Server DB2 interface.

Parameters None

Example `DB2 STATUS` displays the status of the DB2 interface.

RRS Component (RRS) Commands

LIST

Lists all RRS contexts associated with an RRSAF DB2 thread in use by the SequeLink Server.

Parameters None

Example `RRS LIST` lists all RRS contexts, their status, and their age.

RELEASE

Releases RRS contexts associated with *reusable* DB2 threads that are not used for a specified period.

Parameters AGE=*dddd*

Specifies the period, in seconds. When this value is omitted, a value of 300 seconds is used.

Example `RRS RELEASE AGE=60` releases the reusable DB2 threads not used within that last minute, but keeps at least the number of DB2 threads available as specified in ServiceDB2MinThreads in the SequeLink configuration file.

`RRS RELEASE AGE=0` releases all reusable DB2 threads even if their reusable status is current, or the value is specified in the SequeLink configuration file for ServiceDB2MinThreads.

THREADPOOL Component (THPL) Commands

DISPLAY

Shows a summary of the threadpool settings.

Parameters SHOW=STAT

Shows a summary of the threadpool settings and a detailed overview of the threadpool statistics.

Example `THPL,DISPLAY` displays a summary of the Service settings and the DataSource settings for each data source, as shown:

```
ServiceMinThreads(xxx),ServiceMaxThreads(yyy),
ServiceThreadLockThreshold
ServiceDB2MaxThreads(uuu),ServiceDB2MinThreads(ttt)

DataSourceThreadMaxRpc(nnn),
DataSourceThreadRpcTimeOut(nnn)
```

`THPL,DISPLAY,SHOW=STAT` displays an overview of the threadpool statistics, as shown:

```
Threads    NrOfRPC              NrRPC            CPU            %CPU
=======    =======              ========         ===            =====
TWORK003   number of RPC's      % of total RPC   CPU seconds    %of total CPU
TWORK004   handled by this
           workerthread
...

TWORKnnn   ...
```

LIST

Lists the status of all sessions. The maximum number of lines of output for a single invocation of this operator command is 100. The ADDRESS and USERID arguments will further limit the requested output.

Parameters

ADDRESS=*client Ipaddress*

Identifies the client IP address.

USERID=*userid*

Identifies a userid.

Both the optional *client Ipaddress* and the *userid* parameters can be specified generically.

Example

`THPL LIST` lists the status of all sessions, up to a maximum of 100 sessions.

`THPL LIST USERID=MFI*` lists the status of the sessions with a userid prefixed with MFI.

`THPL LIST ADDRESS=10.131.*` lists the status of all sessions for which the client IP address begins with 10.131.

`THPL LIST USERID=MFI*,ADDRESS=10.131.40.59` lists the status of all sessions for userids prefixed with MFI and that have a client IP address of 10.131.40.59.

KILL

Kills a session.

Parameters ID=*session_ID*

Identifies the session that must be stopped. The SessionId can be obtained from the LIST command.

NOTE: This parameter is required.

Example `THPL KILL ID=2` stops session 2.

XA Component (XA) Commands

LIST

Lists all XA transactions and associated RRS contexts in the SequeLink Server.

Parameter None

Example `LIST` lists all XA transactions and associated RRS contexts in the SequeLink Server.

RELEASE

Performs cleanup of XA transactions and associated RRS contexts in the SequeLink Server by releasing contexts without an owner task that may still hold locks on DB2 resources. The transactions can be rolled back or committed to release the locks in the data store held by these contexts.

Parameters

THRDID=*thread_ID*

Identifies the thread ID that holds the resources in the data store.

NOTE: This parameter is required.

TYPE=ROLLBACK | COMMIT

Rolls back or commits the specified transaction. The default is ROLLBACK.

Example

`RELEASE THRDID=T0000003 TYPE=ROLLBACK` rolls back the specified transaction. The RSS context used the last time by T0000003 will end and all resources held in the data store by the context will be rolled back.

D SequeLink® Service Attributes

This appendix describes the SequeLink service attributes you can set to configure and manage your SequeLink services using the SequeLink Manager tools. For each attribute, the following information is listed:

- Valid values for the attribute
- A description of the attribute
- A default value (when an attribute has a default)
- Whether the attribute type is static or dynamic

NOTES:

- Unless noted otherwise, the service attributes apply to all operating systems. Information specific to an operating system is identified as such.
- Attributes beginning with the string "DataSource" are SequeLink data access service attributes that are associated with a server data source.
- *Static attributes* require you to restart a SequeLink service when you add or change the attribute before the change becomes effective. *Dynamic attributes* become effective after the attribute is added or changed and the configuration is saved. Most dynamic attributes affect the behavior of a database connection; therefore, when you add or change an attribute, the new values are used for the next connection, active connections do not use the new values.

  NOTE: Server data source attributes are always dynamic.

- The default value is the value used when it is not overridden during the installation of the SequeLink Server. Some default values are changed when the service is created during the installation process. For example, this appendix states that the default value of the ServiceEventTraceLocation attribute is empty string; however, after the installation of the SequeLink Server, the default value is a directory beneath the installation directory.
- When working with the SequeLink Manager for z/OS, you can access online help for each field by positioning your cursor on the field and pressing F1. The online help will indicate whether the field maps to a SequeLink service attribute.
- Some service attributes are intended for internal use or debugging use only. *Do not* modify these attributes unless you are instructed to do so by DataDirect Technologies technical support.

Summary of Service Attribute Use

The following tables list the service attributes and identify the SequeLink Servers to which they apply:

- Table D-1 "Attributes Used by All SequeLink Servers" on page 493
- Table D-2 "Attributes Used by All SequeLink Servers for Databases Without ROWID Support" on page 496
- Table D-3 "Attributes Used by SequeLink Server for DB2 UDB on Linux, UNIX, and Windows" on page 497
- Table D-4 "Attributes Used by SequeLink Server for DB2 UDB on z/OS" on page 497
- Table D-5 "Attributes Used by SequeLink Server for Informix" on page 499

Table D-1 lists the service attributes that are used by all SequeLink Servers. The attribute name provides a link to the description of the attribute.

Table D-1. Attributes Used by All SequeLink Servers

Table D-1. Attributes Used by All SequeLink Servers *(cont.)*

Table D-1. Attributes Used by All SequeLink Servers *(cont.)*

Table D-1. Attributes Used by All SequeLink Servers *(cont.)*

- "ServiceTCP1stRecvTimeLimit" on page 585
- "ServiceThreadLockThreshold" on page 586
- "ServiceEvQShmMonitorSize" on page 575
- "ServiceEvQShmQMaxResend" on page 575
- "ServiceUnixSyslogFacility" on page 586
- "ServiceUser" on page 587

Table D-2 lists the service attributes that are used by SequeLink Servers for databases that do not have native ROWID support. The attribute name provides a link to the description of the attribute. ROWID is not used for Oracle and Informix.

Table D-2. Attributes Used by All SequeLink Servers for Databases Without ROWID Support

- "DataSourceKeySetName" on page 518
- "DataSourceKeySetDataType" on page 517
- "DataSourceKeySetOptions" on page 518
- "DataSourceKeySetSize" on page 519

In addition to the service attributes listed in Table D-1, SequeLink Server for DB2 UDB on Linux, UNIX, and Windows uses the service attributes listed in Table D-3.

Table D-3. Attributes Used by SequeLink Server for DB2 UDB on Linux, UNIX, and Windows

- "DataSourceDB2ConnectOptions" on page 507
- "DataSourceDB2LogPath" on page 509

In addition to the service attributes listed in Table D-1, SequeLink Server for DB2 UDB on z/OS uses the service attributes listed in Table D-4.

Table D-4. Attributes Used by SequeLink Server for DB2 UDB on z/OS

- "DataSourceDB2DBFilterList" on page 508
- "DataSourceDB2CatalogOwner" on page 506
- "DataSourceDB2CollectionPrefix" on page 507
- "DataSourceDB2DecimalPoint" on page 508
- "DataSourceDB2MaxLobSize" on page 509
- "DataSourceDB2ReportLobsFirst" on page 509
- "DataSourceDB2TranslateCount" on page 510
- "MVSDataSourceDB2Plan" on page 547
- "MVSDataSourceUIDMap" on page 548
- "MVSDB2ExitLibrary" on page 548
- "MVSDB2LoadLibrary" on page 548
- "MVSDB2RootDescription" on page 549
- "MVSDB2SubsystemName" on page 549
- "MVSDB2Version" on page 549
- "MVSGlobalClustername" on page 549
- "MVSGlobalCompTrace" on page 550

Table D-4. Attributes Used by SequeLink Server for DB2 UDB on z/OS *(cont.)*

In addition to the service attributes listed in Table D-1, SequeLink Server for Informix uses the service attributes listed in Table D-5.

Table D-5. Attributes Used by SequeLink Server for Informix

- "DataSourceINFClientLocale" on page 515
- "DataSourceINFDbAnsiWarn" on page 515
- "DataSourceINFDbLang" on page 515
- "DataSourceINFDbLocale" on page 515
- "DataSourceINFDbNls" on page 516
- "DataSourceINFDelimIdent" on page 516
- "DataSourceINFHost" on page 516
- "DataSourceINFInformixDir" on page 516
- "DataSourceINFInformixServer" on page 516
- "DataSourceINFLang" on page 517
- "DataSourceINFService" on page 517
- "ServiceINFMaxNrActStat" on page 578

In addition to the service attributes listed in Table D-1, SequeLink Server for JDBC Socket uses the service attributes listed in Table D-6.

Table D-6. Attributes Used by SequeLink Server for JDBC Socket

- "DataSourceProviderTypesFile" on page 528
- "DataSourceProviderTypesSection" on page 529
- "DataSourceSOCEmptyRowSyntax" on page 533
- "DataSourceSOCJDBCDbProperties Name" on page 535
- "DataSourceSOCJDBCDriverClassName" on page 535
- "DataSourceSOCJDBCLogPath" on page 536

Table D-6. Attributes Used by SequeLink Server for JDBC Socket *(cont.)*

- "DataSourceSOCJDBCReadLongDataLength" on page 536
- "ServiceDBConnectionPooling" on page 569

In addition to the service attributes listed in Table D-1, SequeLink Server for Microsoft SQL Server uses the service attributes listed in Table D-7.

Table D-7. Attributes Used by SequeLink Server for Microsoft SQL Server

- "DataSourceMSSBindAllLOBs" on page 522
- "DataSourceMSSConnectOptions" on page 522
- "DataSourceMSSCursorType" on page 523
- "DataSourceMSSCursorWarnings" on page 523
- "ServiceMSSMergeXaBranches" on page 581

In addition to the service attributes listed in Table D-1, SequeLink Server for ODBC Socket uses the service attributes listed in Table D-8.

Table D-8. Attributes Used by SequeLink Server for ODBC Socket

- "DataSourceProviderTypesFile" on page 528
- "DataSourceProviderTypesSection" on page 529
- "DataSourceSOCEmptyRowSyntax" on page 533
- "DataSourceSOCODBCLogPath" on page 538
- "ServiceDBConnectionPooling" on page 569

In addition to the service attributes listed in Table D-1, SequeLink Server for Oracle uses the service attributes listed in Table D-9.

Table D-9. Attributes Used by SequeLink Server for Oracle

- "DataSourceORAMapTSWTZ" on page 525
- "DataSourceORANumber0IsNumeric" on page 526
- "DataSourceORAPublicSchemaSupp" on page 527
- "DataSourceORARecyclebinSupp" on page 527
- "DataSourceORAReportDBLinkAsCatalog" on page 527
- "DataSourceORAServiceName" on page 528
- "DataSourceORASynDBLinkObjSupp" on page 528
- "ServiceORASerializeLogon" on page 582

In addition to the service attributes listed in Table D-1, SequeLink Server for Sybase uses the service attributes listed in Table D-10.

Table D-10. Attributes Used by SequeLink Server for Sybase

- "DataSourceSybConnectOptions" on page 538
- "DataSourceSYBLogPath" on page 540
- "DataSourceSybNetworkAddress" on page 540

Defining SequeLink® Service Attributes

This section describes SequeLink service attributes in alphabetical order.

DataSourceAllowBatchStatements

Specifies whether the SequeLink data source accepts a batch of SQL statements. A batch is used to submit multiple update commands together as a single unit to the underlying DBMS.

Valid values are:

- FALSE=The SequeLink data source does not accept batch SQL statements. Semicolons are considered to be part of a single SQL statement.
- TRUE=The SequeLink data source accepts batch SQL statements.

The default is FALSE.

Type=Dynamic

NOTE: This service attribute is used by the ODBC Client and the ADO Client.

DataSourceAllowPrefetch

Specifies whether the JDBC Client launches a request to the server to fetch the next set of rows from the database while the application is processing a previous set of retrieved rows.

Valid values are:

- FALSE=Prefetch of result set data is not allowed.
- TRUE=The JDBC Client launches a request to the server to fetch the next set of rows from the database while the application is processing a previous set of retrieved rows.

The default is FALSE.

Type=Dynamic

NOTE: This service attribute is used by the JDBC Client.

DataSourceApplID

Specifies a list of application IDs the SequeLink data access service will accept. A valid application ID is an alphanumeric string with a maximum length of 128 characters.

The default is an empty string.

Type=Dynamic

See Chapter 13 "Configuring SequeLink® Security" on page 291 for more information about using application IDs to limit access to SequeLink services.

DataSourceArrayFetchMaxBytes

Specifies the number of bytes the SequeLink Server uses when doing an arrayfetch from the DBMS. The larger the size of the buffer, the more rows that SequeLink will fetch in one database fetch. This will have a positive effect on performance. However, a larger buffer increases the amount of memory that SequeLink uses.

A value of 0 disables the array fetch mechanism. Valid values are 0 to 65536.

The default is 65536.

Type=Dynamic

DataSourceAutoApplId

Specifies an application ID that is automatically generated by the ODBC Client to identify the client application to the SequeLink service. A valid automatically generated application ID contains exactly 40 hexadecimal digits. Each provided value can have an optional text description. For example:

```
InventoryControl=aaf7798c8c66e6b3a6b1462eff442b39b7be6946
```

The default is an empty string.

See Chapter 13 "Configuring SequeLink® Security" on page 291 for more information about using application IDs to limit access to SequeLink service.

