



SHADOW MAINFRAME ADAPTER CLIENT FOR DB2

SHADOW MAINFRAME ADAPTER CLIENT INSTALLATION AND ADMINISTRATION

POWERED BY
SHADOW

Date: January, 2004

This document is published by the NEON Systems, Inc. Technical Publications Department and applies to Shadow Mainframe Adapter Client for DB2.

Copyright © 1994-2003 NEON Systems, Inc. All rights reserved. Printed in the U.S.A.

Licensee is granted permission to make a limited number of copies of the documentation for its internal business purposes only. All such copies shall bear all copyright, trade secret, trademark and any other intellectual property notices on the original copies. This limited right to reproduce for internal purposes only is not transferable. Furthermore, this limited right DOES NOT include any license to distribute, modify, display or make derivative works from the Copyrighted materials.

NEON, Shadow, Shadow Direct, and Enterprise Direct are registered trademarks, and the NEON logo, Shadow Activity Monitor, Shadow Advanced Controls, Shadow Advanced Scalability, Shadow AutoHTML, Shadow Mainframe Adapter Client, Shadow Enterprise Auditing, Shadow Enterprise Direct, Shadow Enterprise Transactions, Shadow Event Facility, Shadow Enterprise Transactions, Shadow Interface, Shadow JDBC Adapter, Shadow MDI Replacement Module, Shadow REXX/Tools, Shadow Mainframe Adapter Server, Shadow SSL Support Module, Shadow Support Module, Shadow Web Interface, and Shadow Web Mainframe Adapter Server are trademarks of NEON Systems, Inc. in the USA and in other select countries.

The symbols ® and ™ denote USA trademark rights.

All other trademarks are the property of their respective owners.

Throughout this publication, NEON Systems, Inc. is also, for convenience, referred to as "NEON." The Reader should not presume that such use of NEON conflicts with the use of NEON as a registered trademark associated with certain products of NEON Systems, Inc.

This software/documentation contains proprietary information of NEON Systems, Inc.; it is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited.

If this software/documentation is delivered to any U.S. Government Agency, then it is delivered with Restricted Rights and the following legend is applicable:

Restricted Rights Legend

Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in FAR Section 52.227-14 (June 1987) Alt. III(g)(3)(June 1987), FAR Section 52.227-19 (June 1987), or sub-clause (c)(1)(ii) of Rights in Technical Data and Computer Software clause at DFARS 252.227-7013, as applicable.
Contractor is NEON Systems, Inc. 14100 Southwest Freeway, Suite 500, Sugar Land, Texas 77478.

NEON Systems, Inc. does not warrant that this document is error-free. The information in this document is subject to change without notice and does not represent a commitment on the part of NEON Systems, Inc. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose without the express written permission of an authorized representative of NEON Systems, Inc.

Address inquiries to:

NEON Systems, Inc.

14100 SW Freeway, Suite 500
Sugar Land, Texas 77478

World Wide Web: <http://www.neonsys.com>

Phone: 1-800-505-6366
(281) 491-4200 (Corporate Sales, Customer Support)
Fax: (281) 242-3880



Contents

Part I: Introduction

Chapter 1: Introduction	1-1
Overview	1-1
Shadow Mainframe Adapter Client for DB2	1-1

Part II: Planning and Installation

Chapter 2: Planning the Shadow Mainframe Adapter Client Installation	2-1
Prerequisites	2-1
Supported Operating Systems	2-2
Support for Multi-Threaded Applications	2-2
Chapter 3: Installing Shadow Mainframe Adapter Client	3-1
Installation Files	3-1
Installing Shadow Mainframe Adapter Client on Windows	3-2
Installation Steps	3-2
Additional Installation Options	3-9
Installing Shadow Mainframe Adapter Client on UNIX	3-12
Installation Steps	3-12
Additional Installation Options	3-22
Common UNIX Errors	3-25

Part III: Administration

Chapter 4: Shadow Mainframe Adapter Client: Configuration	4-1
Connection Configuration	4-1
Step 1: Configure the Data Source	4-1
Step 2: Configure the Java Connection	4-16
Additional Configuration and Administration	4-18
Using Optimized Fetch	4-18
DBCS Support	4-19
Determining Shadow Mainframe Adapter Client Version Information	4-20
Chapter 5: jDemo	5-1
Overview	5-1
Starting jDemo	5-2
Supported Features	5-3
Establishing a Connection	5-3

Viewing Data Source Information	5-6
Executing SQL	5-13
Closing jDemo	5-14
Chapter 6: Shadow Mainframe Adapter Client Trace Facility	6-1
Overview	6-1
Shadow Mainframe Adapter Client Trace Facility for Windows	6-2
Shadow Mainframe Adapter Client Trace.	6-2
Shadow Mainframe Adapter Client Dynamic JDBC Trace	6-4
Shadow Mainframe Adapter Client Trace Facility for UNIX	6-6
Shadow Mainframe Adapter Client Trace.	6-7
Shadow Mainframe Adapter Client Dynamic JDBC Trace	6-8
Shadow Mainframe Adapter Client Trace Facility Options	6-11
Severity Level Trigger.	6-11
Tracing Options	6-12
Chapter 7: Accessing DB2	7-1
Invoking a Stored Procedure	7-1
Returning Results Sets	7-2
Coding Cursors in Return Result Sets	7-2
Retrieving Column Names	7-3
Using SET CURRENT PATH.	7-3
Default RPCDEFAULTSCHEMA Value	7-3
Setting the RPCDEFAULTSCHEMA Value	7-4
Using the SET CURRENT PATH Statement	7-4
Troubleshooting DB2 Stored Procedures.	7-5
Samples	7-6
Using the Sample Programs	7-6
Part IV: Appendices	
Appendix A: Shadow Mainframe Adapter Client Keywords	A-1
Setting a Keyword	A-1
Keyword Descriptions	A-1
Setting Keywords to Benefit Performance.	A-32
Appendix B: National Language Support	B-1
Languages Supported.	B-1
Setting the Language	B-2
Appendix C: JDBC Compliance	C-1
JDBC Compliance	C-1

About this Publication

This book contains user documentation for Shadow Mainframe Adapter Client, the Mainframe Adapter Client component of the Shadow Connect product.

How this Publication Is Organized

This book contains the following chapters:

Part I: Introduction

- Chapter 1, “Introduction,” provides a brief overview of the Shadow Connect product.

Part II: Planning and Installation

- Chapter 2, “Planning the Shadow Mainframe Adapter Client Installation,” describes the prerequisites for installing Shadow Mainframe Adapter Client.
- Chapter 3, “Installing Shadow Mainframe Adapter Client,” details the installation of Shadow Mainframe Adapter Client on all of the supported operating platforms. This chapter includes step-by-step installation instructions as well as installation considerations.

Part III: Administration

- Chapter 4, “Shadow Mainframe Adapter Client: Configuration,” covers general administrative procedures for configuring Shadow Mainframe Adapter Client.
- Chapter 5, “jDemo,” describes the functionality provided by the jDemo program, a graphical user interface (GUI) to access data structures using Shadow Mainframe Adapter Client.
- Chapter 6, “Shadow Mainframe Adapter Client Trace Facility,” provides information on using the Shadow Mainframe Adapter Client Trace Facility to capture traces to solve problems that can occur during installation and use of Shadow Mainframe Adapter Client.
- Chapter 7, “Accessing DB2,” covers the use of Shadow Mainframe Adapter Client for accessing DB2, including details about using DB2 stored procedures.

Part IV: Appendices

- Appendix A, “Shadow Mainframe Adapter Client Keywords,” provides details about each Shadow Mainframe Adapter Client keyword and shows you how to change a keyword setting.

-
- Appendix B, “National Language Support,” describes the National Language Support (NLS) of Shadow Mainframe Adapter Client.
 - Appendix C, “JDBC Compliance,” provides information about JDBC compliance.

Reader’s Comments

Please e-mail any comments or questions you have about our documentation to support@neonsys.com.

Thank you!

Introduction

This chapter gives a general introduction to the Shadow Connect product, including details about the Shadow Mainframe Adapter Client for DB2 component.

Topics include the following:

- Overview
 - Shadow Mainframe Adapter Client for DB2

Overview

Shadow Connect is an efficient, easy-to-use, and flexible solution for integrating mainframe data sources and transaction environments to Mainframe Adapter Client/Mainframe Adapter Server and n-tier environments. The unique Shadow Connect architecture provides maximum flexibility with minimal impact on CPU cycles.

Shadow Mainframe Adapter Client for DB2

The Shadow Connect product offers various connectivity options, including Shadow Mainframe Adapter Client for DB2. With Shadow Mainframe Adapter Client for DB2, any JDBC enabled application can use standard JDBC facilities to access DB2 data and transactions.

Shadow Mainframe Adapter Client for DB2 consists of the following components:

- Shadow Mainframe Adapter Server
- Shadow Mainframe Adapter Client
- Shadow Interface™ for DB2

Shadow Mainframe Adapter Server

The Shadow Mainframe Adapter Server component, which resides on the mainframe, offers the following benefits:

- Provides native access to ADABAS, CICS, DB2, IMS/DB, IMS/TM, Natural, and VSAM from a single tool.
- Eliminates the need for a mid-tier gateway.
- Installs in less than one day.
- Incorporates centralized online monitoring, control, and diagnostic capabilities.

Shadow Mainframe Adapter Client

The Shadow Mainframe Adapter Client component, which consists of Shadow Mainframe Adapter Client, is a connector that enables Java applications to integrate z/OS data and transactional sources through the JDBC API. The Shadow Mainframe Adapter Client provide applications with transparent access to z/OS data and transactions through standard interfaces, and performs data and SQL dialect conversations, data compression, and network optimization in conjunction with the Shadow Mainframe Adapter Server.

The Shadow Mainframe Adapter Client component is installed on the Mainframe Adapter Client side or the mid-tier Mainframe Adapter Server by using a standard installation process; it is quick and easy to install and to use.

Features

Features and benefits include the following:

- **Robust Java Support.** Provides robust supports of JVM 1.2 (J2EE) and Java servlets and is JDBC 2.0 compliant. This support includes taking advantage of Java capabilities including multi-threading, connection pooling, and batch updates.
- **Multiple Platforms.** Runs on a growing range of platforms including HP-UX, IBM AIX, Linux, Linux/390, Sun Solaris, and Windows.
- **Multiple Communications Protocols.** Operates over multiple communications protocols such as TCP/IP, SNA, or MQ Series.
- **Security Features.** Integrates tightly with enterprise security packages such as RACF, ACF2, and Top Secret. In addition, security is further improved by directly authenticating userids. Finally, reliability and resilience are improved by smoothly handling error messages coming from Security Access Facility (SAF), the z/OS security interface.
- **Two-Phase Commit Support.** Enables robust transactions by supporting two-phase commit.
- **Customizations and Optimization.** Delivers high performance with numerous field-proven customizations and optimizations. For example, data and SQL dialect conversations, data compression, and network optimization is performed in conjunction with the Shadow Mainframe Adapter Server.
- **Diagnostics and Control Features.** Simplifies support of the solution with end-to-end diagnostics. In addition, administration is made easy with comprehensive monitoring and control.
- **DBCS Support.** Permits use of multiple language character sets via DBCS support.

Shadow Interface for DB2

The Shadow Mainframe Adapter Client for DB2 offers access to DB2 data and transactions, providing maximum performance for organizations needing to integrate DB2 data with distributed or Web applications without sacrificing flexibility, reliability, or security. When used with Shadow Connect, the Shadow Interface for DB2 provides a superior solution for integrating DB2 resources in any environment.

Regardless of how the data is initially represented, the Shadow Interface for DB2 can integrate DB2 data and stored procedures without custom coding. In addition, one Shadow Mainframe Adapter Server can access many DB2 subsystems.

Planning and Installation

CHAPTER 2:

Planning the Shadow Mainframe Adapter Client Installation

This chapter discusses the planning considerations required for installing Shadow Mainframe Adapter Client, part of the Mainframe Adapter Client component of the Shadow Connect product.

Topics include the following:

- Prerequisites
- Supported Operating Systems
- Support for Multi-Threaded Applications

Prerequisites

Before you install Shadow Mainframe Adapter Client on your machine, you must meet the following prerequisites:

- **All Users:**
 - A color depth on the target system of at least 256 colors.
 - JVM 1.2 or higher.



Note:

A simple way to verify the availability of JVM 1.2 or higher is to execute the following from within the path:

```
$ java -version
```

- **Windows Users:**
 - Approximately 30 MB of available disk space for installation.
 - A minimum of 16 MB of RAM.
 - A 386 or higher processor.
- **UNIX Users:** Up to 50 MB of available disk space for installation.
- **Linux Users:** glibc v2.1.1 or above.
- **LU 6.2 Users:** One of the following:
 - Wall Data's Rumba for Office
 - Wall Data's Rumba for Mainframe
 - Attachmate's EXTRA! v4.1 (or the 4.0 APPC upgrade) for Windows
 - IBM Network Services/DOS version 1.0
 - Any other LU 6.2 stack that supports APPC

Supported Operating Systems

Shadow Mainframe Adapter Client currently support the following operating platforms:

- AIX
- HP-UX
- Linux
- Linux/390
- Solaris
- Windows

Support for Multi-Threaded Applications

JDBC users require a “thread-safe” environment. Shadow Mainframe Adapter Client is enhanced to work with multi-threaded applications, which include many popular Mainframe Adapter Server products that use concurrent threads or thread-based connection pooling models to invoke operations. Shadow Mainframe Adapter Client is fully thread-safe and does not require thread affinity. No additional parameter or code changes need to be made.

CHAPTER 3: *Installing Shadow Mainframe Adapter Client*

This chapter describes the procedure for installing Shadow Mainframe Adapter Client, part of the Mainframe Adapter Client component of the Shadow Connect product.

Topics include the following:

- Installation Files
- Installing Shadow Mainframe Adapter Client on Windows
- Installing Shadow Mainframe Adapter Client on UNIX



Note:

Before installing Shadow Mainframe Adapter Client, ensure that the prerequisites have been met. For more information, see Chapter 2, “Planning the Shadow Mainframe Adapter Client Installation,” of this guide.

Installation Files

You can install Shadow Mainframe Adapter Client with installation files from the Shadow Connect CD. This CD should have been included in the package you received, containing the Shadow Connect product. When executing the installation file, you will need to pick the appropriate Shadow Mainframe Adapter Client file for your operating system, based on Table 3–1.

Table 3–1. Shadow Mainframe Adapter Client Files

Operating System	Installation File*
AIX 4.3 or higher	IBMShadowMainframe Adapter Client_3_9_nnn_aix.bin
HP-UX 11.0 or higher	IBMShadowMainframe Adapter Client_3_9_nnn_hp.bin
Linux	IBMShadowMainframe Adapter Client_3_9_nnn_linux.bin
Linux/390	IBMShadowMainframe Adapter Client_3_9_nnn_linux390.bin
Solaris 2.6 or higher	IBMShadowMainframe Adapter Client_3_9_nnn_solaris.bin

Table 3–1. Shadow Mainframe Adapter Client Files (Continued)

Operating System	Installation File*
32-bit Windows (XP, 2000, NT, 98, or 95)	IBMShadowMainframe Adapter Client_3_9_nnn_win32.exe

* Note: The “nnn” represents the Shadow Mainframe Adapter Client build number.