Type=Dynamic

DataSourceBatchProcessing

Configures the behavior for batch processing on the SequeLink server. Valid values are:

- Emulated=Each element of the batch is processed individually. Processing of the batch stops on the first error encountered.
- Native=The batch interface of the DBMS is used.

  Handling of error conditions is database-specific:

  - DB2 UDB: None of the statements are processed if an error occurs. No update count information is available.

- Oracle: All elements of the batch are processed. Processing does not stop on errors.
- JDBC Socket Server: Behavior depends on the back-end JDBC driver.
- ODBC Socket Server: Behavior depends on the back-end ODBC driver.
- SQL Server: All elements of the batch are processed. Processing does not stop on errors.
- Sybase: None of the elements of the batch are processed if an error occurs. No update count information is available.

The default is Emulated.

For more information about batch processing, refer to your DBMS documentation.

Type=Dynamic

DataSourceBlockFetchForUpdate

When the isolation level is Read Committed and a SELECT FOR UPDATE statement is issued against some data stores, the SequeLink Client does not lock the expected row. Valid values are:

- 0=The appropriate row will be locked. Specifying 0 will degrade the performance for SELECT FOR UPDATE statements because rows will be fetched one at a time.
- 1=The appropriate row is not locked.

The default is 1.

Type=Dynamic

For more information, refer to the *SequeLink Developer's Reference*.

DataSourceCurrentCatalog

Specifies the default catalog to be used when connected to the SequeLink data access service. The valid value is a defined database catalog name.

The default is an empty string.

Type=Dynamic

DataSourceCursorHold

Specifies the effect of transaction completion on open cursors. Valid values are:

- 0=cursor no hold. The cursors are destroyed when the transaction is committed.
- 1=cursor hold. The cursors are not destroyed when the transaction is committed.

For Informix and for DB2 UDB on z/OS, the default is 1.

For all other DBMS, the default is 0.

Type=Dynamic

z/OS DataSourceDB2CatalogOwner

Specifies the owner of the DB2 catalog. This parameter allows you to limit the meta-information that is returned by using views on the DB2 catalog.

Specify SYSIBM when making selections on the native DB2 catalog tables. Otherwise, specify a different value for the catalog owner to limit the number of tables retrieved by SequeLink. The valid value is a defined database catalog name.

The default is SYSIBM.

Type=Dynamic

z/OS DataSourceDB2CollectionPrefix

Specifies a user-defined prefix, which can be no longer than 15 characters, for the DB2 collection that identifies where the SequeLink package resides.

Based on the isolation level (DataSourceTransactionIsolation attribute), a suffix is added to the DB2 collection name, for example, *CollectionPrefix*_U. Valid values are:

- U= Read Uncommitted
- S= Read Committed
- T= Repeatable Read
- R= Serializable

For example, if you use the default prefix SLD600 for an uncommitted connection, the DB2 collection name is SLD600_U.

The default is SLD600.

Type=Dynamic

DataSourceDB2ConnectOptions

Specifies additional connect options to use when the SequeLink server connects to DB2 on Linux, UNIX, or Windows. This string will be appended to the connect string.

The default is an empty string.

Type=Dynamic

DataSourceDB2DBFilterList

Specifies a list of comma-separated values specifying the databases from which tables can be retrieved by a SQLTables (ODBC), getTables (JDBC), or TABLES (ADO) call. You cannot use quotes.

The default is an empty string.

Type=Dynamic

DataSourceDB2DBMSName

Specifies the value that is returned for SQLGetInfo (SQL_DBMS_NAME). Valid values are:

- 0=DB2/MVS. This value supports the behavior of earlier versions of SequeLink.
- 1=DB2. This value is compliant with the behavior of the IBM DB2 Connect driver, and supports applications such as the IBM remote administration and development tools.
- 2=DB2 UDB for OS/390 and z/OS. This value is compliant with the behavior of the DataDirect Connect *for* DB2 driver.

The default is 0.

Type=Dynamic

z/OS DataSourceDB2DecimalPoint

Specifies whether the decimal point for numbers is a comma (,) or a period(.). This setting allows the data source to overwrite the DB2 installation parameter DECIMAL POINT. If no setting is specified, the DB2 installation parameter is read from the load module of DB2.

Valid values are:

- PERIOD=The service uses a period as the decimal point character.
- COMMA=The service uses a comma as the decimal point character.

The default is PERIOD.

Type=Dynamic

DataSourceDB2LogPath

Defines the log path for all CLI calls executed by the SequeLink service. The valid value is a log path.

Type=Dynamic

z/OS DataSourceDB2MaxLobSize

Specifies the size (in bytes) that is reported for BLOB and CLOB data types, for example, in response to a SQLGetTypeInfo call (ODBC) or the equivalent OLE DB and JDBC calls.

Valid values are from 1 to 2147483647.

The default is 2147483647.

Type=Dynamic

z/OS DataSourceDB2ReportLobsFirst

Specifies the order in which BLOB and CLOB data types are reported, for example, in response to a SQLGetTypeInfo call (ODBC) or the equivalent OLE DB and JDBC calls.

- FALSE=Report LONG VARCHAR and LONG VARCHAR FOR BIT DATA data types before CLOB and BLOB data types.
- TRUE=Report CLOB and BLOB data types before LONG VARCHAR and LONG VARCHAR FOR BIT DATA data types.

The default is FALSE.

Type=Dynamic

z/OS DataSourceDB2TranslateCount

On DB2 UDB V7 and earlier, specifies whether the service will substitute COUNT(columnName) with COUNT(*).

- FALSE=The service will not substitute COUNT(columnName) with COUNT(*).
- TRUE=The service will substitute COUNT(columnName) with COUNT(*).

The default is FALSE.

Type=Dynamic

DataSourceDescribeParam

Defers the fetching of parameter metadata until it is needed by the application.

- Enabled=The service always requests parameter description information, even when the application does not require it. This value minimizes network traffic and should be used when parameter information is required for most statements.
- Deferred=The service requests parameter description information from the server when the client application asks

for it. Use this value when the application requests description information for a small number of statements.

- Disabled=Parameter description requests are disabled.

The default is Enabled for all databases except Sybase, where the Service template contains Disabled.

Type=Dynamic

DataSourceDescription

Specifies a general description of the server data source.

Type=Static

DataSourceDisableWarnings

Turns on and off the filtering of generated warnings.

- TRUE=Turns on filtering
- FALSE=Turns off filtering

Type=Dynamic

DataSourceEnableAccess

Determines whether access to a specified data source is enabled.

- TRUE=Any new connection is accepted.
- FALSE=Any new connection is rejected. This setting does not terminate existing connections, which will be able to finish their work.

The default is TRUE.

Type=Dynamic

DataSourceEnableDescribeParam

Turns on a workaround for connections to Oracle data stores only. For more information, refer to the *SequeLink Developer's Reference*.

- 0=The option is disabled. Support is turned off for SQLDescribeParam.
- 1=Support is turned on for SQLDescribeParam and will describe all parameters as SQL_CHAR with a precision of 999.

The default is 0.

Type=Dynamic

DataSourceFetchNextOnly

Turns on a workaround for Visual Basic/Remote Data Objects (RDO) that circumvents a problem with FORWARD_ONLY cursors when the driver reports other values than FETCH_NEXT for SQLGetInfo(SQL_FETCH_DIRECTION).

For example, if the driver only reports FETCH_NEXT, RDO performs SQLExecDirect, SQLBindCol, and SQLExtendedFetch(NEXT). If the driver supports more than FETCH_NEXT, RDO performs SQLExecDirect, SQLExtendedFetch(NEXT), and SQLGetData. This is only valid when the rowsize is 1, but RDO uses a larger rowsize in this situation.

- TRUE=The driver will incorrectly report that only SQL_FETCH_NEXT is supported, which satisfies RDO.
- FALSE=The workaround is not turned on.

The default is FALSE.

Type=Dynamic

For more information, refer to the *SequeLink Developer's Reference*.

DataSourceFetchTimeStampAsString

Specifies whether a workaround for a Microsoft Access problem with timestamps is turned on.

- TRUE=Yes
- FALSE=No

The default is FALSE.

Type=Dynamic

DataSourceFixCharTrim

Turns on a workaround for applications that have a problem using SQL_CHAR data padded with spaces. The ODBC driver returns SQL_CHAR data padded with spaces as mandated by the ODBC specification. For more information, refer to the *SequeLink Developer's Reference*.

- 0=The workaround is turned off.
- 1=Returns SQL_CHAR data that is not padded with spaces.

The default is 0.

Type=Dynamic

DataSourceGetOutputParams

Turns on a workaround that allows you to control when output parameters of stored procedures are returned to calling applications. For more information, refer to the *SequeLink Developer's Reference*.

The value for this service attribute should be set to the cumulative value of all chosen options added together.

Valid values are:

- 0=The workaround is not turned on.
- 1=Output parameters are returned after an execute.
- 2=Output parameters are returned after a fetch is complete.
- 4=Output parameters are returned after moreresults no more rows.
- 7=Output parameters are returned after all of the above.

The default is 7.

Type=Dynamic

NOTE: Set GetOutputParams=3 when executing stored procedures with output parameters in RDO (Visual Basic 5 and 6).

DataSourceINFClientLocale

Sets the Informix CLIENT_LOCALE environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFDbAnsiWarn

Sets the Informix DBANSIWARN environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFDbLang

Sets the Informix DBLANG environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFDbLocale

Sets the Informix DBLOCALE environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFDbNls

Sets the Informix DBNLS environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFDelimIdent

Sets the Informix DELIMIDENT environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFHost

Sets the Informix INFORMIXHOST environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFInformixDir

Sets the Informix INFORMIXDIR environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFInformixServer

Sets the Informix INFORMIXSERVER environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFLang

Sets the Informix LANG environment variable used by the SequeLink service.

Type=Dynamic

DataSourceINFService

Sets the value of the Informix SERVICE network parameter used by the SequeLink service.

Type=Dynamic

DataSourceKeySetDataType

Specifies the SQL data type of the auto-unique column (ROWID) in a database table. Valid values are:

- SQL_BINARY
- SQL_CHAR
- SQL_DECIMAL
- SQL_DOUBLE
- SQL_FLOAT
- SQL_INTEGER
- SQL_LONGVARCHAR
- SQL_NUMERIC
- SQL_REAL
- SQL_SMALLINT
- SQL_TYPE_DATE
- SQL_TYPE_TIME
- SQL_TYPE_TIMESTAMP
- SQL_VARBINARY
- SQL_VARCHAR

Examples include:

DB2	SQL_BINARY, SQL_TIMESTAMP
Microsoft SQL Server	SQL_VARCHAR
Sybase	SQL_NUMERIC
ODBC Socket and JDBC Socket	Depends on the backend driver

Type=Dynamic

NOTE: This service attribute is not used by the SequeLink Server for Informix and SequeLink Server for Oracle.

DataSourceKeySetName

Specifies the name of the auto-unique column (Rowid) in a database table, as shown in the following examples:

DB2	The name of any column that enforces uniqueness of the column data, such as ROWID.
Microsoft SQL Server	IDENTITYCOL
Sybase	SYB_IDENTITY
ODBC Socket and JDBC Socket	Depends on the backend driver

The default is an empty string.

Type=Dynamic

NOTE: This service attribute is not used by the SequeLink Server for Informix and SequeLink Server for Oracle.

DataSourceKeySetOptions

Specifies the behavior of SQL Select statements with a auto-unique column (Rowid). Valid values are:

- No_Rowid_with_asterisk, which specifies that a Select statement of the format `SELECT *, ROWID` is not allowed.
- Rowid_with_group_by, which specifies that a Select statement with ROWID in the Select clause and in a Group By clause is allowed.

- Rowid_with_order_by, which specifies that a Select statement with ROWID in the Select clause and in a Order By clause is allowed.
- Rowid_with_table_alias, which specifies that a Select statement with ROWID in the Select clause and a table alias in the From clause is allowed.

Type=Dynamic

NOTE: This service attribute is not used by the SequeLink Server for Informix and SequeLink Server for Oracle.

DataSourceKeySetSize

Specifies the bind size of the auto-unique column (ROWID) in a database table.

DB2	40
Microsoft SQL Server	38
Sybase	38
ODBC Socket and JDBC Socket	Depends on the backend driver.

The default is 0.

Type=Dynamic

DataSourceLimitCursorColumnSize

Limits the columnSize of a result set column to an acceptable value. When a default varchar column is created in FileMaker, the database defines it as 1000000 characters. When the application requests a static cursor on a Unicode datasource, the driver reserves 3000000 (UTF-8) bytes in its internal row buffer for this type of column.

NOTE: In earlier SequeLink releases, this attribute was called DataSourceLimitColumnBindSize.

Limiting the column size makes the row buffer size smaller and reduces disk access.

The default is 0. No value is set.

Type=Dynamic

DataSourceLimitParamBindSize

Specifies the bind size to use when exporting a table that contains null data in a memo column from Microsoft Access. Microsoft Access binds the parameter with a columnSize of 4294967295. The SequeLink Client attempts to allocate a buffer of this size, and typically fails due to lack of memory.

To configure the bind size and get around this application bug, set the parameter bind size for SQL_CHAR, SQL_VARCHAR, SQL_BINARY, SQL_VARBINARY values to a reasonable value.

The default is 0.

Type=Dynamic

DataSourceLogoffBehaviour

Specifies the database action that needs to be taken when an abnormal logoff terminates a transaction. When the attribute is not set, the current transaction will be ended according to the database default logoff behavior.

Valid values are:

- COMMIT=The current transaction will be committed.
- ROLLBACK=The current transaction will be rolled back

Type=Dynamic

DataSourceLogonMethod

Specifies the data store logon method to be used to log on to the data store. The values you use depend on the SequeLink security configuration. Valid values are:

- DBMSLogon(*DBUID*,*DBPWD*)
- DBMSLogon(*UID*,*PWD*)
- OSIntegrated

Setting DataSourceLogonMethod=OSIntegrated means that database authentication relies on the operating system user credentials supplied in the Kerberos ticket. See "ServiceAuthMethods" on page 562 for more information.

See "Configuring Data Store Logon" on page 317 for more information about configuring data source logons refer to the *SequeLink Administrator's Guide*.