Installing Shadow Mainframe Adapter Client on Windows

Installing the Shadow Mainframe Adapter Client on Windows involves the following:

- Installation steps.
- Additional installation options.

Installation Steps

To install Shadow Mainframe Adapter Client on Windows, do the following:

1. Install Shadow Mainframe Adapter Client.
2. Verify the search paths.
3. Verify the installation.

Step 1: Install Shadow Mainframe Adapter Client



Note:

The Shadow Mainframe Adapter Client installation can be customized by changing the default settings, as detailed in “Additional Installation Options” on page 3-9.

To install Shadow Mainframe Adapter Client on any 32-bit Windows operating system, do the following:

1. Locate the Shadow Mainframe Adapter Client executable installation file, as listed in Table 3–1 on page 3-1.
2. Launch the Shadow Mainframe Adapter Client executable installation file. An introductory dialog box offering language choices is displayed, as shown in Figure 3–1.



Figure 3–1. Introductory Dialog Box -- Choosing a Language

3. From the drop-down menu, select the language.
4. Click **OK**. After the system prepares the files, the **Introduction** dialog box is displayed, as shown in Figure 3–2. This dialog box offers general information for proceeding with the installation.

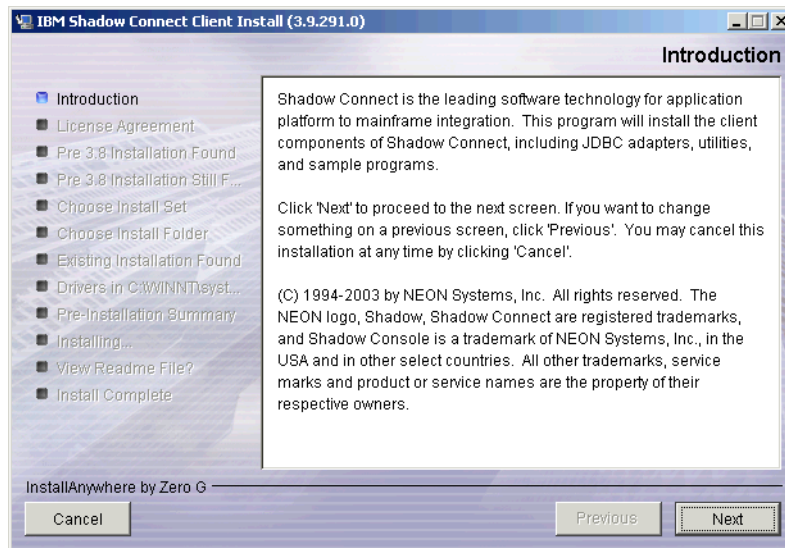


Figure 3–2. Introduction Dialog Box

5. Click **Next**. The **License Agreement** dialog box is displayed, as shown in Figure 3–3.

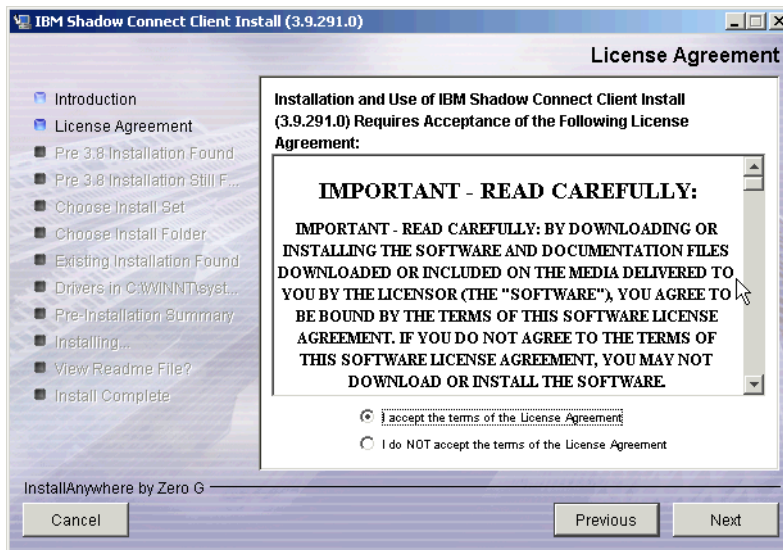


Figure 3–3. License Agreement Dialog Box

6. After reviewing the license agreement, to accept the terms and continue with the installation, click the appropriate radio button.



Note:

If you do not accept the terms of the license agreement, you will not be able to continue with the installation.

7. Click **Next**. The **Choose Install Set** dialog box is displayed, as shown in Figure 3–4.




Figure 3–4. Choose Install Set Dialog Box

8. Choose the appropriate installation option. The installation options are as follows:
 - **Typical:** The typical Shadow Mainframe Adapter Client installation offers a full installation with the setting of user (local) environment variables.
 - **Administrative:** The administrative Shadow Mainframe Adapter Client installation offers a full installation with the setting of system (global) environment variables.
- ▶ **Note:**
The selected option will be displayed in blue text.
9. Click **Next**. The **Choose Install Folder** dialog box is displayed, as shown in Figure 3–5.

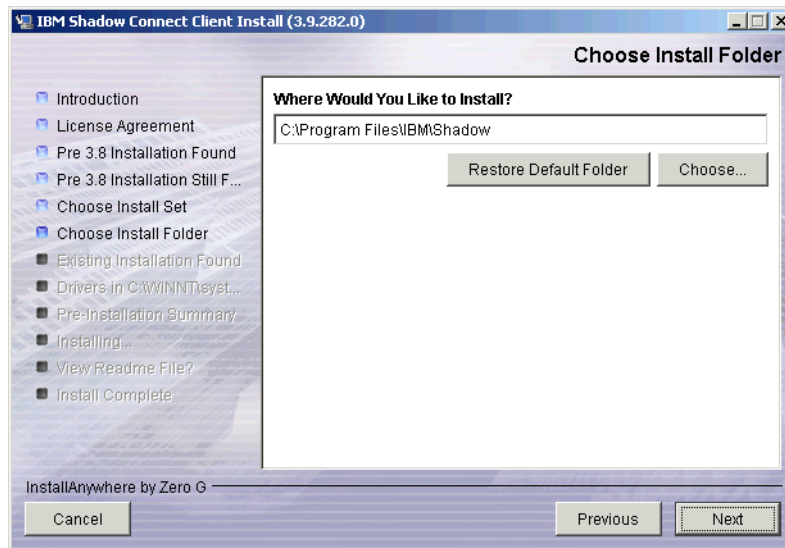


Figure 3–5. Choose Install Folder Dialog Box

10. Choose the destination for the installed files. You may do either of the following:
 - Accept the default destination folder for the installation (C:\Program Files\IBM\Shadow).
 - Select a different location using the **Choose** button.
11. Click **Next**. The installation will search for previously installed versions of Shadow Mainframe Adapter Client.
 - If a previous Shadow Mainframe Adapter Client installation is not found, then the **Pre-Installation Summary** dialog box will be displayed, as shown in Figure 3–7 on page 3-7. See step 12 on page 3-7.

- If a previous Shadow JDBC installation is found, then the **Existing Installation Found** dialog box will be displayed, as shown in Figure 3–6.




Figure 3–6. Existing Installation Found Dialog Box



Note:

Users are only allowed one installation location. For existing installations, users will not be given the option to select the destination for the installed files; the files will be installed to the same location as the existing installation by default.

- Select one of the following options:
 - Completely overwrite the existing Shadow Mainframe Adapter Client installation, regardless of which installation is more current.
 - Upgrade the existing files only if they are older than the ones about to be installed.
- Click **Next**. The **Pre-Installation Summary** dialog box will be displayed, as shown in Figure 3–7.

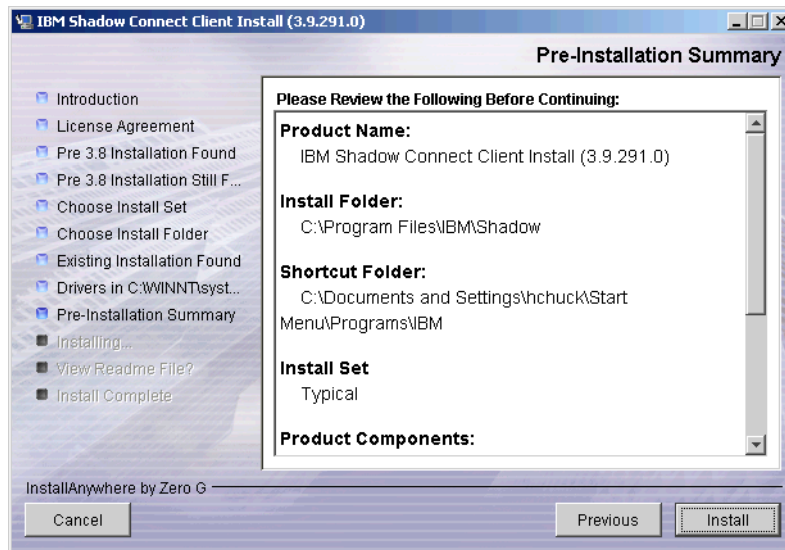


Figure 3–7. Pre-Installation Summary Dialog Box

12. Click **Install**. A screen is displayed that shows the progress of the installation for Shadow Mainframe Adapter Client. After the files have been installed, the **View Readme File** dialog box will be displayed, as shown in Figure 3–8.

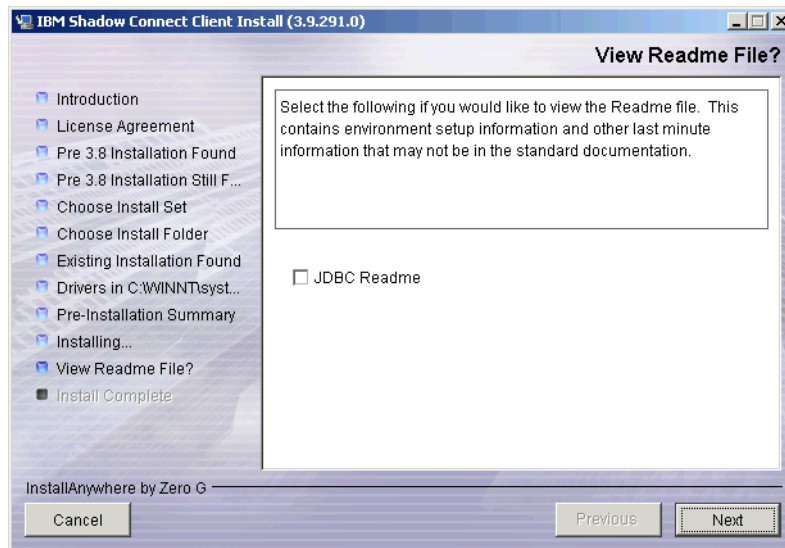


Figure 3–8. View Readme File Dialog Box

13. Review the README file, if desired, and then click **Next**. The installation process is completed, as indicated by the **Install Complete** dialog box shown in Figure 3–9.

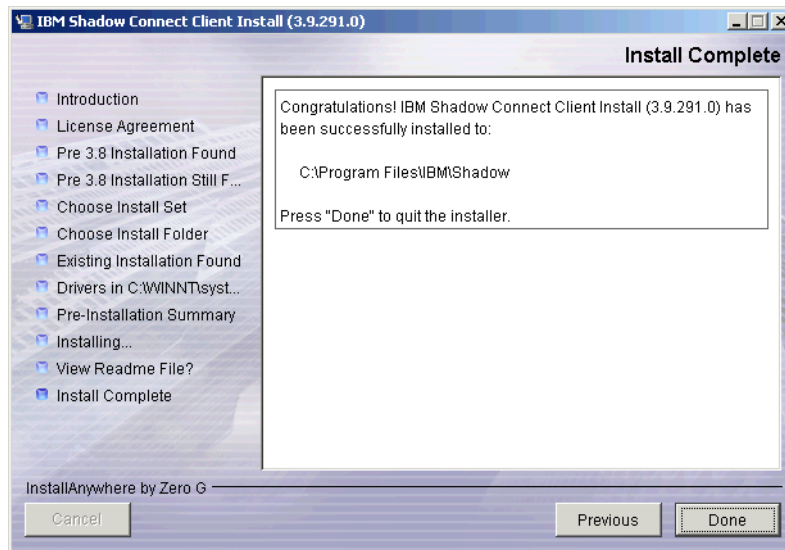


Figure 3–9. Install Complete Dialog Box

14. Click **Done** to exit.



Note:

On Windows 98, the computer must be restarted after installing Shadow Mainframe Adapter Client.

Step 2: Verify the Search Paths

Verify that the CLASSPATH and PATH environment variables are set correctly for the Shadow Mainframe Adapter Client files. The Shadow Mainframe Adapter Client files can be located in any directory as long as the environment variables are set so that they can be found.



Note:

If you chose a typical installation, user environment variables should be defined. If you chose an administrative installation, system level environment variables should be defined. The type of installation was selected in step 8 on page 3-5.

Verifying the CLASSPATH Environment Variable

Ensure that the CLASSPATH environment variable points to the jar file you wish to use. Assuming a default installation, the jar file will be copied to the following location:

`C:\Program Files\IBM\Shadow\jdbc`

The jar file is named as follows:

- **Debug Version:** `scjd12ts.jar`
- **Non-Debug Version:** `scjd12.jar`

Example

If the jar files are installed under `C:\Program Files\IBM\Shadow\jdbc`, to use the non-debug version of Shadow Mainframe Adapter Client, the `CLASSPATH` environment variable should include the following value:

```
C:\Program Files\IBM\Shadow\jdbc\scjd12.jar
```

Verifying the PATH Environment Variable

The `PATH` environment variable should point to the location of the `bin` directory.

Example

Assuming a default installation, the `PATH` environment variable should point to the following location:

```
C:\Program Files\IBM\Shadow\bin
```

Step 3: Verify the Installation

To verify the installation of Shadow Mainframe Adapter Client, you can use the `jDemo` program, a graphical interface demo program, that is installed with Shadow Mainframe Adapter Client.



Doc Reference:

The `jDemo` program is documented in Chapter 5, “`jDemo`,” of this guide.

Additional Installation Options

The Shadow Mainframe Adapter Client installation can be configured and customized by specifying the installer properties. The installer properties are configured within a user-created text file called `installer.properties`. From this file, you can configure default installation properties, such as the installation directory location and the components to be installed, as well as specify the installation mode.

Configuring the Installation Properties

To configure the installation properties, simply do the following:

1. Within the directory where the actual Shadow Mainframe Adapter Client installation file is located, create a text file called `installer.properties`.

**Note:**

To use an `installer.properties` file that is located in a different directory, see “Command Line Installation Option” on page 3-11.

2. Specify the installation properties within the `installer.properties` file. The details of the properties that can be set are described in Table 3–2. An example `installer.properties` file is shown in Figure 3–10.
3. Run the installation program. With the `installer.properties` file located in the same directory, the installation program will automatically use the `installer.properties` file to configure the installation settings.

Table 3–2. Windows Installer Properties

Installer Property	Description
USER_INSTALL_DIR	<p>This property specifies the destination for the installed files. You must provide a complete, absolute path to the location.</p> <p><i>Note:</i> Use forward slashes for the path—even for Windows! See the example in Figure 3–10.</p>
INSTALLER_UI	<p>This property sets the installation mode. Possible values are as follows:</p> <ul style="list-style-type: none"> • GUI: Executes the installation in GUI mode, as illustrated in the steps shown in “Step 1: Install Shadow Mainframe Adapter Client” on page 3-2. • SILENT: Executes the installation in silent mode. In silent mode, the installation GUI is not invoked; thus, no user interaction is required or accepted. If non-default installation settings are required, they must be configured manually in the <code>installer.properties</code> file. <p><i>Note:</i> When Shadow Mainframe Adapter Client is installed in silent mode, it will also be uninstalled in silent mode.</p>
CHOSEN_INSTALL_SET	<p>This property chooses the installation set option. Possible values are as follows:</p> <ul style="list-style-type: none"> • TYPICAL • ADMINISTRATIVE <p>For details about the various installation options, see step 8 on page 3-5.</p>
OVERWRITE_OR_UPGRADE	<p>This property determines how the installation should handle the existence of another Shadow Mainframe Adapter Client install, as follows:</p> <ul style="list-style-type: none"> • OVERWRITE EXISTING INSTALL • UPGRADE OLDER FILES <p>For more information about these options, see step a on page 3-6.</p>

```

# This is a sample installer.properties file that can be used with the
# Shadow Mainframe Adapter Client installation.

#####
# Use the USER_INSTALL_DIR setting to control where Shadow Mainframe Adapter
Client will
# be installed. Provide a complete, absolute path to the location.
# NOTE: Use forward slashes on Windows!
#####
USER_INSTALL_DIR=C:/Program Files/IBM/Shadow

#####
# Use the INSTALLER_UI setting to control whether the installer launches a
# GUI, a text interface, or runs silently. For a GUI, use "GUI" as the value.
# For the text interface, use "CONSOLE" as the value. For a silent install,
# use "SILENT" as the value. This sample is configured to run silently, as
# that is typical when running with a properties file.
#####
INSTALLER_UI=SILENT

#####
# The CHOSEN_INSTALL_SET setting allows you to choose between the high level
# install sets that are available in the install. They are "TYPICAL" and
# "ADMINISTRATIVE".
#
# TYPICAL: Includes Shadow Mainframe Adapter Client, samples, and user env
variables.
# ADMINISTRATIVE: Includes Shadow Mainframe Adapter Client, samples, and
system env
#             variables.
#####
CHOSEN_INSTALL_SET=TYPICAL

#####
# If an existing installation is found on your system, the Windows installer
# will ask you if you would like to overwrite the existing install or upgrade
# the existing installation. Use the OVERWRITE_OR_UPGRADE setting to make
# that choice. Valid values are "OVERWRITE EXISTING INSTALL" or "UPGRADE

```

Figure 3–10. Example Installer Properties File -- Windows

Command Line Installation Option

By executing the installation program via the command line, you can force the installation to use an `installer.properties` file that is located in a different directory. Use the “-f” command line option to accomplish this, as follows:

```
C:\>install_path\install_file -f prop_path\installer.properties
```

**Note:**

Note the “escaped” back-slashes used to specify the path for the `installer.properties` file.

Example

Assume the following:

- The installation program is located at `C:\Installs`.
- The `installer.properties` file is located at `C:\Properties`.

To run the installation with the `installer.properties` file, type the following on the command line:

```
C:\>C:\Installs\IBMShadowMainframe Adapter
Client_3_9_nnn_win32.exe -f C:\\Properties\\installer.properties
```

Installing Shadow Mainframe Adapter Client on UNIX

Installing Shadow Mainframe Adapter Client on UNIX involves the following:

- Installation steps.
- Additional installation options.
- Common UNIX errors.

Installation Steps

To install Shadow Mainframe Adapter Client on UNIX, do the following:

1. Set the file permissions.
2. Install Shadow Mainframe Adapter Client.
3. Verify the search paths.
4. (If using debugging) Set NEONTRACE log file permissions.
5. Verify the installation.

Step 1: Set the File Permissions

Before attempting to install Shadow Mainframe Adapter Client on UNIX, the file permissions for the installation file must be updated to allow execution of the file. The following command will set the appropriate permissions:

```
chmod 777 IBMShadowMainframe Adapter Client_3_9_nnn_xxx.bin
```

Where `nnn` is the build number and `xxx` is the suffix reflecting the UNIX platform (see Table 3–1 on page 3-1 for the list of file names).

Step 2: Install Shadow Mainframe Adapter Client

Considerations

The following should be considered before installing Shadow Mainframe Adapter Client:

- **Upgrade Considerations:** When performing an upgrade, ensure that all files were properly upgraded.



Doc Reference:

For details about resolving upgrade issues, see “Upgrade Issues” on page 3-19.

- **Installation Options:** Shadow Mainframe Adapter Client offers three installation methods:
 - **GUI Installation:** The GUI installation, which is the standard installation method, requires a UNIX system that has X Windows enabled. See “GUI Installation” on page 3-14.
 - **Console Installation:** The console installation, which executes via a text interface instead of the GUI, can be executed by users running the installation via a telnet connection or on any UNIX system that does not have X Windows enabled. See “Console Installation” on page 3-18.
 - **Silent Installation:** The silent installation requires no user input and does not invoke the GUI, so it can also be executed by users running the installation via a telnet connection or on any UNIX system that does not have X Windows enabled. See “Silent Installation” on page 3-19.

Additional customizable options include changing the default installation settings.



Note:

When running the UNIX installation via Windows telnet, you need to install with either the console installation or the silent installation.



Doc Reference:

For details regarding the installation options, see “Additional Installation Options” on page 3-22.

GUI Installation

To install Shadow Mainframe Adapter Client using the GUI installation on a UNIX system with X Windows enabled, do the following:

1. Ensure that your DISPLAY environment variable is set appropriately.
2. Locate the appropriate binary file, as listed in Table 3–1 on page 3-1.
3. Launch the installation file. An introductory dialog box offering language choices is displayed, as shown in Figure 3–11.



Figure 3–11. Introductory Dialog Box -- Choosing a Language

4. From the drop-down menu, select the language.
5. Click **OK**. After the system prepares the files, the **Introduction** dialog box is displayed, as shown in Figure 3–12. This dialog box offers general information for proceeding with the installation.




Figure 3–12. Introduction Dialog Box

6. Click **Next**. The **License Agreement** dialog box is displayed, as shown in Figure 3–13.




Figure 3–13. License Agreement Dialog Box

7. After reviewing the license agreement, to accept the terms and continue with the installation, click the appropriate radio button.



Note:

If you do not accept the terms of the license agreement, you will not be able to continue with the installation.

8. Click **Next**. The **Choose Install Set** dialog box is displayed, as shown in Figure 3–14.




Figure 3–14. Choose Install Set Dialog Box

9. Choose the appropriate installation option. The installation options are as follows:
 - **Typical:** The typical Shadow Mainframe Adapter Client installation offers a full installation with the setting of user (local) environment variables.
 - **Administrative:** The administrative Shadow Mainframe Adapter Client installation offers a full installation with the setting of system (global) environment variables.

▶ **Notes:**

 - The selected option will be displayed in blue text.
 - Shadow Mainframe Adapter Client installations on Linux/390 **do not** set environment variables during the installation process.

10. Click **Next**. The **Choose Install Folder** dialog box is displayed, as shown in Figure 3–15.




Figure 3–15. Choose Install Folder Dialog Box

11. Choose the destination for the installed files. You may do either of the following:
 - Accept the default destination folder for the installation (/home/IBM/Shadow).
 - Select a different location using the **Choose** button.
12. Click **Next**. The **Pre-Installation Summary** dialog box will be displayed, as shown in Figure 3–16.




Figure 3–16. Pre-Installation Summary Dialog Box

13. Click **Install**. A screen is displayed that shows the progress of the installation for Shadow Mainframe Adapter Client. After the files have been installed, the **View Readme File** dialog box will be displayed, as shown in Figure 3–17.

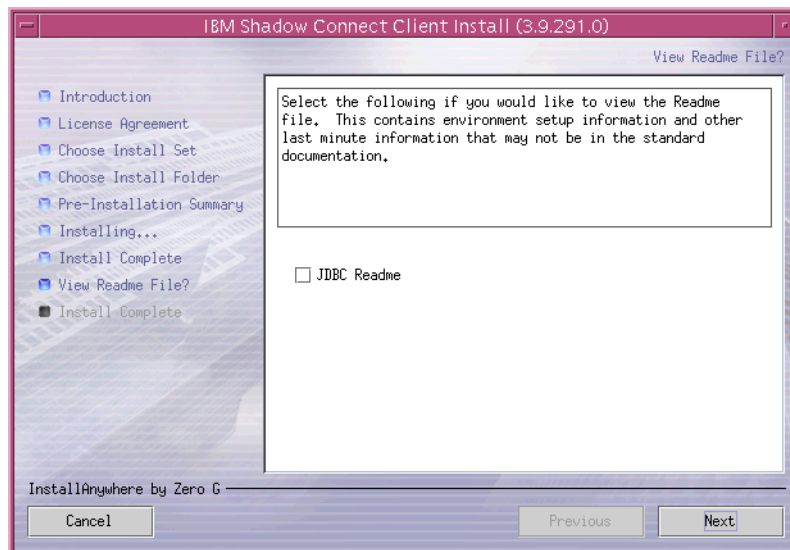


Figure 3–17. View Readme File Dialog Box

14. Review the README file, if desired, and then click **Next**. The installation process is completed, as indicated by the **Install Complete** dialog box shown in Figure 3–18.




Figure 3–18. Install Complete Dialog Box

15. Click **Done** to exit.

Console Installation

The Shadow Mainframe Adapter Client installation can also be run in console mode instead of GUI mode. You will still have the same options and functionality as the GUI version, but everything will be displayed as text.

To install in console mode, do either of the following:

- ▶ Launch the installation from the command line using the “-i console” option, as follows:

```
$ /home/file -i console
```

Where *file* is the name of the Shadow Mainframe Adapter Client installation file (see Table 3–1 on page 3-1).

Or

- ▶ Add the following line to the `installer.properties` file:

```
INSTALLER_UI=CONSOLE
```



Doc Reference:

For details regarding the `installer.properties` file, see “Additional Installation Options” on page 3-22.

Silent Installation

The Shadow Mainframe Adapter Client installation can also be run “silently.” In silent mode, the installation GUI is not invoked; thus, no user interaction is required or accepted. If non-default installation settings are required, they must be configured manually in the `installer.properties` file.



Note:

When Shadow Mainframe Adapter Client is installed in silent mode, it will also be uninstalled in silent mode.

To install in silent mode, do either of the following:

- ▶ Launch the installation from the command line using the “`-i silent`” option, as follows:

```
$ /home/file -i silent
```

Where `file` is the name of the Shadow Mainframe Adapter Client installation file (see Table 3–1 on page 3-1).

Or

- ▶ Add the following line to the `installer.properties` file:

```
INSTALLER_UI=SILENT
```



Doc Reference:

For details regarding the `installer.properties` file, see “Additional Installation Options” on page 3-22.

Upgrade Issues

When upgrading Shadow Mainframe Adapter Client on UNIX, if your installation receives errors indicating that certain files could not be upgraded, check the installation log for errors such as the following:

```
Install File: /path/file
Status: ERROR
Additional Notes: ERROR - ZeroGnd: /path/file
(Cannot open or remove a file containing a running
program.)
```

If such errors were noted, do the following:

1. Ensure that no programs are currently using any of the Shadow Mainframe Adapter Client components.

2. Manually remove any files that cannot be removed by the installation, such as `install.dir.xxx` (an NFS mount file).
3. Rerun the upgrade, checking the installation log to verify a successful installation.

Step 3: Verify the Search Paths

The search paths must be specified for both the jar file and the library file. In addition, if using Shadow Mainframe Adapter Client data sources, the search path for the `shadow.ini` configuration file must be specified. The Shadow Mainframe Adapter Client files can be located in any directory as long as the environment variables are set so that they can be found.



Note:

- If you chose a typical installation, user environment variables should be defined. If you chose an administrative installation, system level environment variables should be defined. The type of installation was selected in step 9 on page 3-16.
- Shadow Mainframe Adapter Client installations on Linux/390 **do not** set environment variables during the installation process; thus, any required environment variables must be set manually by the user.

To use Shadow Mainframe Adapter Client, verify the following search paths:

1. Verify the CLASSPATH environment variable.
2. Verify the search path for the library file.
3. Verify the SHADOW_INI environment variable.

Verifying the CLASSPATH Environment Variable

Ensure that the CLASSPATH environment variable points to the jar file you wish to use. Assuming a default installation, the jar file will be copied to the following location:

```
/home/IBM/Shadow/jdbc
```

The jar file is named as follows:

- **Debug Version:** `scjd12ts.jar`
- **Non-Debug Version:** `scjd12.jar`

Example

If the jar files are installed under `/home/IBM/Shadow/jdbc`, to use the non-debug version of Shadow Mainframe Adapter Client, the `CLASSPATH` environment variable should be set as follows:

```
CLASSPATH=/home/IBM/Shadow/jdbc/scjd12.jar:$CLASSPATH
```

Verifying the Search Path for the Library File

The jar file pointed to by the `CLASSPATH` variable requires a certain library file. Assuming a default installation, the jar file will be copied to the following location:

```
/home/IBM/Shadow/lib
```

The library file is named as follows:

- **Debug Version:** `libscjd12ts.xx`
- **Non-Debug Version:** `libscjd12.xx`

Where `xx` represents the file extension, which depends on your UNIX operating system.

Ensure that the directory name where the library is located is specified with the appropriate environment variable, as follows:

- **Linux, Linux/390, and Solaris:** The `LD_LIBRARY_PATH` environment variable should specify the directory name for the required library file.
- **HP-UX:** The `SHLIB_PATH` environment variable should specify the directory name for the required library file.
- **AIX:** The `LIBPATH` environment variable should specify the directory name for the required library file.

Example (Solaris)

Assuming the appropriate library file is installed at `/home/IBM/Shadow/lib`, the library path variable would be specified as follows for Solaris:

```
LD_LIBRARY_PATH=/home/IBM/Shadow/lib:$LD_LIBRARY_PATH
```

Verifying the SHADOW_INI Environment Variable

Ensure that the `SHADOW_INI` environment variable points to the `shadow.ini` configuration file. Assuming a default installation, the configuration file will be copied to the following location:

```
/home/IBM/Shadow/bin/shadow.ini
```

Example

If the configuration file is installed at `/home/IBM/Shadow/bin`, the `SHADOW_INI` environment variable should be set as follows:

```
SHADOW_INI=/home/IBM/Shadow/bin/shadow.ini:$SHADOW_INI
```

Step 4: Set NEONTRACE Log File Permissions



Note:

This step is only required if debugging will be used.

If debugging is used, the NEONTRACE log file used for debugging must have READ/WRITE permission by the current effective user.

Step 5: Verify the Installation

To verify the installation of Shadow Mainframe Adapter Client, you can use the `jDemo` program, a graphical interface demo program, that is installed with Shadow Mainframe Adapter Client.



Doc Reference:

The `jDemo` program is documented in Chapter 5, “`jDemo`,” of this guide.