Type=Dynamic

z/OS DataSourceMaxCpuAction

Specifies the action to be taken when DataSourceMaxCpuTime is exceeded. Valid values are:

- LOG=Message VAIS208I is written to VAILOG.
- ABORT=The connection will be aborted.

The default is LOG.

Type=Dynamic

DataSourceMaxCpuTime

Specifies (in seconds) the total amount of CPU time that the connection is allowed to consume. If unspecified or set to zero (0), the service is not monitored for exceeding total CPU time. Valid values are:

- 0=Service is not monitored for exceeding total CPU time.
- CPU time_in_seconds=The number of seconds the connection is allowed to consume.

The default is 0.

Type=Dynamic

DataSourceMSSBindAllLOBs

Specifies which columns will be buffered. Valid values are:

- True=The SequeLink Server buffers all columns in the result set.
- False=The SequeLink Server for Microsoft SQL Server does not buffer large object columns that are at the end of a result set. Instead, it uses SQLGetData to get the CLOB/BLOB data as it is requested by the client, which saves memory on the server.

The default is False.

Type=Dynamic

DataSourceMSSConnectOptions

Specifies additional connection options to use when the SequeLink Server connects to Microsoft SQL Server. This string is appended to the connection string. Keyword-value pairs are separated by a semicolon.

You can use this option to connect to a non-default Microsoft SQL Server instance located on another machine. For example:

```
DataSourceMSSConnectOptions=Server=Belg-John\MSSS2000;SnapshotSerializable=1;
```

connects to the Microsoft SQL Server instance named MSSS2000 on the machine Belg-John and enables Snapshot isolation level.

The default is an empty string.

Type=Dynamic

DataSourceMSSCursorType

Specifies the type of cursor the SequeLink Server will use. Valid values are:

- Serverside=Allows multiple concurrent statements (and cursors) to be active at the same time for each connection.
- Serverside-Preserve=Server-side cursors that stay open and positioned when a transaction is committed.
- Clientside=Allows only one active statement/cursor for each connection.

The default is Serverside.

Type=Dynamic

DataSourceMSSCursorWarnings

Microsoft SQL Server sometimes has problems opening server-side cursors for select or stored procedure result sets. This setting masks the warning issued by Microsoft SQL Server each time it encounters this problem. Valid values are:

- True=The warning is not masked from the client application.
- False=The warning is masked from the client application.

The default is False.

Type=Dynamic

DataSourceMSSLogPath

Specifies the log path to be used for all ODBC calls executed by the SequeLink data access service.

Type=Dynamic

DataSourceMSSMapLongToDecimal

Specifies whether the JDBC Client uses the NUMERIC data types instead of the BIGINT data type. Because Microsoft SQL Server 7 does not support the SQL data type BIGINT, a call to setLong results in an exception. Valid values are:

- 0=The JDBC Client uses the BIGINT data type for a call to setLong. An exception results, indicating that an optional feature was not implemented.
- 1=The JDBC Client uses the NUMERIC data type for a call to setLong.

The default value is 0.

Type=Dynamic

DataSourceMSSODBCConnStr

This service attribute is obsolete. Instead, use "DataSourceMSSConnectOptions" on page 522 to specify additional connection options to use when the SequeLink Server connects to Microsoft SQL Server.

DataSourceMSSODBCLogPath

This service attribute is obsolete. Use one of the following service attributes:

- "DataSourceDB2LogPath" on page 509: Specifies the log path for all CLI calls executed for the data access service for SequeLink Server for DB2.
- "DataSourceMSSLogPath" on page 524: Specifies the log path to be used for all ODBC calls executed by the SequeLink data access service.
- "DataSourceSOCODBCLogPath" on page 538: Specifies the log path to be used for all ODBC calls executed by the SequeLink data access service for SequeLink Server for ODBC Socket.
- "DataSourceSYBLogPath" on page 540: Specifies the log path for all CLI calls executed for the data access service for SequeLink Server for Sybase.

DataSourceName

Specifies the name of the server data source. The data source name is defined during configuration and cannot be changed afterward.

The default data source has a fixed name, such as "Default".

Type=Dynamic

DataSourceORAMapTSWTZ

Determines whether columns with a TIMESTAMP WITH TIME ZONE data type are mapped to a timestamp data type when

they are retrieved from an Oracle 9i or higher database. Valid values are:

- TRUE=Columns with a TIMESTAMP WITH TIME ZONE data type are mapped to a timestamp data type when they are retrieved.
- FALSE=An error is generated when TIMESTAMP WITH TIME ZONE columns are retrieved.

The default is FALSE.

Type=Dynamic

DataSourceORANumber0IsNumeric

Specifies how the JDBC Client treats FLOAT columns and columns that are the result of a function (for example, COUNT(*), SUM(...)). Because Oracle describes both types of columns as NUMBER(0,0), the SequeLink Server cannot differentiate between them. SequeLink therefore handles these columns as SQL data type FLOAT. Valid values are:

- 1=SequeLink treats FLOAT columns and all columns that are the result of a function as SQL data type NUMERIC.
- 0=SequeLink treats FLOAT columns and all columns that are the result of a function as SQL data type FLOAT.

The default is 0.

Type=Dynamic

DataSourceORAPublicSchemaSupp

Determines whether the Oracle schema named PUBLIC is supported in catalog statements (for example, SQLTables). Valid values are:

- TRUE=Turns on support.
- FALSE=Turns off support.

Type=Dynamic

DataSourceORARecyclebinSupp

Determines whether database objects in the Oracle recycle bin are shown in the SequeLink catalog statements. Valid values are:

- TRUE=Database objects in the recycle bin are shown in the SequeLink catalog statements.
- FALSE=Database objects in the recycle bin are not shown in the SequeLink catalog statements.

The default is FALSE.

Type=Dynamic

DataSourceORAReportDBLinkAsCatalog

Specifies whether DBLinks are reported as catalogs. Valid values are:

- TRUE=DBLinks are reported as catalogs.
- FALSE=Catalogs are not supported. Use this value for applications such as Microsoft Access that require consistent catalog reporting behavior.

The default is FALSE.

Type=Dynamic

DataSourceORAServiceName

Specifies the Oracle service name to which the application wants to connect. The valid value is a defined Oracle service name.

Type=Dynamic

DataSourceORASynDBLinkObjSupp

Turns on and off support for synonyms of remote Oracle objects in Catalog statements. Valid values are:

- TRUE=Turns on support.
- FALSE=Turns off support.

The default is FALSE.

Type=Dynamic

DataSourceProviderTypesFile

Specifies the name of the file that contains the information returned from the OLE DB ProviderTypes call by an ADO provider data source. The file must be located in the SequeLink installation directory.

Each driver is assigned a specific section name, which is configured in the SequeLink server data source using DataSourceProviderTypesSection.

The default file is swsoc.ini.

Type=Dynamic

DataSourceProviderTypesSection

Specifies the section within the providertypes file that will be used to retrieve OLE DB ProviderTypes information.

The default is Default.

Type=Dynamic

DataSourceReadOnly

Controls read-only access to the SequeLink data access service. Valid values are:

- No=All statements are allowed.
- Select=Only Select statements are allowed.
- Select and batches=Select statements and stored procedures are allowed.
- DBMS=The read-only capabilities of the database are used.

  NOTE: DBMS is not supported on DB2 for z/OS.

The default is No.

Type=Dynamic

DataSourceReportLobsAsLongvar

Enables the JDBC Client to map BLOB and CLOB data types to LONGVARCHAR or LONGVARBINARY.

- TRUE=BLOB and CLOB are mapped to LONGVARCHAR or LONGVARBINARY.
- FALSE=LOBs and getObject on an LOB column use the default behavior for the JDBC Client.

The default is FALSE.

Type=Dynamic

DataSourceSchemaFilterList

A schema name filter for SQLTables and SQLProcedures (ODBC), getTables and getProcedures (JDBC), and TABLES and PROCEDURES (OLE DB/ADO). This attribute is not valid when ServiceCodePage=Database. Valid values are any defined schemas including:

- list_of_schema_names=A comma-separated list of schemas. Only tables owned by the listed schemas are included in the result set. You can use the % or _ character as a search pattern. You cannot use quotes.
- CURRENT SCHEMA=Only tables owned by the current user are returned.

The default is an empty string.

Type=Dynamic

DataSourceSessionToken

Specifies a setting that allows easy correlation of the SequeLink session as displayed in SequeLink monitoring/tracing tools, and the RDBMS session as it can be visualized using the RDBMS-specific monitoring/tracing tools.

Associated with each connection, the SequeLink 6.0 server creates a session token, which is exported to:

- The SequeLink monitor tree, if a monitor profile with session monitoring enabled is active.
- The SequeLink event trace, if an event trace profile that has been configured to trace any session-related events is active.
- An RDBMS-specific session/connection setting, which allows easy correlation of the SequeLink session as displayed in the SequeLink monitor, and the RDBMS session as viewed using RDBMS-specific monitoring and tracing tools.

If the length of the session token exceeds the allowed length of the DBMS field, the session token is truncated and a warning is issued as shown in the following table:

DB2 UDB on Linux, UNIX, and Windows	TP Monitor client accounting string
DB2 on z/OS	Accounting-token (in DB2 statistics and accounting trace records)
Oracle9 or higher	view v$session:*column* CLIENT_IDENTIFIER
SQL Server	Sysprocesses.*program_name*
Sybase	table sysprocesses: *column clientname*

Valid values are a combination of static text and the following placeholders that will be replaced at runtime by the corresponding values. For example:

```
Session from ${ClientInfo} connected to ${DataSourceName}
```

Placeholder Value	Runtime Value
${ClientApi}	One of JDBC, ODBC, ADMIN, or OLEDB
${ClientInfo}	Client IP address or host name
${ClientApplName}	Application name provided by the client application through the ApplicationName connect attribute/keyword
${DataSourceName}	Name of data source for the current session
${ServiceName}	Name of the SequeLink service for the current session
${SessionId}	SequeLink session ID
${ClientApplId}	Application ID being used for the current session
${ClientAutoApplIdDesc}	The text part of the automatic application ID is used for the current session

Type=Dynamic

DataSourceSLKStcCrsrLngCLBuff

Turns on a workaround that allows you to specify the amount of data (in KB) that is buffered for SQL_LONGVARCHAR and SQL_LONGVARBINARY columns with a static cursor. For more information about this client attribute, refer to the *SequeLink Developer's Reference*.

The default is 4.

Type=Dynamic

DataSourceSOCEmptyRowSyntax

Specifies the syntax to use when inserting an empty row into a database table. Valid values are:

- None: Empty rows cannot be inserted.
- SQL92: Empty rows can be inserted using SQL92 syntax, for example:
  `INSERT INTO TABLE default_values`
- EmptyValueList: An empty row with an empty values list can be inserted, for example:
  `INSERT INTO table values ()`
- DefaultValueList: An empty row with an empty values list that contains the default keyword for each column in the target database table can be inserted, for example, for a table with two columns:
  `INSERT INTO table values (default, default)`
- NullValueList, which specifies that an empty row with an empty values list that contains the NULL keyword for each column in the target database table can be inserted, for example, for a table with two columns:
  `INSERT INTO table values (NULL, NULL)`

Type=Dynamic

DataSourceSOCJDBCConnectOptions

Specifies additional connection options to use when the SequeLink Service for JDBC Socket connects to the ODBC-to-JDBC bridge. This string will be appended to the connection string.

NOTE: Set this attribute only when instructed to do so by DataDirect Technologies technical support.

The default value is an empty string.

Type=Dynamic

DataSourceSOCJDBCConnectionURL

Specifies the URL string to pass to the JDBC driver to connect to the database. For example, if you are using the Websphere Information Integrator Classic Federation for z/OS (also known as CrossAccess or DB2 II Classic Federation) driver, the value could be:

```
jdbc:dd-crossaccess30:datasource:tcp/host/port:CODEPAGE=USS
```

The default value is an empty string.

Type=Dynamic

To set up a DataDirect Spy log with the SequeLink Server for JDBC Socket, add the key-value pair spy_log=*Spy log specification* to the DataSourceSOCJDBCConnectionURL to include the DataDirect Spy driver as follows:

```
DataSourceSOCJDBCConnectionURL=DataSourceSOCJDBCConURL;spy_log=(file)filename
```

For example:

```
DataSourceConnectionURL=
jdbc:dd-crossaccess30:EXIMS:tcp/10.30.14.109/9001:CODEPAGE=USS};
spy_log=(file)/tmp/spy.log
```

Other DataDirect Spy options can be specified by adding key-value pairs in which the original Spy-option name is prefixed with spy_. For example:

```
DataSourceConnectionURL=
jdbc:dd-crossaccess30:EXIMS:tcp/10.30.14.109/9001:CODEPAGE=USS;
spy_log= (file)/tmp/spy.log;spy_timestamp_yes
```

DataSourceSOCJDBCDbProperties Name

Specifies the name of a JDBC driver/database-specific properties file that contains the configuration values. Valid values are the file name of a properties file minus the .properties extension.

This properties file is used to set specific database properties that cannot be retrieved by the JDBC driver. In other words, it corrects differences between the ODBC and JDBC APIs. A number of examples are provided, for example, crossaccess30, db2v7mvs, db2v8udb, derby, oracle9i, sqlsrv2000, sybase125.

If you need a properties file for your specific JDBC driver, contact SequeLink technical support.

The default value is an empty string.

Type=Dynamic

DataSourceSOCJDBCDriverClassName

Specifies the name of the JDBC driver being used to connect to the database. For example, if you are using the SequeLink JDBC driver, the value would be:

```
com.ddtek.jdbc.crossaccess30.CrossAccessDriver
```

The default value is an empty string.

Type=Dynamic

DataSourceSOCJDBCLogPath

Specifies the log path to be used for tracing the calls executed by the SequeLink data access service. The valid value is a defined path for a log file.

Type=Dynamic

DataSourceSOCJDBCReadLongDataLength

Specifies whether the SequeLink Server for JDBC Socket can cache long data in order to read the length and provide this information to the application. Valid values are:

- TRUE=The SequeLink Server for JDBC Socket can cache long data. This setting can affect performance.
- FALSE=The SequeLink Server for JDBC Socket cannot cache long data.

Type=Dynamic

DataSourceSOCODBCConnStr

Specifies the connection string used by a SequeLink client application for a connection to an ODBC system data source. When the application provides the optional DBMS user name, password and database values, the values will be appended to the values of the UID, PWD, and DB keywords provided by the client when connecting to the ODBC Socket. The data source specified in the value for this attribute must be a system ODBC data source on Windows.