Additional Installation Options

The Shadow Mainframe Adapter Client installation can be configured and customized by specifying the installer properties. The installer properties are configured within a user-created text file called `installer.properties`. From this file, you can configure default installation properties, such as the installation directory location and the components to be installed, as well as specify the installation mode.

Configuring the Installation Properties

To configure the installation properties, simply do the following:

1. Within the directory where the actual Shadow Mainframe Adapter Client installation file is located, create a text file called `installer.properties`.



Note:

To use an `installer.properties` file that is located in a different directory, see “Command Line Installation Option” on page 3-24.

2. Specify the installation properties within the `installer.properties` file. The various preference settings are listed in Table 3–3. An example `installer.properties` file is shown in Figure 3–19.
3. Run the installation program. With the `installer.properties` file located in the same directory, the installation program will automatically use the `installer.properties` file to configure the installation settings.

Table 3–3. UNIX Installer Properties

Installer Property	Description
USER_INSTALL_DIR	This property specifies the destination for the installed files. You must provide a complete, absolute path to the location.
INSTALLER_UI	<p>This property sets the installation mode. Possible values are as follows:</p> <ul style="list-style-type: none"> • GUI: Executes the installation in GUI mode, as illustrated in the steps shown in “GUI Installation” on page 3-14. • CONSOLE: Executes the installation via a text interface instead of the GUI, as described in “Console Installation” on page 3-18. • SILENT: Executes the installation in silent mode, requiring no user input, as described in “Silent Installation” on page 3-19. In silent mode, the installation GUI is not invoked; thus, no user interaction is required or accepted. If non-default installation settings are required, they must be configured manually in the <code>installer.properties</code> file. <p><i>Note:</i> When Shadow Mainframe Adapter Client is installed in either console or silent mode, it will also be uninstalled in that mode.</p>
CHOSEN_INSTALL_SET	<p>This property chooses the installation set option. Possible values are as follows:</p> <ul style="list-style-type: none"> • TYPICAL • ADMINISTRATIVE <p>For details about the various installation options, see step 9 on page 3-16.</p>

```

# This is a sample installer.properties file that can be used with the
# Shadow Mainframe Adapter Client installation.

#####
# Use the USER_INSTALL_DIR setting to control where Shadow Mainframe Adapter
# Client will
# be installed. Provide a complete, absolute path to the location.
#####
USER_INSTALL_DIR=/usr/local/IBM/Shadow

#####
# Use the INSTALLER_UI setting to control whether the installer launches a
# GUI, a text interface, or runs silently. For a GUI, use "GUI" as the value.
# For the text interface, use "CONSOLE" as the value. For a silent install,
# use "SILENT" as the value. This sample is configured to run silently, as
# that is typical when running with a properties file.
#####
INSTALLER_UI=SILENT

#####
# The CHOSEN_INSTALL_SET setting allows you to choose between the high level
# install sets that are available in the install. They are "TYPICAL" and
# "ADMINISTRATIVE".
#
# TYPICAL: Includes Shadow Mainframe Adapter Client, samples, and user env
# variables.
# ADMINISTRATIVE: Includes Shadow Mainframe Adapter Client, samples, and
# system env

```

Figure 3–19. Example Installer Properties File -- UNIX

Command Line Installation Option

By manually specifying the installer properties file, you can force the installation to use an `installer.properties` file that is located in a different directory. Use the “-f” command line option to accomplish this, as follows:

```
$ /install_path/install_file -f //prop_path//installer.properties
```

► **Notes:**

- Note the “escaped” slashes used to specify the path for the `installer.properties` file.
- Command line options can also be used to change the installation mode. For more information, see “Step 2: Install Shadow Mainframe Adapter Client” on page 3-13.

Example

Assume the following:

- The installation program is located at /installs.
- The installer.properties file is located at /properties.

To run the installation with the installer.properties file, type the following on the command line:

```
$ /installs/file.bin -f //properties//installer.properties
```

Where file is the name of the Shadow Mainframe Adapter Client installation file (see Table 3-1 on page 3-1).

Common UNIX Errors

GUI Installation Failure

Symptom: An attempt to launch the Shadow Mainframe Adapter Client installation in GUI mode fails with the following error:

```
$ ./IBMShadowMainframe Adapter Client_3_9_nnn_xxx.bin
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...
Extracting the installation resources from the installer
archive...
Configuring the installer for this system's environment...

Launching installer...

Invocation of this Java Application has caused an
InvocationTargetException. This application will now exit. (LAX)
```

Stack Trace:

```
java.lang.NoClassDefFoundError
    at java.lang.Class.forName0(Native Method)
    at java.lang.Class.forName(Class.java:115)
    at java.awt.GraphicsEnvironment.getLocalGraphics
Environment(GraphicsEnvironment.java:53)
    at java.awt.Window.<init>(Window.java:183)
    at java.awt.Frame.<init>(Frame.java:310)
    at java.awt.Frame.<init>(Frame.java:257)
    at com.zerog.ia.installer.Main.c(Unknown Source)
    at com.zerog.ia.installer.Main.main(Unknown Source)
    at java.lang.reflect.Method.invoke(Native Method)
    at com.zerog.lax.LAX.launch(Unknown Source)
    at com.zerog.lax.LAX.main(Unknown Source)
```

GUI-

Description: This error indicates that the DISPLAY environment variable is not set appropriately.

Internal Error in scodbcco.c

Symptom: A message warning of an internal error detected in `scodbcco.c`, such as the following:

```
Wed May 14 16:41:25 2003 internal error detected: file scodbcco.c  
line 3416 rc = 0 from yp_get_default_domain  
Wed May 14 16:41:25 2003 internal error detected: file scodbcco.c  
line 957 rc = -1 from scinnsck
```

Description: This message indicates that the Shadow Mainframe Adapter Client Trace Facility is not configured properly; tracing has not been set to a valid log file. For example, in the following configuration, an appropriate path name is not specified:

```
NEONTRACE="INFO JDBCLOG=jdbc.log log=odbc.log"
```

Administration

CHAPTER 4: Shadow Mainframe Adapter Client: Configuration

Shadow Mainframe Adapter Client is part of the Mainframe Adapter Client component of Shadow Connect. Shadow Mainframe Adapter Client enables Java applications to integrate z/OS data and transactional sources through the JDBC API.

This chapter will cover the following topics:

- Connection Configuration
 - Step 1: Configure the Data Source
 - Step 2: Configure the Java Connection
- Additional Configuration and Administration
 - Using Optimized Fetch
 - DBCS Support
 - Determining Shadow Mainframe Adapter Client Version Information

Connection Configuration

To configure the connection, do the following:

1. Configure the data source.
2. Configure the Java connection.

Step 1: Configure the Data Source

To use Shadow Mainframe Adapter Client, the data source settings must be specified. This may be done using one of the following methods:

- Using the JDBC connection string.
- Using the jConfig tool to configure a JDBC data source.

Using the JDBC Connection String



Note:

Any Shadow Mainframe Adapter Client keywords set in the connection string override any keywords set within the JDBC data source settings.

You can make a connection without having to create a data source by appropriately configuring the JDBC connection string.

JDBC Connection String Syntax



Note:

Instead of formatting the JDBC connection string manually, you can also use the jConfig tool graphical user interface (GUI) to create a connection string. For further details, see “Configuring a JDBC Connection String” on page 4-16.

You must specify the following in the JDBC connection string:

- The data source name.
- All the required keywords to complete the connection.
- Any optional keywords.

The general format for the JDBC connection string is as follows:

```
jdbc:neon:<data-source-name>;  
<keyword1>=<value1>;...;<keywordN>=<valueN>
```

Data Source Name

The `jdbc:neon:<data-source-name>` statement must be specified as follows:

- If the connection string is being used *instead of* a JDBC data source, the data source name does not apply and can be set to any arbitrary value.
- If a predefined JDBC data source is being used (possibly in conjunction with a connection string), the data source name must be set appropriately to identify the data source.



Doc Reference:

For more information about configuring a JDBC data source, see “Configuring a JDBC Data Source” on page 4-5.

Required Keywords

Specify the following required keywords in the JDBC connection string:

- **UID:** Specifies the mainframe userid.
- **PWD:** Specifies the mainframe password.
- **HOST:** Specifies the host name or IP address. This field can contain a hostname string or a TCP/IP address in dot-notation format.

- **PORT:** Specifies the port number on which Shadow Mainframe Adapter Server is listening. You can find this number by examining Shadow Mainframe Adapter Server. It is often set to 1200.

**Note:**

Failover Shadow Mainframe Adapter Servers can be specified by listing host/port combinations as follows:

```
HOST=Host1,Host2,Host3;PORT=1200,1210,1500
```

Based on the host/port combinations specified, Shadow Mainframe Adapter Client would try to connect to Host1:1200, Host2:1210, and Host3:1500, in that order.

Specified host names can be the same host. Note that this syntax should *not* be used to specify Shadow Mainframe Adapter Server group directors.

- **SUBSYS:** Specifies the DB2 subsystem.

**Note:**

If set to DFLT, then the DEFAULTDB2SUBSYS parameter that is set on the Shadow Mainframe Adapter Server will be used.

- **PLAN:** Shadow uses plan names on the host to make a connection to DB2. The format for the plan name is as follows:
 - The first two letters of the name are always “SD.”
 - The third letter is always the last letter of the Shadow Mainframe Adapter Server subsystem name, which is “B” by default.
 - The fourth letter designates the isolation level of the plan, as follows:

“C” Cursor Stability	SQL_TXN_READ_COMMITTED
“R” Repeatable Read	SQL_TXN_SERIALIZABLE
“S” Read Stability	SQL_TXN_REPEATABLE_READ
“U” Uncommitted Read	SQL_TXN_READ_UNCOMMITTED

**Note:**

If the fourth character is any character other than “R,” “S,” or “U,” it is assumed the plan was bound with an isolation level of Cursor Stability.

- The final four characters are always a version number (such as “1010”).

For example, the default plan name value is SDBC1010. This default is release-dependent.



Note:

If set to either DFLT or DEFAULT, then the DEFAULTDB2PLANNAME parameter that is set on the Shadow Mainframe Adapter Server will be used for that connection.

Optional Keywords

In addition, any other optional Shadow Mainframe Adapter Client keywords can be set, as appropriate.



Doc Reference:

For a list of the keywords that are valid for the JDBC connection string, see Appendix A, “Shadow Mainframe Adapter Client Keywords,” of this guide.

Using the jConfig Tool

The jConfig tool is installed as part of the Shadow Mainframe Adapter Client installation. The jConfig tool offers an easy-to-use graphical interface that can be used to create and configure JDBC data sources and connection strings.



Note:

JDBC data source settings are stored in the SHADOW.INI configuration file, the location of which is as follows:

- **Windows:** The location is specified in the registry.
- **UNIX:** The file is located as described in “Verifying the SHADOW_INI Environment Variable” on page 3-21 within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide.

You can use the jConfig tool to do the following:

- Configure a JDBC data source.
- Import data sources.
- Configure a JDBC connection string.

Configuring a JDBC Data Source

To use the jConfig tool to create and/or configure a JDBC data source, do the following:

1. Launch the jConfig tool as follows:

- **Windows Users:** Go to **Start → Programs → IBM → Shadow Mainframe Adapter Client → jConfig**.
- **UNIX Users:** Launch the `ShadowConfig.sh` file, which will be installed to the following location by default:

```
/home/IBM/Shadow/bin/
```



Note:

When the jConfig tool is launched, the script will display the command that launches jConfig as follows:

```
Running java -DSHADOW_INI=...
```

The “-DSHADOW_INI=...” parameter defines the location of the `shadow.ini` file. This parameter is based on the `SHADOW_INI` environment variable. If it is incorrect, please verify the `SHADOW_INI` setting as described in “Verifying the `SHADOW_INI` Environment Variable” on page 3-21 within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide.

The system displays the jConfig tool, as shown in Figure 4–1.



Note:

The jConfig tool also accepts the following command line arguments:

- The **browse** command starts jConfig in browse mode to browse and select data sources for use with jDemo (see Chapter 5, “jDemo,” of this guide).
- The **import** command will copy and migrate existing data sources for use with Shadow Mainframe Adapter Client. The user interface *will not* be displayed.

Note: For UNIX, an absolute path for the `odbc.ini` file is required following the **import** command.

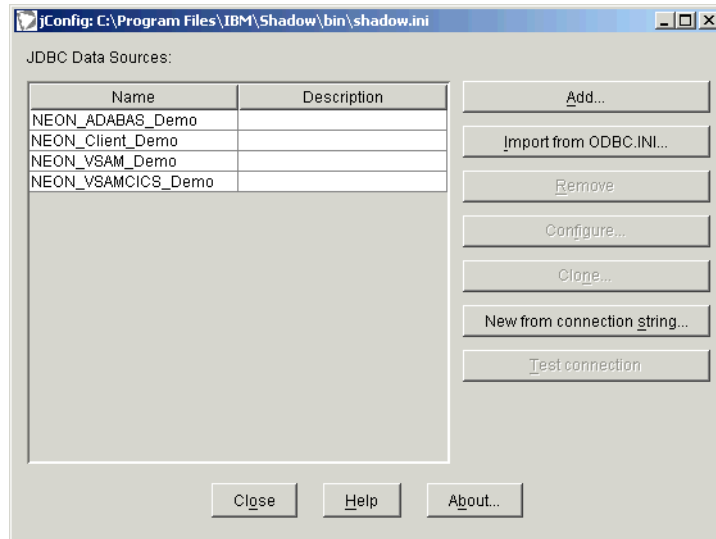


Figure 4–1. Main Window -- jConfig Tool

2. Either create a new data source or edit an existing data source, as follows:
 - Click **Add** to add a new JDBC data source.
 - Click **Import from ODBC.INI** to select existing data sources to import and convert into JDBC data sources. For more information, see “Importing Data Sources” on page 4-14.
 - Click **Configure** to edit an existing JDBC data source.
 - Click **Clone** to make a copy of an existing JDBC data source.
 - Click **New from Connection String** to create a JDBC data source based on a connection string that you provide.

The system displays the **JDBC Data Source Configuration** dialog box in a new window, as shown in Figure 4–2.



Note:

When creating a JDBC data source with the jConfig tool, the Shadow Mainframe Adapter Client version (debug versus non-debug) will be determined by the CLASSPATH setting. For more information, see the following within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide:

- **Windows Users:** See “Step 2: Verify the Search Paths” on page 3-8.
- **UNIX Users:** See “Step 3: Verify the Search Paths” on page 3-20.

Figure 4–2. JDBC Data Source Configuration

3. In the **Data Source Information** section, specify the following:
 - **DSN Name:** Type a name for the data source. For existing data sources, the name will be entered by default; however, you can change the data source name by typing over the existing name.
 - **DSN Description:** (Optional) Type a description for the data source, if desired.

▶ **Note:**
If you enter a description, spaces are allowed.
 - **DBMS Type:** Select **DB2**.
4. In the **TCP/IP Connection Information** section, specify the following:
 - **Host Name: (Required)** The host name or TCP/IP address for the connection. This field can contain a hostname string or a TCP/IP address in dot-notation format.
 - **Port Number: (Required)** The port number on which Shadow Mainframe Adapter Server is listening. You can find this number by examining Shadow Mainframe Adapter Server. It is often set to 1200.