Example 1

Suppose you want to connect to a Microsoft Access database. You use DataSourceSOCODBCConnStr to connect to the Microsoft

Access driver. You specify the backend ODBC data source name for DataSourceSOCODBCConnStr, for example, DSN=MS_Access_2003.

You can also specify other connection options to overwrite the backend ODBC data source settings that were specified in the ODBC Administrator. For example, the data source MS_Access_2003 is set up to access the database D:\DATA\MSACCESS\db1.mdb.

Example 2

Suppose you want to access two Access databases (db1.mdb and db2.mdb) with one SequeLink Service Socket for ODBC.

First, you create an ODBC data source for your application that uses the SequeLink ODBC driver, specifying only the host and port of the SequeLink Server Socket for ODBC in the ODBC administrator. In this case, you connect to the default SequeLink Server data source, which is configured to access db1.mdb.

To access db2.mdb, you create a new SequeLink server data source, for example, SLAccess_db2mdb, using the SequeLink Command Line Administrator or the SequeLink MMC snap-in.

After creating this new SequeLink server data source, the service attribute DataSourceSOCODBCConnStr contains DSN=MS_Access2003.

Change this value to DSN=MS_Access2003;DBQ=D:\DATA\MSACCESS\db2.mdb.

On the client side, specify the SequeLink server data source SLAccess_db2mdb in your connection string or define it in your ODBC data source.

The default is an empty string.

Type=Dynamic

DataSourceSOCODBCLogPath

Specifies the log path to be used for all ODBC calls executed by the SequeLink data access service for SequeLink Server for ODBC Socket. The valid value is a defined path for a log file.

Type=Dynamic

NOTE: This service attribute replaces DataSourceMSSODBCLogPath, which was used in earlier versions of SequeLink.

DataSourceSybConnectOptions

Specifies additional connection options to use when the SequeLink server connects to the Sybase DBMS. This string will be appended to the connect string.

Valid connection options include:

CharSet	Use the CharSet connection option *only* on Windows when ServiceCodePage=OS. Use CS=cp*ACP*, where *ACP* is determined through the registry (see earlier for an explanation on how to find this out). On a typical English Windows installation, use CS=cp1252.
IANAAppCodePage UNIX	Use the IANAAppCodePage connection option *only* on Linux/UNIX when ServiceCodePage=OS. Use IANAAppCodePage=*value*, where *value* refers to an entry in Table D-11, namely IANAAppCode value. In order to find the correct value to use, you must map the value that is specified in the LC_ALL environment variable to one of the entries in Table D-11 using the first column of the table.

Table D-11. Values for DataSourceSybConnectOptions

OS Code Page	Description[1]	IANAAppCodePage Value
ascii_8	US_ASCII	3
big5	Big5	2026
cp437	IBM437	2008
cp850	IBM850	2009
cp852	IBM852	2013
cp855	IBM855	2046
cp857	IBM857	2047
cp860	IBM860	2048
cp864	IBM864	2051
cp866	IBM866	2052
cp869	IBM869	2054
cp1250	WINDOWS_1250	2250
cp1251	WINDOWS_1251	2251
cp1252	WINDOWS_1252	2252
cp1253	WINDOWS_1253	2253
cp1254	WINDOWS_1254	2254
cp1255	WINDOWS_1255	2255
cp1256	WINDOWS_1256	2256
cp1257	WINDOWS_1257	2257
cp1258	WINDOWS_1258	2258
iso_1	ISO_8859_1	4
iso88592	ISO_8859_2	5
iso88595	ISO_8859_5	8
iso88596	ISO_8859_6	9
iso88597	ISO_8859_7	10
iso88598	ISO_8859_8	11
iso88599	ISO_8859_9	12
iso885915	ISO_8859_15	111

Table D-11. Values for DataSourceSybConnectOptions (cont.)

OS Code Page	Description[1]	IANAAppCodePage Value
koi8	KOI8_R	2084
roman8	roman8	2004
sjis	sjis	17
tis620	tis620	2259

1. Depending on the version of Linux/UNIX that you are using, the actual description of the code page will be different.

The default is an empty string.

Type=Dynamic

DataSourceSYBLogPath

Specifies the log path for all CLI calls executed for the data access service for SequeLink Server for Sybase.

Type=Dynamic

NOTE: This service attribute replaces DataSourceMSSODBCLogPath, which was used in earlier versions of SequeLink.

DataSourceSybNetworkAddress

Specifies the Sybase network address. This string will be appended to the connect string.

The default is installation-dependent.

Type=Dynamic

DataSourceTableTypeFilterList

Specifies a table-type filter for SQLTables (ODBC), getTables (JDBC), and TABLES and PROCEDURES (OLE DB/ADO). This attribute is not applicable when ServiceCodePage=database.

The valid value is a comma-separated list of table types. Valid table types on this platform are:

- ALIAS
- GLOBAL
- SYNONYM
- SYSTEM TABLE
- TABLE
- VIEW

The default is an empty string.

Type=Dynamic

DataSourceThreadMaxRpc

Specifies the number of connection requests that will be accepted before the thread allocated to the connection is released to the thread pool. Valid values are from 0 to 1000.

If 0 is specified, the thread is only released when the connection terminates.

The default is 10.

Type=Dynamic

DataSourceThreadRpcTimeOut

Specifies the idle time (in milliseconds) for threads allocated to connections. Once this value is reached, the thread allocated to the connection is released to the thread pool. Valid values are from 0 to 1000000.

Specify 0 to disable the timeout mechanism.

The default is 2000.

Type=Dynamic

DataSourceTransactionIsolation

Specifies the transaction isolation level used for the connection. Valid values are:

Committed	Dirty reads are not possible. Phantoms and reads that cannot be repeated are possible.
Uncommitted	Phantoms, dirty reads, and reads that cannot be repeated are possible.
RepeatableRead	Dirty reads and reads that cannot be repeated are not possible. Phantoms are possible.
Serializable	Transactions can be serialized. Phantoms, dirty reads, and reads that cannot be repeated are not possible.
Not supported	SequeLink cannot set the transaction isolation level. The DBMS default transaction isolation level will be used.

Refer to your database documentation for a definition of each isolation level.

z/OS For DB2 UDB on z/OS: The default is Committed. The Not Supported value is not allowed.

All other DBMS: The default is Not Supported.

Type=Dynamic

DataSourceWorkArounds

Turns on workarounds that allow you to take full advantage of the ODBC driver with ODBC applications requiring nonstandard or extended behavior.

When this attribute is set in the configuration file, it is valid for all clients that connect to SequeLink. When you add this attribute in the connection string, it is set for that client connection only.

Refer to the *SequeLink Developer's Reference* for more information about connecting using a connection string.

IMPORTANT: Each of these options has potential side effects related to its use. An option should only be used to address the specific problem for which it was designed.

- 1= If an ODBC driver reports to Microsoft Access 2.0 that its SQL_CURSOR_COMMIT_BEHAVIOR or SQL_CURSOR_ROLLBACK_BEHAVIOR is 0, Microsoft Access opens tables as read-only. If this option is on, the ODBC driver returns 1, allowing Microsoft Access to open tables as read-write.
- 2=Some applications cannot handle database qualifiers. If this option is on, the driver reports that qualifiers are not supported.
- 4=Visual Basic 4.0 sometimes requires two connections to a DBMS. For DBMSs that support only a single connection, the second attempt fails. If this option is on, the driver detects

when this condition occurs and has the two ODBC connections share a single physical connection to the DBMS.

- 8=If an ODBC driver cannot detect the number of rows affected by an Insert, Update, or Delete statement, it may return -1 in SQLRowCount. Some products cannot handle this. Turning this option on causes the driver to return 1 instead.
- 16=If an ODBC driver in SQLStatistics reports to Microsoft Access 1.1 that an INDEX_QUALIFIER contains a period, Microsoft Access returns a tablename is not a valid name error. If this option is on, the driver returns no INDEX_QUALIFIER, allowing Microsoft Access to open the table.
- 32=This option allows users of flat-file drivers to abort a long-running query by pressing the ESC key.
- 64=This option results in a column name of Cposition where position is the ordinal position in the result set. For example:

  ```
  SELECT col1, col2+col3 FROM table1
  ```

  produces the column names col1 and C2. SQLColAttributes/SQL_COLUMN_NAME returns an empty string for result columns that are expressions. Use this option for applications that cannot handle empty strings in column names.
- 256=Forces SQLGetInfo/SQL_ACTIVE_CONNECTIONS to be returned as 1.
- 512=To prevent ROWID results, this option forces the SQLSpecialColumns function to return a unique index as returned from SQLStatistics.
- 2048=This option results in SQLDriverConnect returning "Database=" instead of "DB=" in the returned connection string.
- 65536=This option strips trailing zeros from decimal results, which prevents Microsoft Access from generating an error

when decimal columns containing trailing zeros are included in the unique index.

- 131072=This option turns all occurrences of the double quote character (" ") into the accent grave character (`). Some applications always quote identifiers with double quotes. Double quoting causes problems for data sources that do not return SQLGetInfo/SQL_IDENTIFIER_QUOTE_CHAR = ".
- 524288=This option overrides the precision and scale settings for SQL_DECIMAL parameters to precision 40 and scale 20.
- 8388608=This option causes SQLGetInfo/SQL_DATABASE_NAME to be returned as an empty string when SQLGetInfo/SQL_MAX_QUALIFIER_NAME_LEN is 0. This option should be used with Inprise/Borland tools, such as Delphi.
- 536870912=This option allows SQLBindParameter to be called after SQLExecute to change the precision of previously bound parameters.
- 1073741824=Microsoft Access assumes that ORDER BY columns do not have to be in the SELECT list. This option provides a workaround for data stores that always use ORDER BY columns.

The default is 0.

Type=Dynamic

DataSourceWorkarounds2

Turns on workarounds that allow you to take full advantage of the ODBC driver with ODBC applications requiring nonstandard or extended behavior.

When this attribute is set in the configuration file, it is valid for all clients that connect to SequeLink. When you add this attribute in the connection string, it is set for that client connection only.

Refer to the *SequeLink Developer's Reference* for more information about connecting using a connection string.

IMPORTANT: Each of these options has potential side effects related to its use. An option should only be used to address the specific problem for which it was designed.

- 2=Some applications incorrectly specify the ColumnSize/DecimalDigits when binding timestamp parameters. This option causes the driver to ignore the ColumnSize/DecimalDigits specified by the application and use the database defaults instead.
- 4=Microsoft Access uses the most recent native type mapping, as returned by SQLGetTypeInfo, for a specific SQL type. This option reverses the order in which types are returned, so that Microsoft Access will use the most appropriate native type. This option is recommended if you are using Microsoft Access against an Oracle data store.
- 32=Microsoft Access does require that the characters "DSN=" are returned by SQLDriverConnect in the connection string output parameter.

Type=Dynamic

DataSourceWorkaroundsXliteration

Turns on multiple transliteration workarounds that resolve transliteration issues between Shift-JIS/Windows-31j and EUC-JP by mapping "look-alike" characters:

- 0=Transliteration workarounds are not enabled.

- 1=This option identifies a database that stores character data from the Shift-JIS/Windows-31j character set.
- 2=This option identifies a database which stores character data from the EUC-JP character set.

The default is 0.

Type=Dynamic

z/OS MVSDataSourceDB2Plan

Specifies the DB2 plan name used by the data source. This attribute is required only when MVSGlobalDB2Attachment=CAF (see "MVSGlobalDB2Attachment" on page 550).

The default is SLD600PL.

Type=Dynamic

z/OS MVSDataSourceMaxStmtLength

Specifies the maximum length of a SQL statement that the SequeLink Server for DB2 on z/OS can handle.

This setting is only effective for DB2 v8 or higher (new function mode). For DB2 v7 or DB2 v8 (compatibility mode), this setting has no effect; the maximum statement is limited to 32700 bytes.

The default is 128 KB.

Type=Dynamic

z/OS MVSDataSourceUIDMap

Specifies the name of a user ID (UID) map. UID maps can be referenced at both the service and data source level. The valid value is a defined UIDMap name.

A UID map specified with MVSDataSourceUIDMap takes precedence over the UID map specified in the SequeLink service (see the attribute "MVSServiceUIDMap" on page 557). If neither MVSDataSourceUIDMap or MVSServiceUIDMap is specified, no UID map is used.

The default is an empty string. The default for the data source level UIDMap is the UID map specified for the SequeLink service.

Type=Dynamic

z/OS MVSDB2ExitLibrary

Specifies the name of the fully qualified DB2 exit-library (without quotes) for this DB2 interface used to generate JCL. Valid values are z/OS data set names.

NOTE: This library must be APF-authorized.

Type=Static

z/OS MVSDB2LoadLibrary

Specifies the name of the fully qualified DB2 load-library (without quotes) for this DB2 interface used to generate JCL. Valid values are z/OS data set names.

NOTE: This library must be APF-authorized.

Type=Static

z/OS MVSDB2RootDescription

Specifies a general description of the DB2 interface. The valid value is a defined DB2 interface.

Type=Static

z/OS MVSDB2SubsystemName

Specifies the subsystem ID of the DB2 address space. The valid value is a defined interface ID.

The default is MVSDB2InterfaceID.

Type=Static

z/OS MVSDB2Version

Specifies the DB2 version of the DB2 interface. Valid values are:

- V710
- V810
- V910

Type=Static

z/OS MVSGlobalClustername

Specifies the clustername to which the SequeLink Server must register for DNS/WLM-supported Sysplex connection distribution.

Type=Static

z/OS MVSGlobalCommandChar

This attribute is obsolete.

z/OS MVSGlobalCompTrace

Turns on and off tracing for server core components. Valid values are ALL or one of the following component_IDs:

C - Core components
D - DB2 component
L - Generic Log component
S - Server controller component
T - Threadpool component

If unspecified, no tracing will be performed.

The default is no component trace.

Type=Static

z/OS MVSGlobalDB2Attachment

Specifies the type of DB2 attachment to be used for the DB2 service.

MVSGlobalDB2Attachment=RRSAF can only be used when RRS is active. Valid values are:

- CAF
- RRSAF

The default is RRSAF.

NOTE: When ServiceConnectionModel=Threadpool is selected, RRSAF must be selected.

Type=Static

z/OS MVSGlobalDescCode

Specifies the WTO descriptor code for all messages directed to a console by the message logging task. The valid value is a defined WTO descriptor code.

The default is 06.

Type=Static

z/OS MVSGlobalID

NOTE: Set this attribute only when instructed to do so by DataDirect Technologies technical support.

Specifies an identifier to make this "section" unique. This attribute must be set to GLOBAL and must not be changed.

The default is GLOBAL.

Type=Static

z/OS MVSGlobalRouteCode

Specifies the WTO route code for all messages directed to a console by the message logging task. The valid value is a defined WTO route code.

The default is 11.

Type=Static

z/OS MVSGlobalSMFRecordType

Specifies the System Management Facility (SMF) record type to be used for SMF records produced by the server. Valid values are from 128 to 255.

If 0, SMF recording is turned off.

If unspecified, no SMF records are produced.

If you set this value to a positive number, you must inform the system of these records by changing the z/OS SMF parameters in the SMFPRMxx member of the SYS1.PARMLIB data set. You can activate the collection of these records using the SET command.

The recommended value is 197.

Type=Static

z/OS MVSGlobalSosLimit

Specifies the global short-on-storage (SOS) limit, which is the total amount of private storage (in KB) that must be available in the address space to accept further connections. If this amount of storage is unavailable, the incoming connections will be rejected. To turn this feature off, specify 0 or delete this attribute. Valid values are from 0 to 2048.

The default is 1024.

Type=Static

z/OS MVSGlobalSubSysID

Specifies an ID that the SequeLink Server will use as a suffix to construct a unique ResourceManager name to register to RRS. This is a required parameter when MVSGlobalDB2Attachment=RRSAF.

The default is SQLK.

Type=Static

z/OS MVSGlobalWLMEnclaves

Specifies whether the server address space will use WLM enclaves. Valid values are:

- NONE=The server address space will not use WLM enclaves.
- CONNECTION=One new WLM enclave is created for the duration of the connection.
- RPC=A new WLM enclave is created for every SequeLink RPC (each network access to the SequeLink Server).

The default is NONE.

Type=Static

z/OS MVSServiceAdminAuthorizationClass

Specifies a general resource class name to be used by the server.

The default is FACILITY.

Type=Dynamic

z/OS MVSServiceAdminAuthorizationEnable

Enables authorization for the SequeLink Agent service. Valid values are:

- TRUE=MVSServiceAdminAuthorizationClass and/or MVSServiceAdminAuthorizationResource take effect for authorization of your resources with the security manager on z/OS.
- FALSE=Authorization is disabled.

See Chapter 13 "Configuring SequeLink® Security" on page 291 for more information about configuring SequeLink security.

The default is FALSE.

Type=Dynamic

z/OS MVSServiceAdminSecurity

This attribute is obsolete. Instead, use MVSServiceAdminAuthorization.

z/OS MVSServiceAdminSecurityClass

This attribute is obsolete. Instead, use MVSServiceAdminAuthorizationClass.

z/OS MVSServiceAdminSecurityResource

This attribute is obsolete. Instead, use MVSServiceAdminAuthorizationResource.

z/OS MVSServiceAuthorizationClass

Specifies a general resource class name used by the SequeLink Server when authorization is enabled. The valid value is a string.

The default is FACILITY.

Type=Dynamic

z/OS MVSServiceAuthorizationEnable

Enables authorization for the data access service. Valid values are:

- TRUE=MVSServiceAuthorizationClass and/or MVSServiceAuthorizationResource take effect for authorization of your resources with the security manager on z/OS.
- FALSE=Authorization is disabled.

See Chapter 13 "Configuring SequeLink® Security" on page 291 for more information about configuring SequeLink security.

Type=Dynamic

z/OS MVSServiceAuthorizationResource

If MVSServiceAuthorization is enabled, the resource name is used to validate a connection request against the z/OS security system. If the attribute's value is blank or an empty string, the server uses the service name as the resource name to be checked.

NOTE: The value specified for the data source takes precedence over the value specified for the SequeLink service.

The default is the service name.

Type=Dynamic

z/OS MVSServiceDataSourceAuthorization

Specifies that RACF resource security is used for this server data source. When the resource-based authorization has been set up, you can specify resource-based authorization for your server-side data source by setting this attribute to TRUE. See "MVSServiceAuthorizationResource" on page 555 for information about defining the name of the SequeLink server data source.

The default is FALSE.

Type=Dynamic

z/OS MVSServiceDataSourceSecurity

This attribute is obsolete. Instead, use MVSServiceDataSourceSecurity.

z/OS MVSServiceDB2InterfaceID

Specifies the ID of a DB2 interface. A DB2 interface is a set of parameters that describe a DB2 subsystem which can be accessed by the SequeLink Server for z/OS.

Type=Static

z/OS MVSServiceLoadModule

Specifies the name of the load module for the service. Valid values are uppercase.

The default is VAISTHRD (for SequeLink Server for DB2 services).

Type=Static

z/OS MVSServiceSecurity

This attribute is obsolete. Instead, use MVSServiceAuthorizationEnable.

z/OS MVSServiceSecurityClass

This attribute is obsolete. Instead, use MVSServiceAuthorizationClass.

z/OS MVSServiceSecurityResource

This attribute is obsolete. Instead, use MVSServiceAuthorizationResource.

z/OS MVSServiceUIDMap

Specifies the name of a user ID (UID) map. The valid value is a defined UIDMap name. A UID map specified in the server data source takes precedence over the UID specified in the SequeLink service (see "MVSDataSourceUIDMap" on page 548).

If this attribute is not specified, it is possible to have some data sources that use UID maps while other data sources do not.

The default is an empty string.

Type=Dynamic

z/OS MVSUID

Specifies an entry in the user ID (UID) map. The format of this entry is: user=*mapped_user* or *=*mapped_user*, where:

- user is a valid user or user group for the z/OS security system
- * is a wildcard for any user
- *mapped_user* is a valid DB2 authorization ID

*=*mapped_user* is required when anonymous authentication is configured for the z/OS server (ServiceAuthMethods=Anonymous).

Type=Dynamic

z/OS MVSUIDDefaultAccess

Specifies the default action for a user ID (UID) map. A UID map set in the server data source takes precedence over a UID map set in the SequeLink service. If the server data source does not specify a UID map, the UID map defined in the SequeLink service is used.

Valid values are:

- PERMIT=If user ID mapping is set for the SequeLink service or data source and the user ID cannot be found in the UID map, the connection is accepted.
- DENY=If user ID mapping is set for the SequeLink service or server data source and the user ID cannot be found in the UID map, the connection is refused.

Type=Dynamic

z/OS MVSUIDMapDescription

Specifies a general description of the user ID (UID) mapping table. The valid value is a defined user ID map name.

Type=Dynamic

ServiceAdminAuthMethods

Specifies one or multiple authentication mechanisms that the SequeLink Manager can use to authenticate itself to the server. Valid values are:

- anonymous
- OSLogon(*UID*,*PWD*)
- kerberos
- integrated_nt

See Chapter 13 “Configuring SequeLink® Security” on page 291 for more information about SequeLink security features.

Type=Dynamic

ServiceAdministrator

Sets authorization for users who are allowed to manage SequeLink services using the SequeLink Manager.

NOTES:

- On Windows, users who are allowed to manage SequeLink services using the SequeLink Manager must have administrator rights.
- On Linux, UNIX, and Windows, specify *everyone* when ServiceAuthMethod=anonymous.

Valid values for the ServiceAdministrator attribute are:

- *user_name*=The user ID of a user who is allowed to use the SequeLink Manager. To configure authorization for multiple users, you must set this attribute multiple times, one instance of the attribute for each user. For example:

  ```
  ServiceAdministrator=RSMITH
  ServiceAdministrator=DJONES
  ServiceAdministrator=TCONRAD
  ```

  NOTE: On Windows servers, you must prefix the user ID with the Windows server name or the Windows domain name, for example, SALES\DJONES. When connecting, the user must also prefix the user ID with the Windows server name, if connecting to a local server, or the Windows domain name.

  NOTE: Alternatively, you can set the ServiceAdministratorGroup attribute to configure authorization for groups of users defined on Linux, UNIX, or Windows. See "ServiceAdministratorGroup" on page 561 for more information about configuring authorization for user groups on Linux, UNIX, and Windows.

- authenticated=Any user who can provide a valid host user ID and password or who uses Integrated authentication will receive the same authorization.

- everyone=All connections will receive the same authorization level, regardless of how they are authenticated.

See Chapter 13 "Configuring SequeLink® Security" on page 291 for more information about configuring SequeLink security.

Type=Dynamic

ServiceAdministratorGroup

Sets authorization for defined Linux, UNIX, and Windows user groups who are allowed to manage SequeLink services using the SequeLink Manager. Valid values are user groups defined on Linux, UNIX, and Windows.

To configure authorization for multiple user groups, you must set the ServiceAdministrator attribute multiple times, one time for each user group. For example:

```
ServiceAdministratorGroup=SLUSERG1
ServiceAdministratorGroup=SLUSERG2
ServiceAdministratorGroup=SLUSERG3
```

NOTE: On Windows servers, you must prefix the user group ID with the Windows server name or the Windows domain name where the group is defined, for example, SALES\SLUSERG1.

See Chapter 13 “Configuring SequeLink® Security” on page 291 for more information about configuring SequeLink security.

Type=Dynamic

ServiceAdminKerberosPrincipalName

Specifies the case-sensitive service principal name to be used for Kerberos authentication for the SequeLink Agent service. Use this attribute to specify another service principal name other than the default service principal name. The name must match the service principal name in the KDC.

For example:

```
SLDB2UDB60/server1.eu.datadirect.com@XYZ.COM
DB2SRV/IRONFLEX.DATADIRECT.COM@IRONFLEX.TEST
```

Default=*host/fully_qualified_domain_name@REALM*

Type=Dynamic

For information about specifying a service principal name to be used for Kerberos authentication to a SequeLink data access service, see "ServiceKerberosPrincipalName" on page 579.

ServiceAuthMethods

Specifies one or multiple authentication mechanisms the service accepts. The client must select the supported mechanism to authenticate itself to the server.

Valid values are:

Windows
- anonymous
- OSLogon(*UID*,*PWD*)
- OSLogon(*HUID*,*HPWD*)
- OSLogon(*UID*,*PWD*,*NPWD*)
- OSLogon(*HUID*,*HPWD*,*NPWD*)
- kerberos
- integrated_nt

Linux/UNIX
- anonymous
- OSLogon(*UID*,*PWD*)
- OSLogon(*HUID*,*HPWD*)
- kerberos

z/OS
- Anonymous
- Kerberos
- OSLogon(*UID*,*PWD*)
- OSLogon(*HUID*,*HPWD*)
- OSLogon(*UID*,*PWD*,*NPWD*)
- OSLogon(*HUID*,*HPWD*,*NPWD*)

See Chapter 13 "Configuring SequeLink® Security" on page 291 for more information about configuring SequeLink security.

NOTE: UID mapping is required when ServiceAuthMethods=Anonymous.

Type=Dynamic

ServiceAuthorizedAdminClient

Specifies one or multiple client TCP/IP network identifiers that are allowed to access the service using an administrator client. Use configuration values to limit the clients that are allowed administrative connections to the agent service.

This attribute supports both address and name formats as shown in the following examples:

Client TCP/IP host name	`burner.ddtek.com`
Client TCP/IP domain names through the use of a wild card	`192.16.2.*` or `*.ddtek.com`
Client TCP/IP address	`127.0.0.1??`
Client TCP/IP address range	`192.16.*.*`

NOTE: When using host names, ServiceResolveHostNames must be set to True, and only primary domain names can be used.

To configure multiple TCP/IP location filters, you must set the ServiceAuthorizedAdminClient attribute multiple times, one instance for each location filter. For example:

```
ServiceAuthorizedAdminClient=192.16.*.*
ServiceAuthorizedAdminClient=192.17.*.*
ServiceAuthorizedAdminClient=192.18.*.*
```

See "TCP/IP Location Filters" on page 296 for more information.

Type=Dynamic

ServiceAuthorizedClient

Specifies one or multiple client TCP/IP network identifiers that are allowed to access the service using a client application. Use configuration values to limit the number of clients that are allowed data access connections to data access services.

The attribute supports both address and name formats as shown in the following examples:

Client TCP/IP host name	`burner.ddtek.com`
Client TCP/IP domain names through the use of a wild card	`192.16.2.*` or `*.ddtek.com`
Client TCP/IP address	`127.0.0.1`
Client TCP/IP address range	`192.16.*.*`

NOTE: When using host names, ServiceResolveHostNames must be set to True, and only primary domain names can be used.

To configure multiple TCP/IP location filters, you must set the ServiceAuthorizedClient attribute multiple times, one instance for each location filter. For example:

```
ServiceAuthorizedClient=192.16.*.*
ServiceAuthorizedClient=192.17.*.*
ServiceAuthorizedClient=192.18.*.*
```

See “TCP/IP Location Filters” on page 296 for more information.

Type=Dynamic

ServiceCancelEnabled

Specifies whether cancelling connection requests using SQLCancel is supported. Valid values are:

- TRUE=yes
- FALSE=no

The default is TRUE.

Type=Static

ServiceCatchExceptions

Specifies how the SequeLink service will handle exceptions. Valid values are:

- TRUE=The service will attempt to recover from unexpected exceptions.
- FALSE=The exception is passed to the operating system resulting in an error or core dump.

The default is TRUE.

Type=Static

ServiceCodePage

Controls transliteration for the SequeLink service. Valid values are:

- Default=SequeLink supports standard ASCII/EBCDIC transliteration for ODBC and OLE DB/ADO Clients, and transliteration to UTF-16 for JDBC and .NET Clients.
- OS=The client transliterates character data from the code page the SequeLink service is using to the code page of the client application/system.
- Database=Internationalization support features are enabled. The SequeLink Client uses the optimal encoding to exchange character information between client and server. This value minimizes the amount of transliteration required. Graphic data types are always supported.

  NOTE: The Database value is supported only for the ODBC, JDBC, and .NET Clients. This value cannot be used with SequeLink Server for Informix.

The default is Default.

Type=Dynamic.