**Note:**

Failover Shadow Mainframe Adapter Servers can be specified by listing multiple host/port values, separated by commas. Shadow Mainframe Adapter Client will try to connect to each listed host/port pair in order.

Specified host names can be the same host. Note that this syntax should *not* be used to specify Shadow Mainframe Adapter Server group directors.

- **User ID:** Mainframe userid.
 - **Password:** Mainframe password. This is automatically encrypted (both on the screen and in the data source settings).
5. In the **Optional Information** section, click **Advanced**. The system displays the **Advanced Information** dialog box, as shown in Figure 4–3.

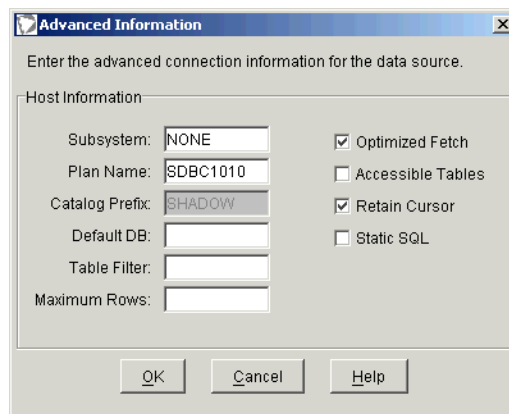


Figure 4–3. Advanced Information

6. Specify the following information:
- **Subsystem:** The DB2 subsystem to which you want to connect.

**Note:**

If the Mainframe Adapter Client sets **Subsystem** to DFLT, then the DEFAULTDB2SUBSYS parameter that is set on the Shadow Mainframe Adapter Server will be used for that connection.

- **Plan Name:** Shadow uses plan names on the host to make a connection to DB2. The format for the plan name is as follows:
 - The first two letters of the name are always “SD.”

- The third letter is always the last letter of the Shadow Mainframe Adapter Server subsystem name, which is “B” by default.
- The fourth letter designates the isolation level of the plan, as follows:
 - “C” Cursor Stability SQL_TXN_READ_COMMITTED
 - “R” Repeatable Read SQL_TXN_SERIALIZABLE
 - “S” Read Stability SQL_TXN_REPEATABLE_READ
 - “U” Uncommitted Read SQL_TXN_READ_UNCOMMITTED

**Note:**

If the fourth character is any character other than “R,” “S,” or “U,” it is assumed the plan was bound with an isolation level of Cursor Stability.

- The final four characters are always a version number (such as “1010”).

For example, the default plan name value is SDBC1010. This default is release-dependent.

**Note:**

If the Mainframe Adapter Client sets **Plan Name** to either DFLT or DEFAULT, then the DEFAULTDB2PLANNNAME parameter that is set on the Shadow Mainframe Adapter Server will be used for that connection.

- **Catalog Prefix:** The PROCEDURE QUALIFIER (in ODBC 2.0 terminology) or the PROCEDURE CAT (in ODBC 3.0 terminology). Specify one of the following values:
 - **SYSIBM:** (Default) Accept the default value to use the standard system catalogs (SYSIBM).
 - **SHADOW:** Specify SHADOW if you are using the Shadow-optimized catalog feature, which requires DB2. In the latter case, SHADOW should represent the table owner (in DB2 terminology, the authorization ID) for a set of DB2 tables that are optimized to support catalog queries.

**Note:**

Shadow-optimized catalogs are considered essential for smooth operation of pre-packaged applications with the Shadow product.

- **Default DB:** If you need to create tables in DB2 using your application, you must specify the name of the database in which your userid has authority to create tables. This value is added to CREATE TABLE statements that do not specify a database in which to create the table. If the statement did not specify the IN DATABASE clause, Shadow Mainframe Adapter Client then appends this clause. This is *required* for users who do not have authority to create tables in DB2's default database, which includes most users.
- **Table Filter:** This field, along with the accessible tables feature described later, is used to restrict the information returned by catalog inquiries (SQLTables, SQLTablePrivileges, and SQLColumns). When your application requests a list of all the tables in your RDBMS instance, Shadow Mainframe Adapter Client constructs this list based on the value of the table filter. If this field is left blank, only tables that fall under the first-level name (userid) are returned. This filter can include the two SQL wildcard characters “%” and “_”, where “%” substitutes for zero or more characters and “_” for exactly one character. This field is useful because the RDBMS typically sends thousands of rows if it is not limited.

**Notes:**

- If the **Table Filter** field is left blank, the filter defaults to the userid value, which means Shadow Mainframe Adapter Client only returns tables that you own or created under your userid. If you are sure that no filter is needed, specify “%” in the field. This will return all tables. For further filtering, for example, ai38% will return tables that begin with “ai38”, such as ai38og.staff, ai38pds.staff, and ai38pds.xyz.
- More than one table filter can be specified by comma delimiting the entries.
- **Maximum Rows:** This field contains an integer that limits the number of rows returned from a single query. If no value is entered, no restriction is placed on the number of rows returned.
- **Optimized Fetch:** This box should be checked if you want to use the block fetch feature of Shadow Mainframe Adapter Client to speed up your queries. The increase in performance is significant.

**Note:**

When the block fetch feature is enabled, cursor-based deletes (DELETE WHERE CURSOR OF) cannot be used.

**Doc Reference:**

For more information about the optimized fetch feature, see “Using Optimized Fetch” on page 4-18.

- **Accessible Tables:** Many applications require information about the tables in the connected data source. However, in DB2’s case, this list of tables is unmanageably large unless you provide restrictions on the size. To do this, select the **Accessible Tables** box. Only those tables for which you have SELECT authority are returned. This includes those in PUBLIC and PUBLIC AT ALL LOCATIONS but not those where you have SELECT authority by virtue of a secondary authorization ID.
- **Retain Cursor:** This box controls how DB2 COMMIT operations affect cursors in the data source. **Retain Cursor** should be selected to preserve cursors in the same position as before the COMMIT operation took place. This enables the application to execute or fetch without preparing the statement again. The default is set to ON. If you uncheck the box (this sets the value to OFF), the application will close and delete cursors after a COMMIT operation, meaning you must prepare and execute the next statement from scratch. If you are not sure about this value, set it to ON (checked).

**Note:**

The **Retain Cursor** option also causes Shadow Mainframe Adapter Server to issue COMMITs after every SELECT statement as long as auto-commit is enabled by the application.

- **Static SQL:** Not applicable.
7. Click **OK** to return to the **JDBC Data Source Configuration** dialog box, as shown in Figure 4–2 on page 4-7.
 8. If you wish to enable base tracing, click **Debug**. The system will display the **Debug Information** dialog box, as shown in Figure 4–4. You can use this dialog box to configure the Shadow Mainframe Adapter Client Trace Facility options.

**Notes:**

- To enable tracing, you must be using a debug version of Shadow Mainframe Adapter Client (as specified via the CLASSPATH setting).
- The **Debug Information** dialog box will *only* configure base tracing. To configure JDBC traces, you must set the NEONTRACE environment variable appropriately.

**Doc Reference:**

For more information about the Shadow Mainframe Adapter Client Trace Facility, see Chapter 6, “Shadow Mainframe Adapter Client Trace Facility,” of this guide.

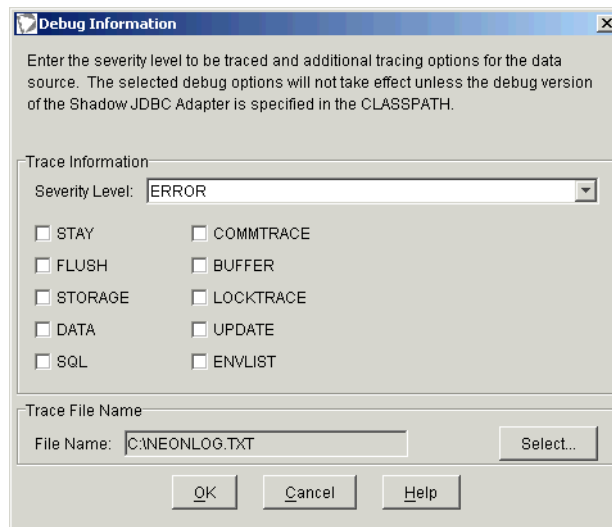


Figure 4–4. Debug Information

9. Click **OK** to return to the **JDBC Data Source Configuration** dialog box, as shown in Figure 4–2 on page 4-7.
10. In the **Optional Information** section, click **Parameters**. The system will display the **Optional Parameter Information** dialog box, as shown in Figure 4–5. You can use this dialog box to set additional Shadow Mainframe Adapter Client keyword values.

**Note:**

Some of these keyword values may need to be set to non-default values. The keywords and their values are detailed in Appendix A, “Shadow Mainframe Adapter Client Keywords,” of this guide.

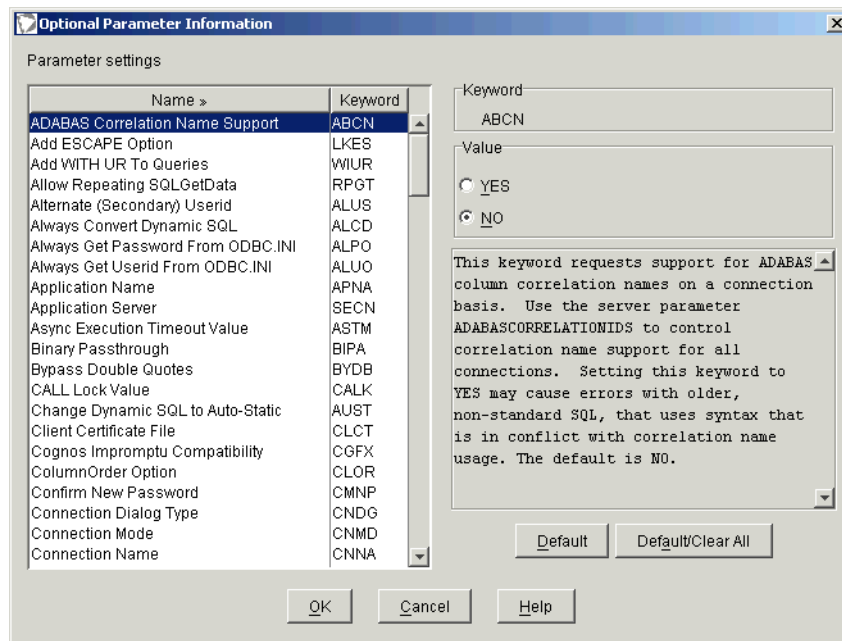


Figure 4–5. Optional Parameter Information

**Note:**

You can sort the list either by **Name** (the long name) or **Keyword** by simply clicking the appropriate column heading.

11. Click **OK** to return to the **JDBC Data Source Configuration** dialog box, as shown in Figure 4–2 on page 4-7.
12. Click **OK** to return to the main window of the jConfig tool (see Figure 4–1 on page 4-6). Your entered values will be written into the SHADOW.INI settings for that JDBC data source.

Testing a JDBC Data Source Connection

Once you have created a JDBC data source, you can verify the configuration by testing the JDBC data source connection as follows:

1. Within the main window of the jConfig tool (see Figure 4–1 on page 4-6), from the **JDBC Data Sources** list, select the data source.
2. Click **Test Connection**. The system displays the **Test Connection** dialog box, indicating the results of the test, as shown in Figure 4–6.




Figure 4–6. Test Connection -- Successful Connection

If the connection test is not successful, the resulting error message will be displayed to offer assistance in problem diagnosis.

Importing Data Sources

When Shadow Mainframe Adapter Client is installed, any existing data sources that have been determined to be Shadow Mainframe Adapter Client data sources¹ will be automatically copied and migrated for use with the updated version of Shadow Mainframe Adapter Client.



Note:

After the data sources are imported, existing data sources will *not* be removed; it will be left to the user to delete unused data sources.

In addition, the jConfig tool also offers a feature to allow users to manually select existing data sources to be imported for use with Shadow Mainframe Adapter Client. To import data sources, do the following:

1. From the main window of the jConfig tool (see Figure 4–1 on page 4-6), click **Import from ODBC.INI**. The system will locate existing Shadow Mainframe Adapter Client data sources as follows:
 - **Windows:** The jConfig tool uses the registry to find Shadow Mainframe Adapter Client data sources defined in the ODBC.INI settings.

¹ “Shadow Mainframe Adapter Client data sources” are defined as ODBC/TCPIP data sources with the APPL, DBTY, HOST, and PLAN keywords configured (empty string values are acceptable). For details about Shadow Mainframe Adapter Client keywords, see Appendix A, “Shadow Mainframe Adapter Client Keywords,” of this guide.

- **UNIX:** The jConfig tool will display the **Select ODBC.INI File to Import Data Sources** dialog box, allowing users to navigate to and select the appropriate file from which to import existing data sources.

Upon finding the existing Shadow Mainframe Adapter Client data sources, the system will display the **Import Data Sources** dialog box, as shown in Figure 4–7.

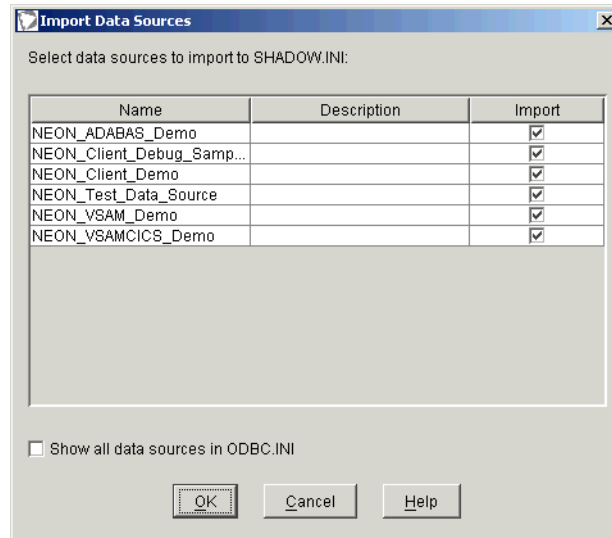


Figure 4–7. Import Data Sources



Note:

By default, only the data sources that have been determined to be Shadow Mainframe Adapter Client data sources will be displayed. If you wish to display all existing data sources, select the **Show all data sources in ODBC.INI** check box.

2. From the list of data sources, select the ones to be imported.
3. Click **OK**.
4. The data source will be imported and added to the list of data sources shown in the main window of the jConfig tool.



Note:

If an imported ODBC data source has the same name as an existing JDBC data source, the imported data source name will be suffixed by a “2”.

Configuring a JDBC Connection String

To use the jConfig tool to create and/or configure a JDBC connection string, do the following:

1. Use the jConfig tool to configure a JDBC data source as described in “Configuring a JDBC Data Source” on page 4-5.



Note:

In order to create a connection string, you do *not* need to save the JDBC data source.

2. From the **JDBC Data Source Configuration** dialog box (see Figure 4–2 on page 4-7), click **Show Connection String**. The system will display the **Connection String** dialog box, as shown in Figure 4–8.

From the **Connection String** dialog box, you can copy the connection string to paste into Java source codes or configuration dialog boxes for application Mainframe Adapter Servers.




Figure 4–8. Connection String

3. Close the dialog box as follows:
 - Click **Return to Configuration** to return to the **JDBC Data Source Configuration** dialog box (see Figure 4–2 on page 4-7).
 - Click **Exit Configuration** to return to the main window of the jConfig tool (see Figure 4–1 on page 4-6).