ServiceCodePageMap

Controls transliteration for the SequeLink service when the code page retrieved from the operating system or database is not recognized. When ServiceCodePage is set to OS or Database, operating system- or database-specific code pages can be mapped to the IANA character set name, overwriting or extending the service built-in mappings. For more information about the IANA character set, refer to the IANA Web site.

The format of the value depends on the type of mapping that is being extended. In general, the value is a semi-colon separated

list of values, where the first or second value specified is the code page value retrieved from either the operating system or the database. The ianaName is the official IANA identifier for the code page.

Valid values are:

- OS=The code page value returned from the operating system. The syntax is

  OS code page identifier;*ianaName*;

- Oracle9=The code page value returned from an Oracle 9.x or higher database. The syntax is:

  charSetId;*oracleName*;*ianaName*;[maxByteSz;fixedWidth;]

Type=Dynamic

ServiceConnectInfo

Specifies the TCP/IP port on which the service is listening for connection requests. The port is specified using the syntax:

tcp://host.port

where host is the name of the host on which the SequeLink service runs and port is an available TCP port.

Type=Static

ServiceConnectionModel

Specifies the connection model to be used for connections to the SequeLink service. If you change to another connection model, you must delete and recreate any monitoring or event trace profiles. Valid values are:

- ThreadPool=A pool of threads is created with ServiceMinThreads prestarted threads and maximum ServiceMaxThreads threads. This thread pool is used to service client connection requests.
- Process/Connection=SequeLink creates a separate OS process to service client connection requests. (Not valid for z/OS.)
- Thread/Connection=SequeLink creates a separate thread for each client connection requests.

The default is ThreadPool.

Type=Static

z/OS ServiceDB2MaxThreads

Limits the number of DB2 threads used by the SequeLink DB2 service. This value must be higher than the value of the ServiceMaxThreads parameter and lower than the value of the MAX_BATCH_CONNECT parameter in the Thread Management panel (DSNTIPE) of the DB2 installation.

A value of 0 means that the number of threads will not be limited by the SequeLink DB2 service.

The default is 0.

Type=Dynamic

z/OS ServiceDB2MaxTransactions

Specifies the number of DB2 transactions that a DB2 thread is allowed to service before it is closed. This ensures that when the DB2 thread terminates, the locks set on temporary tables are released.

A value of 0 means that the DB2 thread is not terminated.

The default is 0.

Type=Dynamic

z/OS ServiceDB2MinThreads

Specifies the number of reusable DB2 threads that will not be released by the release timer routine or by the RRS RELEASE operator command.

A value of 0 means that all reusable DB2 threads that exceed an age of 300 seconds will be released.

The default is 0.

Type=Dynamic

ServiceDBConnectionPooling

Enables the connection pooling feature in the Driver Manager that is used for SequeLink Server for ODBC Socket and SequeLink Server for JDBC Socket on Windows. Valid values are:

- Disabled=The Driver Manager's connection pooling feature is disabled.
- Enabled=Connection pooling is enabled. Using connection pooling can result in significant performance improvements.

The default is Disabled.

Type=Static

ServiceDeadClntDetInt

Specifies the interval (in seconds) between requests that are sent from the SequeLink Server to the SequeLink Client to verify the availability of the SequeLink Client. Valid values are 0 and any integer from 61 to 1000000.

A value of 0 disables the dead client detection feature.

The default is 600.

Type=Static

ServiceDebugLogLevel

Specifies the level of detail for messages logged in the debug log file. One or multiple message levels can be enabled/disabled. Valid values are:

- Fatal
- Error
- Warnings
- Information
- Debug
- SSP Packet log
- SSP Requests

The default is Fatal.

NOTE: The value of this attribute is a bitmask with each bit having the following decimal values when turned on:

Fatal:	Bit 0=1
Error:	Bit 1=2
Warnings:	Bit 2=4
Information:	Bit 3=8
Debug:	Bit 4=16
SSP Packet Log:	Bit 5=32
SSP Requests:	Bit 6=64

If you set this attribute using a SequeLink Manager command that prompts for a decimal value, such as the ServiceAttributeReplace command, the value you set must equal the total decimal value of the bits you want to turn on. For example, if you want to turn all bits on, meaning all options would be logged, you would set the attribute to 127 (1+2+4+8+16+32+64=127). To turn off all bits, set the attribute to 0.

Type=Dynamic

ServiceDebugLogPath

Specifies the directory where debug log files are written. If the log path is fully specified, the log is written to the specified directory. If only a filename is specified, the log file is written in the *install_dir*\logging directory.

Contains the directory in which debug log files will be written.

z/OS Specifies a UNIX System Services HFS directory to which Log files will be written.

The default is none.

Type=Static

ServiceDescription

Specifies a general description of the SequeLink service.

Type=Static

ServiceDetailedOSLogonErrors

Specifies what type of error will be returned when the OSLogon based authentication fails. Valid values are:

- TRUE=A detailed error will be returned.
- FALSE=A generic error will be returned.

The default is TRUE.

Type=Dynamic

ServiceEncryptionAlgorithm

Specifies the data scrambling algorithm used when sending requests or replies across the network between client and server. Valid values are:

- None
- DES
- 3DES
- Byteswap

The default is None, which means cleartext.

See Chapter 13 "Configuring SequeLink® Security" on page 291 for more information about configuring SequeLink security.

Type=Static

ServiceEnvironmentVariable

Specifies a list of variables that will be set before the SequeLink service is started. The syntax for valid values is varname=value. For example, ORACLE_SID=ORA10g.

To define more than one variable, you must add more than one instance of the attribute to your service.

Examples:

- To specify the Shared library path for the ODBC-to-JDBC bridge and the JVM *LIBPATH* , where *LIBPATH* is:
  - LIBPATH on AIX and z/OS USS
  - LD_LIBRARY_PATH on Solaris and Linux
  - SHLIB_PATH on HP-UX
  - PATH on Windows

  For example, on Windows:

```
ServiceEnvironmentVariable=PATH=D:\Programs\java\j2sdk1.4.2_05\jre\bin\
server;D:\Programs\java\j2sdk1.4.2_05\jre\bin
```

- To specify options used by the backend JDBC Driver:

```
ServiceEnvironmentVariable=SL_JAVA_OPTIONS=-Xms32m -Xms64m -Xcomp -verbose
```

- To specify the directory for the license ini file on Linux or UNIX:

```
ServiceEnvironmentVariable=ISLVINI=/usr/sl60/server/ipe
```

- To configure the SequeLink Oracle service to use the Oracle Net Service on Windows:

  ```
  ServiceEnvironmentVariable=LOCAL=NetServiceName
  ```

  where *NetServiceName* is the Oracle Net service name. See "Configuring a SequeLink® Service for Oracle" on page 426 for detailed instructions on configuring this SequeLink service.

Type=Static

ServiceEventTraceLocation

Specifies the directory where the event trace file is located. The valid value is an existing directory.

The default is an empty string.

Type=Static

ServiceEventTraceSize

Specifies the size (in bytes) of the event trace file. Valid values are from 10000 to 2000000000. On z/OS, the specified value is rounded to a 4 KB page value.

The default is 1000000.

Type=Static

ServiceEvQGetNrEventsMax

Specifies the number of events to get from the event queue in one read operation.

The default is 10.

Type=Static

NOTE: Set this attribute only when instructed to do so by DataDirect Technologies technical support.

ServiceEvQPingTimeout

Specifies the timeout (in milliseconds) the SequeLink Agent will wait for a SequeLink data access service response while determining whether the service is started.

The default is 1000.

Type=Static

ServiceEvQShmMonitorSize

Specifies the size of the shared memory segment (in bytes) to be used for the shared monitor counters. The shared monitor counters are monitor values that are accessible using the Windows performance monitor integration.

The default is 4096.

Type=Static

ServiceEvQShmQMaxResend

Specifies the number of attempts while sending events on the event queue before generating a fatal error.

The default is 5.

Type=Static

ServiceEvQShmQSize

Specifies the size of the shared memory segment for the event queue. The event queue is used as a communication path between SequeLink services.

The default is 65536.

Type=Static

ServiceEvQShmQWaitResend

Specifies the time (in milliseconds) to wait between successive attempts to send events on the event queue.

The default is 50.

Type=Static

NOTE: Set this attribute only when instructed to do so by DataDirect Technologies technical support.

ServiceExecPath

Specifies the path of a SequeLink Server executable and is used differently depending on platform. The valid value is a defined location of a SequeLink server executable.

- On Windows: The value of the ServiceExecPath is used when registering the service.
- On Linux and UNIX: The value of ServiceExecPath is used by the SequeLink Agent when starting the service.

Do not alter this attribute.

Type=Static

ServiceExecPath2

Specifies the path of a SequeLink server executable to start by SequeLink service starter. The valid value is a defined location of a SequeLink server executable.

Type=Static

ServiceHost

Specifies the name of the host on which the SequeLink service is installed.

Type=Dynamic for newly-created services

Type=Read-only for existing services

ServiceIdleTimeInt

Specifies time (in seconds) a connection can be idle before it is terminated by the SequeLink Server. Because ServiceIdleTimeInt is static, you must restart the SequeLink service to make the change effective.

Valid values are 0 and any integer from 60 to 1000000.

- If ServiceIdleTimeInt is greater than ServiceDeadClntDetInt, the value will be rounded down to the nearest multiple of ServiceDeadClntDetInt.
- If ServiceIdleTimeInt is less than ServiceDeadClntDetInt, the value of ServiceIdleTimeInt will be used.

The default is 0. Idle time checks are disabled.

Type=Static

ServiceIIOPObjectKey

Specifies IIOP object key in IIOP header.

The default is IIOP:slx::.

Type=Static

NOTE: Set this attribute only when instructed to do so by DataDirect Technologies technical support.

ServiceIIOPOperationTarget

NOTE: Set this attribute only when instructed to do so by DataDirect Technologies technical support.

Specifies IIOP operation target in IIOP header.

The default is SSP.

Type=Static

ServiceINFMaxNrActStat

Specifies the maximum number of active statements that are allowed for each connection to an Informix database.

The default is 250.

Type=Dynamic

ServiceInternalTimeout

Specifies the number (in milliseconds) that thread pool synchronization actions block before generating an internal error. Valid values are positive numbers.

The default is 60000.

Type=Static

ServiceKerberosPrincipalName

Specifies the case-sensitive service principal name to be used for Kerberos authentication to the SequeLink data access service. Use this attribute to specify another service principal name other than the default service principal name. The name must match the service principal name in the KDC.

For example:

```
SLDB2UDB60/server1.eu.datadirect.com@XYZ.COM
DB2SRV/IRONFLEX.DATADIRECT.COM@IRONFLEX.TEST
```

Default=*host/fully_qualified_domain_name@REALM*

Type=Dynamic

On Linux/UNIX, the service principal name must exist in the keytab file.

For information about specifying a service principal name to be used for Kerberos authentication to a SequeLink Agent service, see "ServiceAdminKerberosPrincipalName" on page 561.

ServiceLanguage

The specified language for SequeLink messages. The only valid value and the default is 1= English.

Type=Static

ServiceMaxSessions

Specifies the maximum number of sessions a multithreaded SequeLink Server will accept.

On z/OS, verify that this parameter does not exceed the value of MAXFILEPROC in BPXPRMxx. The MAXFILEPROC parameter can be changed dynamically.

- On z/OS, the default value is 2000.
- On Linux, UNIX, and Windows, the default value is 0. A value of 0 means unlimited number of sessions is accepted.

On Linux, UNIX, and Windows, the default value is 0. A value of 0 means unlimited number of sessions is accepted.

Type=Dynamic

ServiceMaxThreads

The maximum number of threads that can be started in the thread pool. This attribute is ignored when ServiceConnectionModel is set to Thread/Connection or Process/Connection.

Valid values are:

- For DB2 UDB on z/OS, valid values are from 6 to 256.
- For DB2 UDB on AIX, valid values are from 6 to 512.
- For all others, valid values are from 6 to 64000.

The default is 64.

Type=Static

ServiceMessageFile

The location of the service message file:

On Windows: The valid value is the path to a resource-only DLL.

On Linux/UNIX: The valid value is the path to a .cat file.

Type=Static

ServiceMinThreads

The number of pre-started threads that will be started in the thread pool. The value should be equal to or less than the value of ServiceMaxThreads. This attribute is ignored when ServiceConnectionModel is set to Thread/Connection or Process/Connection.

Valid values are:

- For DB2 UDB on z/OS, valid values are from 6 to 256.
- For DB2 UDB on AIX, valid values are from 6 to 512.
- For all others DBMSs, valid values are from 6 to 64000.

The default is 8.

Type=Static

ServiceMSSMergeXaBranches

When accessing Microsoft SQL Server, enables the application to merge all XA branches as one transaction branch, emulating tightly coupled branch transactions. When disabled, multiple

connections within the same distributed transaction ignore each other's locks.

Valid values are:

- TRUE=SequeLink merges all XA branches as one transaction branch.
- FALSE=All XA branches are kept separately.

The default is FALSE.

Type=Static

ServiceName

Specifies the name of the service to be specified during service creation. After a service is created, the name of a service cannot be changed.

Type=Dynamic for new services

Type=Read-only for existing services

ServiceORASerializeLogon

Specifies whether all Oracle API calls that are executed to establish a connection with the Oracle database are serialized.

Valid values are:

- TRUE=The Oracle API calls are serialized.
- FALSE=The Oracle API calls are not serialized.

The default is TRUE.

Type=Static

ServiceRegisterTCPPort

Specifies whether the TCP/IP port will be registered automatically in the operating systems services file (/etc/services on Linux/UNIX or %SystemRoot%\system32\drivers\etc\services on the Microsoft Windows platforms on which SequeLink Server runs, for example).

- TRUE=The TCP/IP port used by the SequeLink service will be registered automatically in the operating system services file.
- FALSE=The TCP/IP port used by the SequeLink service will not be registered automatically in the operating system services file.

The default is FALSE.

Type=Static

ServiceResolveHostNames

Specifies how host names are resolved. Valid values are:

- TRUE=Information about connected clients is displayed using symbolic system names.
- FALSE=Information about connected clients is displayed in an IP format.

The default is FALSE.

Type=Dynamic

ServiceSSLCipherSuites

Specifies the SSL cipher suites that SequeLink supports for data transfers between the SequeLink Client and SequeLink Server. Valid values are:

- SSL_DH_anon_EXPORT_WITH_RC4_40_MD5
- SSL_DH_anon_WITH_RC4_128_MD5
- SSL_DH_anon_EXPORT_WITH_DES40_CBC_SHA
- SSL_DH_anon_WITH_DES_CBC_SHA
- SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
- TLS_DH_anon_WITH_AES_128_CBC_SHA
- TLS_DH_anon_WITH_AES_256_CBC_SHA

The default is TLS_DH_anon_WITH_AES_128_CBC_SHA.

Type=Static

ServiceSSLEnabled

Specifies whether SSL is enabled for the SequeLink service. Valid values are:

- TRUE=The SequeLink service is enabled for SSL.
- FALSE=The SequeLink service is not enabled for SSL.

The default is FALSE.

Type=Static

ServiceSSLSessionCacheSize

Specifies the maximum number of SSL (or TLS) session identifiers that are cached on the SequeLink Server before the sessions are flushed from the cache. Caching SSL session identifiers so that clients can reuse them in a subsequent connection improves

performance because the CPU-intensive computations required to create an SSL session are eliminated.

Valid values are from 0 to 100000.

The default is 10000.

Type=Static

ServiceSSLSessionCacheTimeout

Specifies the maximum time (in seconds) to keep a SSL (or TLS session in the SequeLink Server session cache.

Valid values are from 0 to 100000.

The default is 300.

Type=Static

ServiceSSLVersions

Specifies the version of the SSL standard that is used for encryption. Valid values are:

- SSL 3.0
- TLS 1.0

The default is TLS 1.0.

Type=Static

ServiceTCP1stRecvTimeLimit

Maximum amount of time (in seconds) the server will wait to receive the first network packet.

The default is 30.

Type=Static

NOTE: Set this attribute only when instructed to do so by DataDirect Technologies technical support.

ServiceThreadLockThreshold

Specifies a percentage of the value of ServiceMaxThreads. When the number of active threads is less than this percentage, a connection that has executed more RPCs than DataSourceThreadMaxRpc on the current thread is allowed to lock this thread for another DataSourceThreadRpcTimeOut period.

This attribute is required when ServiceConnectionModel=Threadpool (see "ServiceConnectionModel" on page 568 for more information).

Valid values are from 0 to 50.

On Linux, UNIX, and Windows, the default is 0 (the thread switching mechanism is disabled, improving performance in times of low system activity).

z/OS On z/OS, the default is 10.

Type=Dynamic

ServiceUnixSyslogFacility

Facility with which all SequeLink syslog messages are logged. Valid values are:

- USER
- LOCAL0-LOCAL7

The default is USER.

Type=Dynamic

ServiceUser

Sets authorization for users who are allowed to access the service for data access.

On Linux, UNIX, and Windows: Specify everyone when ServiceAuthMethods=anonymous. To configure authorization for multiple users, you must set the ServiceUser attribute for each user. Valid values are:

- *user_id*=The user ID of a user who is allowed to use the SequeLink service. To configure authorization for more than one user, you must configure this attribute multiple times, one instance for each user.

  NOTES:

  - On Windows servers, you must prefix the user ID with the Windows server name or the Windows domain name, for example, SALES\DJONES.
  - Alternatively, you can set the ServiceUserGroup attribute to configure authorization for groups of users defined on Linux, UNIX, or Windows. See "ServiceUserGroup" on page 588 for more information about configuring authorization for user groups on Linux, UNIX, and Windows.

- authenticated=Any user who can provide a valid host user ID and password or who uses Integrated NT authentication will receive the same authorization.

- everyone=All connections will receive the same authorization level, regardless of how they are authenticated.

See Chapter 13 "Configuring SequeLink® Security" on page 291 for more information about configuring SequeLink security.

Type=Dynamic

ServiceUserGroup

Sets authorization for defined Linux, UNIX, and Windows user groups who are allowed to access the service for data access. Valid values are user groups defined on Linux, UNIX, and Windows.

To configure authorization for multiple user groups, you must set the ServiceAdministrator attribute multiple times, one time for each user group. For example:

```
ServiceUserGroup=SLUSERG1
ServiceUserGroup=SLUSERG2
ServiceUserGroup=SLUSERG3
```

NOTE: On Windows servers, you must prefix the user group ID with the Windows server name or the Windows domain name where the user group is defined, for example, SALES\SLUSERG1.

See Chapter 13 "Configuring SequeLink® Security" on page 291 for more information about configuring SequeLink security.

Type=Dynamic

E SequeLink® Events

This appendix lists the SequeLink events, describes the events, lists the attributes associated with the events, and explains how to write a filter for an event.

SequeLink® Events and their Attributes

Table E-1 lists SequeLink events and the attributes that apply to each attribute. See Table E-2 for the definition of the attributes.

Table E-1. SequeLink Events

Event	Description
DataSourceSet Attributes: SessionId, ClientInfo, DataSourceName	Data Source for the specified session.
DbmsSessionStarted Attributes: SessionId, ClientInfo, ServiceUser, DbmsSessionId, DbmsUser	DBMS Connection opened.
DbmsSessionStopped Attributes: SessionId, ClientInfo, ServiceUser, DbmsSessionId, DbmsUser	DBMS Connection closed.

Table E-1. SequeLink Events *(cont.)*

Event	Description
ErrorInternal Attributes: ErrorCode, ErrorMessage, [SessionId, ClientInfo, ServiceUser]	Fatal error occurred.
ErrorOccurred Attributes: ErrorCode, ErrorMessage, [SessionId, ClientInfo, ServiceUser]	Error occurred.
EvProcStarting Attributes: none	Event Processing is started.
EvProcStopping Attributes: none	Event Processing is stopped.
NetPacketRead Attributes: SessionId, ClientInfo, ServiceUser, NumberOfBytes	Packet read from network (from client).
NetPacketWrite Attributes: SessionId, ClientInfo, ServiceUser, NumberOfBytes	Packet written to network (to client).
RowsFetched Attributes: SessionId, ClientInfo, ServiceUser, StatementId, RowsFetched	Rows fetched from the DBMS.

Table E-1. SequeLink Events *(cont.)*

Event	Description
ServiceParams Attributes: DebugLogLevel[a]	Service Parameters.
ServiceStarted Attributes: ProcessId	Service started.
ServiceStopping Attributes: SwitchCntR, SwitchCntT	Service stopped.
SessionAuthenticated Attributes: SessionId, ClientInfo, ServiceUser, Authorization	Authentication succeeded.
SessionIdentification Attributes: SessionId, ClientInfo, DataSourceName, ClientApplName, ClientApplId, ClientAutoApplIdDesc, SessionToken	Sets session identification.
SessionParams Attributes: SessionId, ClientInfo, DebugLogLevel*	Session Parameters.
SessionStarted Attributes: SessionId, ClientInfo, Authorization, DataSourceName, ClientApi	Session started.

Table E-1. SequeLink Events *(cont.)*

Event	Description
SessionStopped Attributes: SessionId, ClientInfo, ServiceUser, ReceivedPackets*, SumReceivedPackets*, MinReceivedPackets*, MaxReceivedPackets*, SentPackets*, SumSentPackets*, MinSentPackets*, MaxSentPackets*, RowsFetched, ExecuteCount, TxnPrepare, TxnCommit, TxnRollback	Session stopped.
StatementClosed Attributes: SessionId, ClientInfo, ServiceUser, StatementId, RowsFetched, RowsAffected, ExecuteCount	DBMS Statement closed.
StatementExecuted Attributes: SessionId, ClientInfo, ServiceUser, StatementId, Statement, RowsAffected, ReturnCode, Statement	DBMS Statement executed.
StatementOpened Attributes: SessionId, ClientInfo, ServiceUser, StatementId	DBMS Statement opened.
TransactionCommit Attributes: SessionId, ClientInfo, ServiceUser	DBMS Transaction Committed.

Table E-1. SequeLink Events (cont.)

Event	Description
TransactionPrepare Attributes: SessionId, ClientInfo, ServiceUser	DBMS Transaction prepared.
TransactionRollback Attributes: SessionId, ClientInfo, ServiceUser	DBMS Transaction rollback occurred.

a. These attributes cannot be specified within a filter.

SequeLink® Event Attributes

Table E-2 describes each event attribute.

Table E-2. SequeLink Event Attributes

Event Attribute	Description	Type
Authorization	Session authorization (Administrator, User).	String
ClientApi	Type of client application (for example, Administrator, ODBC, ADO/OLE DB, .NET, or JDBC).	String
ClientApplI	Application ID used to authorize the current sessions.	String
ClientApplName	Application name as provided by the client application when connecting (ApplicationName property).	String
ClientAutoApplId	Application ID used to authorize the current sessions.	String
ClientAutoApplIdDesc	Text portion of the application ID used to authorize the current sessions.	String
ClientInfo	IP address or host name of the client system.	String
DataSourceName	Identification of the data source used for the session.	String

Table E-2. SequeLink Event Attributes *(cont.)*

Event Attribute	Description	Type
DbmsSession	Identification of the DBMS session.	String
DbmsUser	User used to open the DBMS session.	String
ErrorArgument	One or more arguments completing the ErrorMessage attribute.	String
ErrorCode	Numeric error code associated with an error event.	Integer
ErrorMessage	Error text.	String
EventId	Numeric identification of the event.	Integer
ExecuteCount	Number of SQL statements executed.	Integer
NumberOfBytes	Number of bytes sent or received.	Integer
ProcessId	Service process ID assigned by the operating system.	Integer
ReturnCode	Result of statement execution: 0 when OK, -1 when Error, -2 when Warning.	String
RowsAffected	Rows affected by SQL statement(s).	Integer
RowsFetched	Number of rows fetched.	Integer
ServiceName	Name of the service generating the event.	String
ServiceUser	Authenticated user.	String
SessionId	Numeric identification of the session.	Integer
SessionToken	Contains the value of DataSourceSessionToken after evaluation of placeholders and the session token for this session after placeholder replacement.	Integer
Statement	SQL statement.	String
StatementId	Numeric identification of the session.	Integer
SwitchCntR	Thread switches due to DataSourceMaxRPC.	Integer
SwitchCntT	Thread switches due to DataSourceTimeOut.	Integer
Timestamp	Timestamp when event occurred.	String
TxnCommit	Number of transactions committed.	Integer
TxnPrepare	Number of transactions prepared.	Integer
TxnRollback	Number of transaction rollbacks.	Integer

Filtering Events

You can place a filter on the attributes of any event. For example, if you want to monitor and trace only sessions that are started by users, not administrators, you would write the following filter for the Session Started event:

```
${Authorization} = "user"
```

The syntax for a filter placed on an event is:

[not] [(]${*event_attribute*} *comparison_operator* *filter_value*[)] [*boolean_operator* (${*event_attribute*} *comparison_operator* *filter_value*)...]

where:

comparison_operator is one of the following comparison operators:

=	>	equals	bigger	nsmaller
!=	<=	nequals	nbigger	
<	>=	contains	smaller	

You can use the symbols or the words; they are equivalent (for example, using != is the same as using nequals).

filter_value is the value for the attribute. The value can be a string or an integer, depending on the attribute type (see Table E-2 on page 593). Strings must be quoted.

boolean_operator is one of the following operators: and, nand, or, nor, xor, nxor

For example:

```
[not] (${ServiceUser} equals "sluser")
```

```
(${ClientInfo} contains "196.72") and (${RowsFetched} >
1000)
```

```
(${Statement} contains "insert") and (${ReturnCode} != 0)
```

```
(${Statement} contains "insert" and (${ReturnCode} = -1 or
${ReturnCode} = -2)
```

NOTE: Strings must be quoted, and the comparison operator "contains" can only be used in combination with string constants.

F Internationalization, Localization, and Unicode

This appendix provides an overview of how internationalization, localization, and Unicode relate to each other. It also provides a background on Unicode use in SequeLink, and how Unicode is accommodated by Unicode ODBC drivers.

Internationalization and Localization

Software that has been designed for *internationalization* is able to manage different linguistic and cultural conventions transparently and without additional modification. The same binary copy of an application should run on any localized version of an operating system, without requiring source code changes.

Software that has been designed for *localization* includes language translation (such as text messages, icons, and buttons), cultural data (such as dates, times, and currency), and other components (such as input methods and spell checkers) for meeting regional market requirements.

Properly designed applications can accommodate a localized interface without extensive modification. The software should be designed, first, to run internationally, and, second, to accommodate the language- and cultural-specific elements of a designated locale.

Locale

A locale represents the language and cultural data chosen by the user and dynamically loaded into memory at run time. The locale settings are applied to the operating system and to subsequent application launches.

While language is a fairly straightforward item, cultural data is a little more complex. Dates, numbers, and currency are all examples of data that is formatted according to cultural expectations. Because cultural preferences are bound to a geographic area, country is an important element of locale. Together these two elements (language and country) provide a precise context in which information can be presented. Locale presents information in the language and form that is best understood and appreciated by the local user.

Language

A locale's language is specified by the ISO 639 standard. The following table lists some language codes in the standard.

Language Code	Language
en	English
nl	Dutch
fr	French
es	Spanish
zh	Chinese
ja	Japanese
vi	Vietnamese

Because language is correlated with geography, a language code might not capture all the nuances of usage in a particular area. For example, French and Canadian French may use different

phrases and terms to mean different things even though basic grammar and vocabulary are the same. Language is only one element of locale.

Country

The locale's country identifier is also specified by an ISO standard, ISO 3166, which describes valid two-letter codes for all countries. ISO 3166 defines these codes in uppercase letters. The following table lists some language codes in the standard.

Country Code	Country
US	United States
FR	France
IE	Ireland
CA	Canada
MX	Mexico

The country code provides more contextual information for a locale and affects a language's usage, word spelling, and collation rules.

Variant

A variant is an optional extension to a locale. It identifies a custom locale that is not possible to create with just language and country codes. Variants can be used by anyone to add additional context for identifying a locale. The locale en_US represents English (United States), but en_US_CA represents even more information and might identify a locale for English (California, U.S.A). Operating system or software vendors can use these variants to create more descriptive locales for their specific environments.

Unicode Character Encoding

In addition to locale, the other major component of internationalizing software is the use of the Universal Codeset, or Unicode. Most people know that Unicode is a standard encoding that can be used to support multi-lingual character sets. Unfortunately, understanding Unicode is not as simple as its name would indicate. Software developers have used a number of character encodings, from ASCII to Unicode, to solve the many problems that arise when developing software applications that can be used worldwide.