Step 2: Configure the Java Connection

There are two methods to make a Java connection with Shadow Mainframe Adapter Client:

- (Recommended) Using `javax.sql.DataSource` to connect.
- Using `java.sql.Driver` to connect.

Using javax.sql.DataSource to Connect

The recommended way of getting the correct DataSource object is via the Java Naming and Directory Interface (JNDI) service, which involves using javax.sql.DataSource to connect, as follows:

1. The DataSource object will depend on the environment, as follows:

- **Users in a J2EE compatible environment:** Register the com.neon.jdbc.DataSource object with JNDI as follows:

```
com.neon.jdbc.DataSource ds = new
com.neon.jdbc.DataSource();
ds.setDataSourceName ("XYZ");
Context ctx = new InitialContext ();
ctx.bind ("jdbc/MyDB", ds);
```

Where XYZ specifies the name of the data source created in “Step 1: Configure the Data Source” on page 4-1.

- **Users in a non-J2EE compatible environment:** Use the com.neon.jdbc.SimpleDataSource object after including the jdbc2_0-stdext.jar file in the CLASSPATH environment variable as follows (assuming default installations):

- **Windows:** SET CLASSPATH=C:\Program Files\IBM\Shadow\jdbc\jdbc2_0-stdext.jar;%CLASSPATH%
- **UNIX:** CLASSPATH=/home/IBM/Shadow/jdbc/jdbc2_0-stdext.jar:\$CLASSPATH

2. Retrieve the DataSource object as follows:

```
Context ctx = new InitialContext ();
DataSource ds = (DataSource) ctx.lookup ("jdbc/MyDB");
```

3. Connect to the data source, as follows:

```
Connection conn = ds.getConnection ("username", "password");
```

Where username specifies the userid and password specifies the password.

Using java.sql.Driver to Connect

To use javax.sql.DataSource to connect, do the following:

1. Create the following driver class:

```
Class cls = Class.forName("com.neon.jdbc.Driver");
```

2. Create a driver instance as follows:

```
java.sql.Driver driver = (java.sql.Driver) cls.newInstance();
```

3. Connect to the data source using the JDBC connection string as follows:

```
java.sql.Connection conn =  
driver.connect ("jdbc:neon:XYZ;UID=xxx;PWD=yyy", null);
```

Where XYZ specifies the name of the data source created in “Step 1: Configure the Data Source” on page 4-1, xxx specifies the userid, and yyy specifies the password.

Additional Configuration and Administration

Other configuration options and administrative features include the following:

- Using optimized fetch.
- DBCS support.
- Determining Shadow Mainframe Adapter Client version information.

Using Optimized Fetch

When Shadow Mainframe Adapter Client is instructed to use optimized fetch, it adds the FOR FETCH ONLY clause to every SELECT statement it receives that does not include a FOR UPDATE clause. The statements are then sent to the Shadow Mainframe Adapter Server, which detects the FOR FETCH ONLY statement and begins to asynchronously pre-extract rows ahead of the current row. Shadow Mainframe Adapter Server sends these rows back to Shadow Mainframe Adapter Client in large blocks. Each block can hold as much as 32 kilobytes of data. The use of optimized fetch is highly recommended because it enhances performance in two ways:

- Network traffic is minimized, since blocks of information are sent.
- Subsequent fetches done with Shadow Mainframe Adapter Client can nearly always be satisfied by data that has already been transmitted.

In general, block fetch improves performance for queries that process large portions of tables without materially affecting single-row queries. However, if block fetch is used with a Repeatable Read (RR) plan, many more pages may be locked for update than without block fetch, especially if the number of rows normally extracted by the query is small.

Enabling Optimized Fetch

Optimized fetch is enabled as follows:

- **Using the Connection String:** Enable optimized fetch by setting RO=YES in the connection string to set the appropriate Shadow Mainframe Adapter Client keyword.
- **Using a JDBC Data Source:** Enable optimized fetch via the jConfig tool by navigating to the **Configure** → **Advanced** option and selecting **Optimized Fetch**.

**Doc References:**

- For detailed descriptions of Shadow Mainframe Adapter Client keywords, see Appendix A, “Shadow Mainframe Adapter Client Keywords,” of this guide.
- For more information about setting Shadow Mainframe Adapter Client keywords, see “Step 1: Configure the Data Source” on page 4-1.

Also, block fetch can be used without making configuration changes to Shadow Mainframe Adapter Client. By manually adding the FOR FETCH ONLY clause to the SELECT statement, you instruct Shadow Mainframe Adapter Server to use block fetch for that series of operations.

Optimized Fetch Restriction

The only restriction on using optimized fetch is that cursor-based deletes cannot be used while optimized fetch is turned on. In other words, you cannot use DELETE WHERE CURRENT OF statements.

This restriction can be circumvented by passing a SELECT statement that includes the FOR UPDATE clause. In this case, Shadow Mainframe Adapter Client does not add the FOR FETCH ONLY clause to the end of the statement and cursor-based deletes can be used as needed. It is not necessary to perform an UPDATE when using this clause.

DBCS Support

Shadow Mainframe Adapter Client supports DBCS on all of its supported platforms *except* Windows 95/98.

**Note:**

In order to support DBCS on Linux, you must have glibc2.1.1 on your Linux machine.

**Doc Reference:**

For a list of the platforms supported by Shadow Mainframe Adapter Client, see Chapter 2, “Planning the Shadow Mainframe Adapter Client Installation,” of this guide.

Determining Shadow Mainframe Adapter Client Version Information

Shadow Mainframe Adapter Client offers a feature to enable users to quickly and easily determine information about their version of Shadow Mainframe Adapter Client.

The `DriverInfo` class will return information about Shadow Mainframe Adapter Client.

Assuming a default location, the `DriverInfo` class will be located as follows:

- **Windows:** `C:\Program Files\IBM\Shadow\jdbc`
- **UNIX:** `/home/IBM/Shadow/jdbc`


Execute the `DriverInfo` class to determine Shadow Mainframe Adapter Client information as follows:

```
java DriverInfo -jdbc
```

The Shadow Mainframe Adapter Client name, version, build date, and other relevant information will be returned.

Example

Figure 4–9 shows an example of executing the `DriverInfo` class on Windows to return information about Shadow Mainframe Adapter Client.



```

C:\Program Files\IBM\Shadow\jdbc>java DriverInfo -jdbc
Driver Name       : NEON JDBC Driver
Driver Version    : 3.9.291 Debug
Driver Build Date : Jan  8 2004 23:38:13
Native Library    : scjd12ts.dll
UM Name           : Java HotSpot(TM) Client VM
UM Version        : 1.4.1_03-b02
UM Info           : mixed mode
OS Name           : Windows 2000
OS Version        : 5.0
Architecture      : x86
Language          : en
Region            : Unknown
Class Path        : ;C:\Program Files\IBM\Shadow\jdbc\scjd12ts.jar;C:\j2sdk1.4.1_03\lib\tools.jar;C:\j2sdk1.4.1_03\lib\dt.jar;C:\j2sdk1.4.1_03\jre\lib\rt.jar;C:\j2sdk1.4.1_03\lib\j2ee.jar
Library Path      : C:\j2sdk1.4.1_03\bin;.C:\WINNT\system32;C:\WINNT;C:\j2sdk1.4.1_03\bin;C:\j2sdk1.4.1_03\jre\bin;C:\WINNT\system32;C:\WINNT;C:\WINNT\System32\Wbem;C:\WINNT\Crystal;C:\PROGRA~1\MAPINF~1\Program;C:\PROGRA~1\SEAGAT~1\Shared;C:\Program Files\IBM\Shadow\bin
Boot Class Path   : C:\j2sdk1.4.1_03\jre\lib\rt.jar;C:\j2sdk1.4.1_03\jre\lib\i18n.jar;C:\j2sdk1.4.1_03\jre\lib\sunrsasign.jar;C:\j2sdk1.4.1_03\jre\lib\jsse.jar;C:\j2sdk1.4.1_03\jre\lib\jce.jar;C:\j2sdk1.4.1_03\jre\lib\charsets.jar;C:\j2sdk1.4.1_03\jre\classes
Boot Library Path : C:\j2sdk1.4.1_03\jre\bin
C:\Program Files\IBM\Shadow\jdbc>

```

Figure 4–9. Execute `DriverInfo` for Shadow Mainframe Adapter Client

CHAPTER 5: *jDemo*

The jDemo program is a graphical user interface (GUI) provided with Shadow Mainframe Adapter Client, part of the Mainframe Adapter Client component of the Shadow Connect product. The jDemo interface is used to access data structures using Shadow Mainframe Adapter Client.

Topics include the following:

- Overview
- Starting jDemo
- Supported Features
 - Establishing a Connection
 - Viewing Data Source Information
 - Executing SQL
- Closing jDemo

Overview

jDemo was installed onto your computer as part of the standard Shadow Mainframe Adapter Client installation. jDemo can be used to demonstrate some basic database functionality through Shadow Mainframe Adapter Client. It provides a GUI to allow the user to do the following:

- Make a database connection.
- Retrieve general data source information.
- Execute SQL statements to access ADABAS, CICS, DB2, IMS, and VSAM data, as well as Natural/ACI programs.
- Collect query results.

Prerequisites

jDemo can be run on any Mainframe Adapter Client machine meeting the following prerequisites:

- Shadow Mainframe Adapter Client must be installed.
- Java Runtime Environment v1.2 or greater must be available.

Starting jDemo

To start jDemo, do the following:

- **Windows Users:** From the Mainframe Adapter Client PC Windows **Start** menu, select **Programs** → **IBM** → **Shadow Mainframe Adapter Client** → **jDemo**.
- **UNIX Users:** Launch the `NeonJavaDemo.sh` file, which will be installed to the following location by default:

```
/home/IBM/Shadow/bin/
```



Note:

When the jDemo tool is launched, the script will display the command that launches jDemo as follows:

```
Running java -DSHADOW_INI=...
```

The “-DSHADOW_INI=...” parameter defines the location of the `shadow.ini` file. This parameter is based on the `SHADOW_INI` environment variable. If it is incorrect, please verify the `SHADOW_INI` setting as described in “Verifying the `SHADOW_INI` Environment Variable” on page 3-21 within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide.

The main jDemo window opens, as shown in Figure 5–1.




Figure 5–1. Main Window -- jDemo

Supported Features

jDemo supports the following features:

- **Establishing a Connection:** Users can establish single or multiple connections to data sources through various means:
 - Establish connections using a connection string.
 - Establish connections using a defined JDBC data source.
- **Viewing Data Source Information:** Users can view the following data related to a data source once a connection is established:
 - Display Shadow Mainframe Adapter Client name and version information for an established connection.
 - Display database name and version information for an established connection.
 - Display tables and stored procedures for the data source.
- **Executing SQL:** Users can execute SQL statements against and display query results.

Establishing a Connection

To establish a connection using jDemo, do one of the following:

- In the jDemo window, go to the **Connections** menu and select **Add New Connection**.

Or

- In the jDemo window, click **New**.

The system displays the **New Database Connection** dialog box, as shown in Figure 5–2.



Note:

You *can* have more than one connection open at the same time.

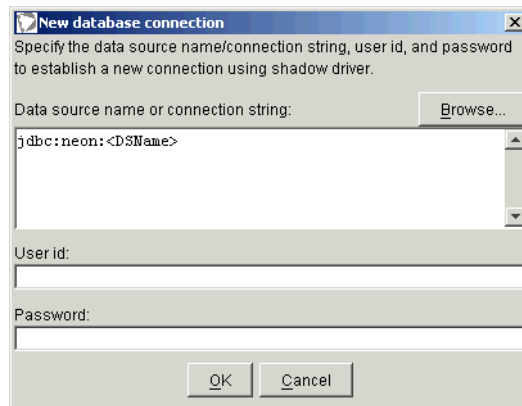


Figure 5–2. New Database Connection

From the **New Database Connection** dialog box, you can do either of the following:

- Establish a connection using a connection string.
- Establish a connection using a defined JDBC data source.

Establishing a Connection Using a Connection String

To establish a connection using a connection string, you can do either of the following:

- Type a valid connection string into the text box as follows:

```
jdbc:neon:DataSourceName;keyword1=value1;keyword2=value2;...
```

Where `DataSourceName` can be set to any arbitrary value because it does not apply and **must not** specify an existing JDBC data source. The keyword and value pairs should specify the desired configuration.



Doc Reference:

For more information about the connection string syntax, see “Using the JDBC Connection String” on page 4-1 within Chapter 4, “Shadow Mainframe Adapter Client: Configuration,” of this guide.

Or

- Click **Browse** to use the jConfig tool to create a connection string.

**Doc Reference:**

For more information about using the jConfig tool to create a connection string, see “Configuring a JDBC Connection String” on page 4-16 within Chapter 4, “Shadow Mainframe Adapter Client: Configuration,” of this guide.

**Note:**

The **Userid** and **Password** fields can be used to specify the userid and password, respectively. This will prevent the password from being displayed in clear text in the connection string. Any values specified in those fields will override any values specified in the connection string.

Establishing a Connection Using a JDBC Data Source

To establish a connection using a JDBC data source, click **Browse**. The system opens the jConfig tool, which allows users to select an existing, defined JDBC data source or create a new JDBC data source.

**Doc Reference:**

For more information about creating and configuring a JDBC data source, see “Configuring a JDBC Data Source” on page 4-5 within Chapter 4, “Shadow Mainframe Adapter Client: Configuration,” of this guide.

**Note:**

The **Userid** and **Password** fields can be used to specify the userid and password, respectively, overriding any values specified in the JDBC data source configuration.

Established Connections

For each established connection, the system issues a message to indicate the connection was established, as shown in Figure 5–3.




Figure 5–3. Connection Established

In addition, after the connection is made, the connection is displayed in the **Current Connection** list as “data-source-name -- com.neon.jdbc.Driver”.

Viewing Data Source Information

Once a connection is established, the following data related to a data source can be viewed:

- Shadow Mainframe Adapter Client name and version information.
- Database name and version information.
- Tables, table contents, and/or table columns for the data source.
- Stored procedures for the data source.

Displaying Shadow Mainframe Adapter Client Information

To display the Shadow Mainframe Adapter Client name and version information, select the **Driver** tab, as shown in Figure 5–4.




Figure 5–4. jDemo -- Driver Tab

Displaying Database Information

To display the database name and version information, select the **Database** tab, as shown in Figure 5–5.




Figure 5–5. jDemo -- Database Tab

Displaying Tables and Related Information

To display the tables, table contents, and/or table columns for the data source, select the **Tables** tab. When the tab is first selected, no tables are listed.

Displaying Tables

To display the tables for the data source, do the following:

1. From the **Tables** tab, click **Load Tables**. The system displays the **Load Tables -- Parameters** dialog box, as shown in Figure 5–6.

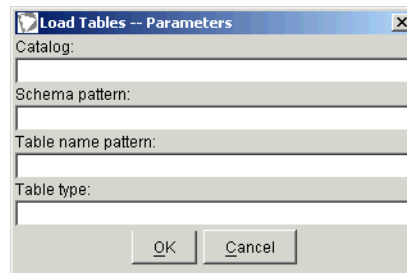


Figure 5–6. Load Tables -- Parameters

2. Specify criteria in the available fields to filter the tables that will be returned.



Note:

The tables that are returned also depend on the table filter, which is set as follows:

- **Connection string:** Set via the DP keyword.
- **JDBC data source:** Set via the **Table Filter** field from the **Advanced** options.

For more information, see “Step 1: Configure the Data Source” on page 4-1 within Chapter 4, “Shadow Mainframe Adapter Client: Configuration,” of this guide.

3. Click **OK**. jDemo will load and display the tables that meet the specified criteria, as shown in Figure 5–7.

Driver	Database	Tables	Stored Procedures		
	TABLE_CAT	TABLE_SCHEM	TABLE_NAME	TABLE_TYPE	REMARKS
1	SEQ	SEQ	FLATFILE	TABLE	
2	SEQ	SEQ	GDGDSN	TABLE	
3	SEQ	SEQ	HEATHER	TABLE	
4	SEQ	SEQ	PDSMEM	TABLE	
5	SEQ	SEQ	PVFLAT	TABLE	
6	VSAM	VSAM	ANDFILEA	TABLE	
7	VSAM	VSAM	ANDYSTAF	TABLE	
8	VSAM	VSAM	FH001R	TABLE	
9	VSAM	VSAM	FILEA	TABLE	
10	VSAM	VSAM	FILEB	TABLE	
11	VSAM	VSAM	FILEC	TABLE	
12	VSAM	VSAM	FILED	TABLE	
13	VSAM	VSAM	FILEE	TABLE	
14	VSAM	VSAM	FILEF	TABLE	
15	VSAM	VSAM	FILEG	TABLE	
16	VSAM	VSAM	FILEH	TABLE	
17	VSAM	VSAM	FILEI	TABLE	
18	VSAM	VSAM	FILEJ	TABLE	
19	VSAM	VSAM	FILEK	TABLE	
20	VSAM	VSAM	FILEL	TABLE	
21	VSAM	VSAM	FILEM	TABLE	
22	VSAM	VSAM	FILEN	TABLE	
23	VSAM	VSAM	FILEO	TABLE	
24	VSAM	VSAM	FILEP	TABLE	
25	VSAM	VSAM	FILEQ	TABLE	
26	VSAM	VSAM	FILER	TABLE	
27	VSAM	VSAM	FILES	TABLE	
28	VSAM	VSAM	FILET	TABLE	
29	VSAM	VSAM	FILEU	TABLE	
30	VSAM	VSAM	FILEV	TABLE	
31	VSAM	VSAM	FILEW	TABLE	
32	VSAM	VSAM	FILEX	TABLE	
33	VSAM	VSAM	FILEY	TABLE	
34	VSAM	VSAM	FILEZ	TABLE	

Load tables... Table: Show table contents Show table columns...

Figure 5–7. jDemo -- Tables Tab

Displaying Table Contents

To display the table contents for a table, do the following:

1. Do one of the following to select a table:
 - If the table information for the desired table has already been loaded and displayed (as shown in Figure 5–7), then from the **Tables** tab, within the **TABLE_NAME** column, simply select the desired table name. The row will be highlighted, as shown in Figure 5–8.

Or

- Within the **Table** field, type the name of the table.

Driver	Database	Tables	Stored Procedures				
		TABLE_CAT	TABLE_SCHEM	TABLE_NAME	TABLE_TYPE	REMARKS	
21	VSAM		VSAM	PVTEST1	TABLE		
22	VSAM		VSAM	QAVSAM1	TABLE		
23	VSAM		VSAM	SDCOVS	TABLE		
24	VSAM		VSAM	SDSVDNUM	TABLE		
25	VSAM		VSAM	TESTCHAR	TABLE		
26	VSAM		VSAM	VSAM1	TABLE		
27	VSAM		VSAM	VSAM1PT1	TABLE		
28	VSAM		VSAM	VSAM1RM	TABLE		
29	VSAM		VSAM	VSAM2	TABLE		
30	VSAM		VSAM	VSAM2TST	TABLE		
31	VSAM		VSAM	VSAM3	TABLE		

Load tables... Table: VSAM.VSAM1 Show table contents Show table columns...

Figure 5–8. jDemo -- Tables Tab with Selected Table

2. Click **Show Table Contents**. The system displays the table contents in the **Results** pane, as shown in Figure 5–9.

Current connection: JDBC_Data_Source_Demo -- com.neon.jdbc.Driver

TABLE_CAT	TABLE_SCHEM	TABLE_NAME	TABLE_TYPE	REMARKS	
873	DSQDBCTL	Q	PROFILES	TABLE	QMF USER PROFILE TABLE
874	DSQDBCTL	Q	RESOURCE_TABLE	TABLE	RESOURCE TABLE FOR Q...
875	DSQ1STBB	Q	APPLICANT	TABLE	QMF SAMPLE TABLE
876	DSQ1STBB	Q	INTERVIEW	TABLE	QMF SAMPLE TABLE
877	DSQ1STBB	Q	ORG	TABLE	QMF SAMPLE TABLE
878	DSQ1STBB	Q	PARTS	TABLE	QMF SAMPLE TABLE
879	DSQ1STBB	Q	PRODUCTS	TABLE	QMF SAMPLE TABLE
880	DSQ1STBB	Q	PROJECT	TABLE	QMF SAMPLE TABLE
881	DSQ1STBB	Q	SALES	TABLE	QMF SAMPLE TABLE
882	DSQ1STBB	Q	STAFF	TABLE	QMF SAMPLE TABLE

Table: Q.STAFF

SQL:

Query

Results:

	ID	NAME	DEPT	JOB	YEARS	SALARY	COMM
1	10	SANDERS	25	MGR	7	18357.50	NULL
2	20	PERNAL	20	SALES	8	18171.25	612.45
3	30	MARENGHI	38	MGR	5	17506.75	NULL
4	40	O'BRIEN	38	SALES	6	18006.00	846.55
5	50	HANES	15	MGR	10	20659.80	NULL
6	60	QUIGLEY	38	SALES	NULL	16808.30	650.25
7	70	ROTHMAN	15	SALES	7	16502.83	1152.00
8	80	JAMES	20	CLERK	NULL	13504.60	128.20
9	90	KOONITZ	42	SALES	6	18001.75	1386.70
10	100	PLOTZ	42	MGR	7	18352.80	NULL
11	110	NGAN	15	CLERK	5	12508.20	206.60
12	120	NAUGHTON	38	CLERK	NULL	12954.75	180.00
13	130	YAMAGUCHI	42	CLERK	6	10505.90	75.60
14	140	FRAYE	51	MGR	6	21150.00	NULL

Figure 5–9. Table Contents

Displaying Table Columns

To display the columns for a table, do the following:

1. If the table information for the desired table has already been loaded and displayed (as shown in Figure 5–7 on page 5-7), then from the **Tables** tab, within the **TABLE_NAME** column, simply select the desired table name. The row will be highlighted (see Figure 5–8).
2. Click **Show Table Columns**. The system displays the **Load Table Columns -- Parameters** dialog box, as shown in Figure 5–10.

Load Table Columns -- Parameters

Catalog:

Schema pattern:
VSAM

Table name pattern:
VSAM1

Column name pattern:

OK Cancel Help

Figure 5–10. Load Table Columns -- Parameters

3. Specify criteria in the available fields to filter the tables that will be returned.

**Note:**

If a table was selected in step 1, then the **Schema Pattern** and **Table Name Pattern** fields will be specified according to the selected table; otherwise, you must specify the criteria appropriately.

- Click **OK**. The system displays the table columns in the **Results** pane, as shown in Figure 5–11 on page 5-10.

The screenshot shows the jDemo application window. The 'Tables' tab is active, displaying a list of tables. The table 'STAFF' is selected. Below the table list, the 'Results' pane shows the columns for the 'STAFF' table.

TABLE_CAT	TABLE_SCHEM	TABLE_NAME	TABLE_TYPE	REMARKS	
873	DSQDBCTL	Q	PROFILES	TABLE	QMF USER PROFILE TABLE
874	DSQDBCTL	Q	RESOURCE_TABLE	TABLE	RESOURCE TABLE FOR Q...
875	DSQ1STBB	Q	APPLICANT	TABLE	QMF SAMPLE TABLE
876	DSQ1STBB	Q	INTERVIEW	TABLE	QMF SAMPLE TABLE
877	DSQ1STBB	Q	ORG	TABLE	QMF SAMPLE TABLE
878	DSQ1STBB	Q	PARTS	TABLE	QMF SAMPLE TABLE
879	DSQ1STBB	Q	PRODUCTS	TABLE	QMF SAMPLE TABLE
880	DSQ1STBB	Q	PROJECT	TABLE	QMF SAMPLE TABLE
881	DSQ1STBB	Q	SALES	TABLE	QMF SAMPLE TABLE
882	DSQ1STBB	Q	STAFF	TABLE	QMF SAMPLE TABLE

TABLE_CAT	TABLE_SCHEM	TABLE_NAME	COLUMN_NAME	DATA_TYPE	TYPE_NAME	COLUMN_LENGTH	BUFFER_LENGTH	DECIMAL_DIGITS	NUM_PREC_RADIX	NULLABLE	REMARKS	
1	DSQ1STBB	Q	STAFF	COMM	3	NUMERIC	7	9	2	10	1	EMI
2	DSQ1STBB	Q	STAFF	DEPT	5	SMALLINT	5	2	0	10	1	EMI
3	DSQ1STBB	Q	STAFF	ID	5	SMALLINT	5	2	0	10	0	EMI
4	DSQ1STBB	Q	STAFF	JOB	1	CHAR	5	5		NULL	1	EMI
5	DSQ1STBB	Q	STAFF	NAME	12	VARCHAR	9	9		NULL	1	EMI
6	DSQ1STBB	Q	STAFF	SALARY	3	NUMERIC	7	9	2	10	1	EMI
7	DSQ1STBB	Q	STAFF	YEARS	5	SMALLINT	5	2	0	10	1	NUI

Figure 5–11. Table Columns

Displaying Stored Procedures and Related Information

To display the stored procedures and/or stored procedure columns for the data source, select the **Stored Procedures** tab. When the tab is first selected, no tables are listed.

Displaying Stored Procedures

To display the stored procedures for the data source, do the following:

- From the **Stored Procedures** tab, click **Load Stored Procedures**. The system displays the **Load Stored Procedures -- Parameters** dialog box, as shown in Figure 5–12.




Figure 5–12. Load Tables -- Parameters

2. Specify criteria in the available fields to filter the stored procedures that will be returned.
3. Click **OK**. jDemo will load and display the stored procedures that meet the specified criteria, as shown in Figure 5–13.

Driver	Database	Tables	Stored Procedures			
	PROCEDURE_CAT	PROCEDURE_SCHEM	PROCEDURE_NAME	REMARKS	PROCEDURE_TYPE	
68	NULL	SYSPROC	DB2PROC1		1	
69	NULL	AI38DJP	DB2PROC2		1	
70	NULL	SYSPROC	DB2PROC2		1	
71	NULL	SYSPROC	DB2PROC3		1	
72	NULL	SYSPROC	DB2PROC4		1	
73	NULL	AI38PHV	DB2PROC5		1	
74	NULL	SYSPROC	DB2PROC5		1	
75	NULL	AI38PHV	DB2PROC6		1	
76	NULL	AI38DJP	DB2PROC9		1	
77	NULL	SYSPROC	DB2R0003		1	
78	NULL	AI38PHV	DB2PROC4		1	

Load stored procedures... Stored procedure: Show stored procedure columns...

Figure 5–13. jDemo -- Stored Procedures Tab

Displaying Stored Procedure Columns

To display the columns for a stored procedure, do the following:

1. Do one of the following to select a stored procedure:
 - If the information for the desired stored procedure has already been loaded and displayed (as shown in Figure 5–13 on page 5-11), then from the **Stored Procedure** tab, within the **PROCEDURE_NAME** column, simply select the desired stored procedure name. The row will be highlighted, as shown in Figure 5–14.

Or


- Within the **Stored Procedure** field, type the name of the stored procedure.

Driver	Database	Tables	Stored Procedures		
	PROCEDURE_CAT	PROCEDURE_SCHEM	PROCEDURE_NAME	REMARKS	PROCEDURE_TYPE
68	NULL	SYSPROC	DB2PROC1		1
69	NULL	AI38DJP	DB2PROC2		1
70	NULL	SYSPROC	DB2PROC2		1
71	NULL	SYSPROC	DB2PROC3		1
72	NULL	SYSPROC	DB2PROC4		1
73	NULL	AI38PHV	DB2PROC5		1
74	NULL	SYSPROC	DB2PROC5		1
75	NULL	AI38PHV	DB2PROC6		1
76	NULL	AI38DJP	DB2PROC9		1
77	NULL	SYSPROC	DB2R0003		1
78	NULL	AI38PHV	DB2R0004		1

Load stored procedures... Stored procedure: SYSPROC.DB2PROC1 Show stored procedure columns...

Figure 5–14. jDemo -- Stored Procedure Tab with Selected Stored Procedure

- Click **Show Stored Procedure Columns**. The system displays the **Load Stored Procedure Columns -- Parameters** dialog box, as shown in Figure 5–15.



The dialog box titled "Load Stored Procedure Columns -- Parameters" contains the following fields:

- Catalog: [Empty]
- Schema pattern: SYSPROC
- Stored procedure name pattern: DB2PROC1
- Column name pattern: [Empty]

Buttons: OK, Cancel, Help

Figure 5–15. Load Stored Procedure Columns -- Parameters

- Specify criteria in the available fields to filter the stored procedures that will be returned.



Note:

If a table was selected in step 1, then the **Schema Pattern** and **Stored Procedure Name Pattern** fields will be specified according to the selected stored procedure; otherwise, you must specify the criteria appropriately.

- Click **OK**. The system displays the stored procedure columns in the **Results** pane, as shown in Figure 5–16 on page 5-13.

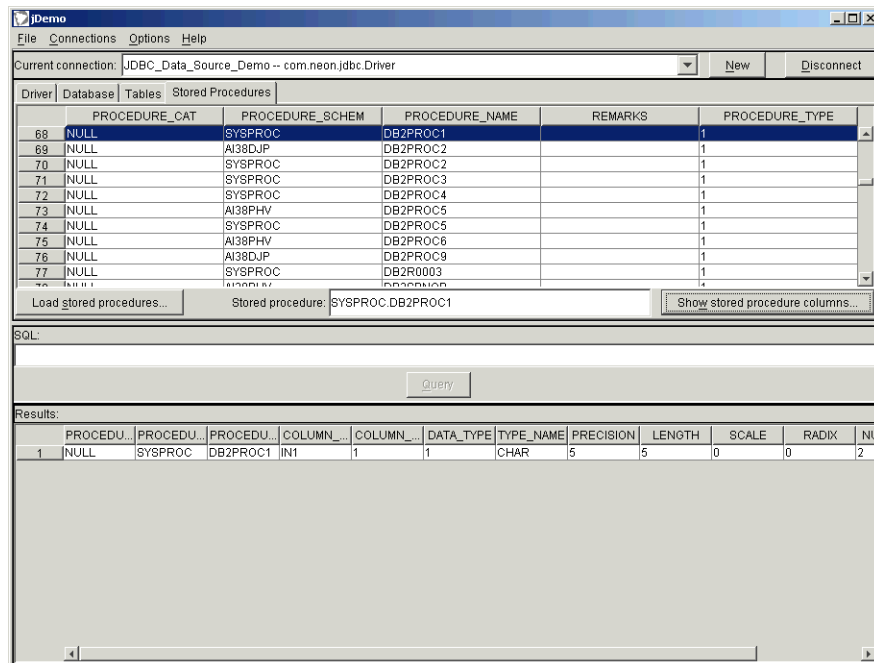


Figure 5–16. Stored Procedure Columns

Executing SQL

jDemo is a flexible tool for executing SQL queries.