Background

Most legacy computing environments have used ASCII character encoding developed by the ANSI standards body to store and manipulate character strings inside software applications. ASCII encoding was convenient for programmers because each ASCII character could be stored as a byte. The initial version of ASCII used only 7 of the 8 bits available in a byte, which meant that software applications could use only 128 different characters. This version of ASCII could not account for European characters, and was completely inadequate for Asian characters. Using the eighth bit to extend the total range of characters to 256 added support for most European characters. Today, ASCII refers to either the 7-bit or 8-bit encoding of characters.

As the need increased for applications with additional international support, ANSI again increased the functionality of ASCII by developing an extension to accommodate multi-lingual software. The extension, known as the Double-Byte Character Set or DBCS, allowed existing applications to function without change, but provided for the use of additional characters, including complex Asian characters. With DBCS, characters map to either one byte (such as American ASCII characters) or two bytes (for example, Asian characters). The DBCS environment also

introduced the concept of an operating system code page that identified how characters would be encoded into byte sequences in a particular computing environment. DBCS encoding provides a cross-platform mechanism for building multi-lingual applications; however, using variable-width codes is not ideal.

Many developers felt that there was a better way to solve the problem. A group of leading software companies joined forces to form the Unicode Consortium. Together, they produced a new solution to building worldwide applications—Unicode. Unicode was originally designed as a fixed-width, uniform two-byte designation that could represent all modern scripts without the use of code pages. The Unicode Consortium has continued to evaluate new characters, and the current number of supported characters is over 95,200.

Although it seemed to be the perfect solution to building multi-lingual applications, Unicode started off with a significant drawback—it would have to be retrofitted into existing computing environments. To use the new paradigm, all applications would have to change. This was clearly unacceptable, and several standards-based transliterations were designed to convert two-byte fixed Unicode values into more appropriate character encodings, including, among others, UTF-8, UCS-2, and UTF-16.

UTF-8 is a standard method for transforming Unicode values into byte sequences that maintain transparency for all ASCII codes. UTF-8 is endorsed by the Unicode Consortium as a standard mechanism for transforming Unicode values and is popular for use with HTML, XML, and similar protocols. UTF-8 is, however, currently used primarily on AIX, HP-UX, Solaris, and Linux.

UCS-2 encoding is a fixed two-byte encoding sequence and is a method for transforming Unicode values into byte sequences for Microsoft Windows platforms.

UTF-16 is a superset of UCS-2, with the addition of some special characters in surrogate pairs. UTF-16 is the standard encoding for Windows 2000, Windows XP, and Windows Server 2003.

Unicode Support in Databases

Recently, database vendors have begun to support Unicode data types natively in their systems. With Unicode support, one database can hold multiple languages. For example, a large multinational corporation could store expense data in the local languages for the Japanese, U.S., English, German, and French offices in one database.

Not surprisingly, the implementation of Unicode data types varies from vendor to vendor. For example, the Microsoft SQL Server 2000 implementation of Unicode provides data in UTF-16 format, while Oracle provides Unicode data types in UTF-8 and UTF-16 format. A consistent implementation of Unicode not only depends on the operating system, but also on the database itself.

Unicode Support in ODBC

Prior to the ODBC 3.5 standard, all ODBC access to function calls and string data types was through ANSI encoding (either ASCII or DBCS). Applications and drivers were both ANSI-based.

The ODBC 3.5 standard specified that the ODBC Driver Manager (on both Windows and UNIX) be capable of mapping both Unicode function calls and string data types to ANSI encoding as transparently as possible. This meant that ODBC 3.5-compliant Unicode applications could use Unicode function calls and string data types with ANSI drivers because the Driver Manager could convert them to ANSI. Because of character limitations in ANSI, however, not all conversions are possible.

The ODBC Driver Manager version 3.5 or later, therefore, supports the following configurations:

- ANSI application with an ANSI driver
- ANSI application with a Unicode driver
- Unicode application with a Unicode driver
- Unicode application with an ANSI driver

A Unicode application can work with an ANSI driver because the Driver Manager provides limited Unicode-to-ANSI mapping. The Driver Manager makes it possible for a pre-3.5 ANSI driver to work with a Unicode application. What distinguishes a Unicode from a non-Unicode driver is the Unicode driver's capacity to interpret Unicode function calls without the intervention of the Driver Manager, as described in "For More Information" on page 604.

Unicode Support in JDBC

Multi-lingual applications can be developed on any operating system platform with JDBC using the SequeLink *for* JDBC Client to access both Unicode and non-Unicode enabled databases. Internally, Java applications use UTF-16 Unicode encoding for string data. When fetching data, the SequeLink *for* JDBC Client automatically performs the conversion from the character encoding used by the database to UTF-16. Similarly, when inserting or updating data in the database, the JDBC driver automatically converts UTF-16 encoding to the character encoding used by the database.

Unicode Support in .NET

Internally, .NET applications use UTF-16 Unicode encoding for string data. When fetching data, the SequeLink *for* .NET Client automatically performs the conversion from the character encoding used by the database to UTF-16. Similarly, when inserting or updating data in the database, the driver automatically converts UTF-16 encoding to the character encoding used by the database.

For More Information

For more information about the differences between Unicode and non-Unicode drivers, and about developing ODBC applications on UNIX that use Unicode, refer to the *SequeLink Developer's Reference*.

G Values for IANAAppCodePage Connection String Attribute

You need to set the ODBC Client IANAAppCodePage connection string attribute (see "Connecting Using a Connection String" on page 195) if your application is not Unicode-enabled and/or if your database character set is not Unicode. The value you specify must match the database character encoding and the system locale. See Appendix F "Internationalization, Localization, and Unicode" on page 597 for information about Unicode.

To determine the correct numeric value (the MIBenum value) for the IANAAppCodePage connection string attribute, use the ivcheckcp utility (see "Using the ivcheckcp Utility" on page 606).

Alternatively, do the following:

1. Determine the code page of your database.
2. Determine the MIBenum value that corresponds to your database code page. You can check the value at:

   http://www.iana.org/assignments/character-sets

   On this web page, search for the name of your database code page. This name will be listed as an alias or the name of a character set and will have a MIBenum value associated with it.
3. Check Table G-1 to make sure that the MIBenum value you looked up on the IANA Web page is supported by ODBC Client. If the value is not listed, contact SupportLink to request that support for that value be added.

The appendix contains the following sections:

Using the ivcheckcp Utility

The ivcheckcp utility, located in the *installdir*/tools subdirectory, checks the codepage of your system and displays the locale, code page, iananame and IANAAppCodePage. The IANAAppCodePage number is the transliteration configuration parameter for the DataDirect Driver Manager on Linux/UNIX.

Check Table G-1 in this appendix to make sure that the IANAAppCodePage value you looked up on the IANA Web page is supported by the ODBC Client. If the value is not listed, contact SupportLink to request that support for that value be added.

The syntax of this command is

```
ivcheckcp <locale>
```

Example 1 (on AIX)

```
$ ./ivcheckcp en_US.ISO8859-1
                  locale         codepage       iananame IANAAppCodePage
         en_US.ISO8859-1        ISO8859-1       ISO-8859-1 4
```

Example 2 (on AIX)

```
$ ./ivcheckcp `locale -a`
                 locale         codepage       iananame IANAAppCodePage
                      C        ISO8859-1       ISO-8859-1 4
                  POSIX        ISO8859-1       ISO-8859-1 4
                  en_US        ISO8859-1       ISO-8859-1 4
          en_US.8859-15       ISO8859-15       ISO-8859-15 111
       en_US.ISO8859-1         ISO8859-1       ISO-8859-1 4
```

Example 3 (on z/OS)

```
                 locale         codepage         iananame IANAAppCodePage
       en_US.ISO8859-1          IBM-1047          IBM1047 2102
```

IANAAppCodePageValues

Table G-1 lists supported values, along with a description, for the IANAAppCodePage connection string attribute at the time of this publication. Support for additional values may have been added since publication time; therefore, for up-to-date values, go to:

http://www.datadirect.com/support/troubleshooting/su-faq-iana/index.ssp

Table G-1. IANAAppCodePage Values

Value (MIBenum)	Description
3	US_ASCII
4	ISO_8859_1
5	ISO_8859_2
6	ISO_8859_3
8	ISO_8859_5
9	ISO_8859_6

Table G-1. IANAAppCodePage Values *(cont.)*

Value (MIBenum)	Description
10	ISO_8859_7
11	ISO_8859_8
12	ISO_8859_9
17	Shift_JIS
18	EUC_JP
38	EUC_KR
106	UTF-8
109	ISO_8859_13
111	ISO_8859_15
113	GBK
2004	HP_ROMAN8
2009	IBM850
2025	GB2312
2026	Big5
2084	KOI8_R
2088	KOI8_U
2251	WINDOWS_1251
2252	WINDOWS_1252
2258	WINDOWS_1258
2259	TIS_620
10001	IBM-856
10003	IBM-921
10004	IBM-922
10012	IBM-943
10024	IBM-1046
10030	IBM-1124

Supported Code Pages on Solaris

Table G-2 lists supported code pages, IANA character set name and IANAAppCodePage for the SequeLink Client 6.0 or higher for ODBC on the Solaris platform.

Table G-2. Code Pages Supported on Solaris

Code page	IANA Character Set Name	IANAAppCodePage
646	US_ASCII	3
ISO8859-1	ISO_8859_1	4
ISO8859-2	ISO_8859_2	5
ISO8859-5	ISO_8859_5	8
ISO8859-6	ISO_8859_6	9
ISO8859-7	ISO_8859_7	10
ISO8859-8	ISO_8859_8	11
ISO8859-9	ISO_8859_9	12
PCK	Shift_JIS	17
eucJP	EUC_JP	18
ISO8859-13	ISO_8859_13	109
ISO8859-15	ISO_8859_15	111
GBK	GBK	113
UTF-8	UTF-8	106
gb2312	GB2312	2025
BIG5	Big5	2026
KOI8-R	KOI8_R	2084
TIS620.2533	TIS_620	2259

Supported Code Pages on HP

Table G-3 lists supported code pages, IANA character set name and IANAAppCodePage for the SequeLink Client 6.0 or higher for ODBC on the HP-UX platform.

Table G-3. Code Pages Supported on HP-UX

Code page	IANA Character Set Name	IANAAppCodePage
iso88591	ISO_8859_1	4
iso88592	ISO_8859_2	5
iso88595	ISO_8859_5	8
iso88596	ISO_8859_6	9
iso88597	ISO_8859_7	10
iso88598	ISO_8859_8	11
iso88599	ISO_8859_9	12
utf8	UTF-8	106
big5	Big5	2026
tis620	TIS_620	2259
roman8	HP_ROMAN8	2004

Supported Code Pages on AIX

Table G-4 lists supported code pages, IANA character set name and IANAAppCodePage for the SequeLink Client 6.0 or higher for ODBC on the AIX platform.

Table G-4. Code Pages Supported on AIX

Code Page	IANA Character Set Name	IANAAppCodePage
kIANACS_MIB_US_ ASCII	US_ASCII	3
ISO8859-1	ISO_8859_1	4
ISO8859-2	ISO_8859_2	5
ISO8859-5	ISO_8859_5	8
ISO8859-6	ISO_8859_6	9
ISO8859-7	ISO_8859_7	10
ISO8859-8	ISO_8859_8	11
ISO8859-9	ISO_8859_9	12
IBM-eucJP	EUC_JP	18
IBM-eucKR	EUC_KR	38
ISO8859-15	ISO_8859_15	111
GBK	GBK	113
UTF-8	UTF-8	106
IBM-850	IBM-850	2009
big5	Big5	2026
IBM-1252	WINDOWS_1252	2252
TIS-620	TIS_620	2259
IBM-856	IBM-856	10001
IBM-921	IBM-921	10003
IBM-922	IBM-922	10004
IBM-943	IBM-943	10012
IBM-1046	IBM-1046	10024
IBM-1124	IBM-1124	10030

Supported Code Pages on Linux

Table G-5 lists supported code pages, IANA character set name and IANAAppCodePage for the SequeLink Client 6.0 or higher for ODBC on the Linux platform.

Table G-5. Code Pages Supported on Linux

Code Page	IANA Character Set Name	IANAAppCodePage
ANSI_X3.4-1968	US_ASCII	3
ISO-8859-1	ISO_8859_1	4
ISO-8859-2	ISO_8859_2	5
ISO-8859-3	ISO_8859_3	6
ISO-8859-5	ISO_8859_5	8
ISO-8859-5	ISO_8859_6	9
ISO-8859-7	ISO_8859_7	10
ISO-8859-8	ISO_8859_8	11
ISO-8859-9	ISO_8859_9	12
EUC-JP	EUC_JP	18
EUC-KR	EUC_KR	38
ISO-8859-13	ISO_8859_13	109
ISO-8859-15	ISO_8859_15	111
GBK	GBK	113
GB18030	GB18030	114
UTF-8	UTF-8	106
GB2312	2025 GB2312	2025
BIG5	Big5	2026
KOI8-R	KOI8_R	2084
KOI8-U	KOI8-U	2088
CP1251	WINDOWS_1251	2251
TIS-620	TIS_620	2259

Supported Code Pages on Windows

Table G-6 lists supported code pages, code page name, MIBenum value, and IANA character set name for the SequeLink Client 6.0 or higher for ODBC on the Windows platform.

Table G-6. Code Pages Supported on Windows

Codepage	Code Page Name	MIBenum	IANA Character Set Name
932	Japanese Shift-JIS	17	Shift_JIS
936	Simplified Chinese GBK	113	GBK
949	Korean	10015	WINDOWS_949
1250	Central Europe	2250	WINDOWS_1251
1250	Central Europe	2250	WINDOWS_1251
1251	Cyrillic	2251	WINDOWS_1251
1252	Latin I	2252	WINDOWS_1252
1253	Greek	2253	WINDOWS_1253
1254	Turkish	2254	WINDOWS_1254
1255	Hebrew	2255	WINDOWS_1255
1256	Arabic	2256	WINDOWS_1256
1257	Baltic	2257	WINDOWS_1257
1258	Vietnam	2258	WINDOWS_1258

Index

Symbols

Numerics

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

X

Z