Query Line Options

jDemo offers two query line options:

- **Single Line Queries:** (Default) By going to the **Options** menu and selecting **Single Line Query**, users can only enter single-lined SQL statements in the SQL text box.
- **Multi-Line Queries:** By going to the **Options** menu and selecting **Multi-Line Query**, users can only enter SQL statements spanning multiple lines in the SQL text box.

Executing a SQL Query

To execute a SQL query using jDemo, do the following:

1. In the **SQL** text box, type a SQL statement such as the following:

```
CALL * FROM Q.STAFF
```
2. Click **Query**. The query is executed and the results are displayed in the **Results** pane, as shown in Figure 5–17.

SQL:							
SELECT * FROM Q.STAFF							
Query							
Results:							
	ID	NAME	DEPT	JOB	YEARS	SALARY	COMM
1	10	SANDERS	25	MGR	7	18357.50	NULL
2	20	PERNAL	20	SALES	8	18171.25	612.45
3	30	MARENCHI	38	MGR	5	17506.75	NULL
4	40	O'BRIEN	38	SALES	6	18006.00	846.55
5	50	HANES	15	MGR	10	20659.80	NULL
6	60	QUIGLEY	38	SALES	NULL	16808.30	650.25
7	70	ROTHMAN	15	SALES	7	16502.83	1152.00
8	80	JAMES	20	CLERK	NULL	13504.60	128.20
9	90	KOONITZ	42	SALES	6	18001.75	1386.70
10	100	PLOTZ	42	MGR	7	18352.80	NULL
11	110	NGAN	15	CLERK	5	12508.20	206.60
12	120	NAUGHTON	38	CLERK	NULL	12954.75	180.00
13	130	YAMAGUCHI	42	CLERK	6	10505.90	75.60
14	140	FRAYE	51	MGR	6	21150.00	NULL

Figure 5–17. jDemo -- Executing a Query

Closing jDemo

To close jDemo, perform the following steps:

1. To disconnect from the data source, do one of the following:
 - In the jDemo window, go to the **Connections** menu and select **Remove Current Connection**.
 - Or**
 - In the jDemo window, click **Disconnect**.
2. Go to the **File** menu and select **Exit**.

CHAPTER 6: **Shadow Mainframe Adapter Client Trace Facility**

Debugging versions of Shadow Mainframe Adapter Client, part of the Mainframe Adapter Client component of the Shadow Connect product, include the Shadow Mainframe Adapter Client Trace Facility, which can produce trace log data. The output from this trace facility can be very useful in solving problems that can occur during installation and use.

This chapter will cover the following topics:

- Overview
- Shadow Mainframe Adapter Client Trace Facility for Windows
- Shadow Mainframe Adapter Client Trace Facility for UNIX
- Shadow Mainframe Adapter Client Trace Facility Options

Overview

The Shadow Mainframe Adapter Client Trace Facility traces selected Shadow events. It creates a trace file in a selected directory and records the traced events in that file. If a trace file by that name is already present, the trace facility adds any newly recorded events to it.

The format of the trace output data always includes a date, time, and pertinent information about each event. An example of trace output data is shown in Figure 6-1.

```

Fri Oct 01 22:19:30 1993  pcbColName      = 0x0a5f:759a
Fri Oct 01 22:19:30 1993  pfSqlType     = 0x0a5f:759e
Fri Oct 01 22:19:30 1993  pcbColDef     = 0x0a5f:75a0
Fri Oct 01 22:19:30 1993  pibScale      = 0x0a5f:7594
Fri Oct 01 22:19:30 1993  pfNullable    = 0x0a5f:759c
Fri Oct 01 22:19:30 1993  SQLDescribeCol exiting - return = SQL_SUCCESS(0)
Fri Oct 01 22:19:30 1993  szColName     = 'REMARKS'
Fri Oct 01 22:19:30 1993  *pcbColName   = 7
Fri Oct 01 22:19:30 1993  *pfSqlType    = SQL_CHAR(1)
Fri Oct 01 22:19:30 1993  *pcbColDef    = 64
Fri Oct 01 22:19:30 1993  *pibScale     = 9999
Fri Oct 01 22:19:30 1993  *pfNullable   = SQL_NULLABLE_UNKNOWN(2)
Fri Oct 01 22:19:30 1993  SQLFetch entered
Fri Oct 01 22:19:30 1993  lpstmt        = 0x095f:0000
Fri Oct 01 22:19:30 1993  internal error detected: file scodbcrc.c line 1228
rc = 0 from scclxlat
Fri Oct 01 22:19:30 1993  SQLFetch exiting - return = SQL_SUCCESS(0)
Fri Oct 01 22:19:30 1993  SQLGetData entered
Fri Oct 01 22:19:30 1993  lpstmt        = 0x095f:0000

```

Figure 6–1. Example Shadow Mainframe Adapter Client Trace Output Data

Shadow Mainframe Adapter Client Trace Facility for Windows



Note:

In order to enable tracing, you must have specified the debug version of Shadow Mainframe Adapter Client (*scjdl2ts.jar*) via the CLASSPATH environment variable. See “Step 2: Verify the Search Paths” on page 3-8 within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide.

The Shadow Mainframe Adapter Client Trace Facility offers the following features:

- **Shadow Mainframe Adapter Client Trace:** A tracing facility to create base traces and JDBC traces.
- **Shadow Mainframe Adapter Client Dynamic JDBC Trace:** Dynamic control of JDBC traces.

Shadow Mainframe Adapter Client Trace

You can create a base trace and/or JDBC trace with the Shadow Mainframe Adapter Client Trace Facility. You can enable and control the Shadow Mainframe Adapter Client Trace Facility by by appropriately setting the NEONTRACE environment variable.

When you use the NEONTRACE system environment variable to set the tracing options for the Shadow Mainframe Adapter Client Trace Facility, the trace starts at connection time. To set this variable, use the following form:

```
NEONTRACE=options JDBCLOG=jdbcpath LOG=path
```

Where:

options

Specifies the tracing options, delimited with spaces. Possible options are listed in “Shadow Mainframe Adapter Client Trace Facility Options” on page 6-11.

JDBCLOG

Creates a JDBC trace file.

jdbcpath

(If creating a JDBC trace) Specifies the full path name and file name of the JDBC trace file.

LOG

Creates a base trace file.

**Notes:**

- If the LOG keyword is present in the NEONTRACE environment variable, any other keywords must *precede* the LOG keyword.
- If you do not want to create a base trace file, set the LOG keyword to NULL as follows:

```
LOG=NULL
```

path

(If creating a base trace) Specifies the full pathname and file name of the base trace file.

Considerations

In addition, users creating traces with the Shadow Mainframe Adapter Client Trace Facility should do the following:

- If your base tracing options are not specified in the NEONTRACE environment variable setting, then any tracing options specified via the data source settings (either in a connection string or in a data source defined with the jConfig tool) will set the base tracing options.

**Note:**

Tracing options set via the data source settings will *only* affect the base tracing, not the JDBC tracing.

- Ensure that the CLASSPATH environment variable points to the debug jar file, which is named `scjd12ts.jar`. Also, ensure that it appears before any other Shadow Mainframe Adapter Client jar file on the CLASSPATH.



Doc Reference:

For more information about setting the CLASSPATH environment variable, see “Step 2: Verify the Search Paths” on page 3-8 within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide.

Examples

For example, to generate an INFO level JDBC trace to `C:\JDBCLOG.TXT` and a driver trace to `C:\NEONLOG.TXT`:

- **For Windows 95 or higher:** Add the following line to your `AUTOEXEC.BAT`:

```
NEONTRACE=INFO JDBCLOG=C:\jdbclog.txt LOG=C:\neonlog.txt
```

- **For Windows NT:** Open the **Control Panel**, select the **System** icon, and add the following line to the environment variable set:

```
NEONTRACE=INFO JDBCLOG=C:\jdbclog.txt LOG=C:\neonlog.txt
```

Shadow Mainframe Adapter Client Dynamic JDBC Trace

Dynamic JDBC tracing is supported with the `scjdlog` utility, which receives the tracing messages via TCP/IP and writes them to the log file on disk.

The `scjdlog` utility allows users to dynamically control the trace by supporting the commands shown in Table 6–2.

Table 6–1. scjdlog Utility Commands

Command	Description
<code>-p port-number</code>	Change the listening port of the <code>scjdlog</code> utility. The default value for the port number is 1201.
<code>-l NEONTRACE-options</code>	Change the tracing options for the JDBC dynamic trace. A trace option of NONE effectively “turns off” the JDBC dynamic trace.
<code>-f filename</code>	Change the file name of the JDBC trace log file. The default value for the file name is <code>neonj.txt</code> .
<code>-q</code>	Quit the <code>scjdlog</code> utility.

Enabling Shadow Mainframe Adapter Client Dynamic JDBC Tracing

You can perform dynamic JDBC tracing for Windows as follows:

1. Ensure that the CLASSPATH environment variable points to the debug jar file, which is named `scjdl2ts.jar`. Also, ensure that it appears before any other Shadow Mainframe Adapter Client jar file on the CLASSPATH.



Doc Reference:

For more information about setting the CLASSPATH environment variable, see “Step 2: Verify the Search Paths” on page 3-8 within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide.

2. Set an environment variable called NEONTRACE as follows:

```
NEONTRACE=options JDBCLOG=:
```

Where `options` consists of the options you initially select (see “Shadow Mainframe Adapter Client Trace Facility Options” on page 6-11), delimited with spaces, and the colon (“:”) in the JDBCLOG keyword specifies the use of dynamic tracing.

Optionally, you can also specify an IP address and a listening port, as follows:

```
NEONTRACE=options JDBCLOG=ip-address:port-number
```



Note:

If you do not want to capture a base trace, set the LOG keyword of the NEONTRACE environment variable to NULL as follows:

```
NEONTRACE=options JDBCLOG=: LOG=NULL
```

3. From a **Command Prompt** window, run `scjdllog.exe`. In a default installation, this file is located as follows:

```
C:\Program Files\IBM\Shadow\bin
```

The `scjdllog` utility is now waiting for tracing information to be written to it. Leave it running while you continue the required steps.



Note:

This window will become unusable because it is running the dynamic trace.

4. Open a second **Command Prompt** window and do the following to verify that the trace information is being captured:
 - a. Assuming a default installation, go to the following directory:

```
C:\Program Files\IBM\Shadow\bin
```
 - b. Verify that the size of the `neonj.txt` file is increasing, indicating that the trace information is being written to this file.
5. From the second **Command Prompt** window, you may issue commands to control the trace dynamically with the `scjdlog` utility. For a description of the commands supported by the `scjdlog` utility, see Table 6–1.

Examples:

- To generate a dynamic INFO level JDBC trace:

```
scjdlog -l INFO
```

- To stop writing the trace log:

```
scjdlog -l NONE
```

- To quit the `scjdlog` utility:

```
scjdlog -q
```

Shadow Mainframe Adapter Client Trace Facility for UNIX



Note:

In order to enable tracing, you must have specified the debug version of Shadow Mainframe Adapter Client (`scjd12ts.jar`) via the `CLASSPATH` environment variable. See “Step 3: Verify the Search Paths” on page 3-20 within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide.

The Shadow Mainframe Adapter Client Trace Facility offers the following features:

- **Shadow Mainframe Adapter Client Trace:** A tracing facility to create base traces and/or JDBC traces.
- **Shadow Mainframe Adapter Client Dynamic JDBC Trace:** Dynamic control of JDBC traces.

Shadow Mainframe Adapter Client Trace

You can create a base trace and/or JDBC trace with the Shadow Mainframe Adapter Client Trace Facility. You can enable and control the Shadow Mainframe Adapter Client Trace Facility by appropriately setting the NEONTRACE environment variable.

When you use the NEONTRACE environment variable to set the tracing options for the Shadow Mainframe Adapter Client Trace Facility, the trace starts at connection time. To set this variable, use the following form:

```
NEONTRACE=options JDBCLOG=/jdbcpath LOG=/odbcpath
```



Note:

When setting the NEONTRACE environment variable via the command line, you must enclose the settings in quotes.

Where:

options

Specifies the tracing options, delimited with spaces. Possible options are listed in “Shadow Mainframe Adapter Client Trace Facility Options” on page 6-11.

JDBCLOG

Creates a JDBC trace file.

jdbcpath

(If creating a JDBC trace) Specifies the full path name and file name for the JDBC trace file.

LOG

Creates a base trace file.



Notes:

- If the LOG keyword is present in the NEONTRACE environment variable, any other keywords must *precede* the LOG keyword.
- If you do not want to create a base trace file, set the LOG keyword to NULL as follows:

```
LOG=NULL
```

odbcpath

(If creating a base trace) Specifies the full path name and file name for the base tracing file.

Considerations

In addition, users creating traces with the Shadow Mainframe Adapter Client Trace Facility should do the following:

- If your base tracing options are not specified in the NEONTRACE environment variable setting, then any tracing options specified via the data source settings (either in a connection string or in a data source defined with the jConfig tool) will set the base tracing options.



Note:

Tracing options set via the data source settings will *only* affect the base tracing, not the JDBC tracing.

- Ensure that the CLASSPATH environment variable points to the debug jar file, which is named `scjd12ts.jar`. Also, ensure that it appears before any other Shadow Mainframe Adapter Client jar file on the CLASSPATH.



Doc Reference:

For more information about setting the CLASSPATH environment variable, see “Step 3: Verify the Search Paths” on page 3-20 within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide.

Example

For example, to generate an INFO level JDBC trace to `/tmp/jdbclog.txt` and a driver trace to `/tmp/neonlog.txt`, set the following environment variable:

```
NEONTRACE=INFO JDBCLOG=/tmp/jdbclog.txt LOG=/tmp/neonlog.txt
```

Shadow Mainframe Adapter Client Dynamic JDBC Trace

Dynamic JDBC tracing is supported with the `scjdlog` utility, which receives the tracing messages via TCP/IP and writes them to the log file on disk.

The `scjdlog` utility allows users to dynamically control the trace by supporting the commands shown in Table 6–2.

Table 6–2. scjdllog Utility Commands

Command	Description
-p port-number	Change the listening port of the scjdllog utility. The default value for the port number is 1201.
-l NEONTRACE-options	Change the tracing options for the JDBC dynamic trace. A trace option of NONE effectively “turns off” the JDBC dynamic trace.
-f filename	Change the file name of the JDBC trace log file. The default value for the file name is neonj.txt.
-q	Quit the scjdllog utility.

Enabling Shadow Mainframe Adapter Client Dynamic JDBC Tracing

You can perform dynamic JDBC tracing for UNIX as follows:

1. Ensure that the CLASSPATH environment variable points to the debug jar file, which is named `scjdl2ts.jar`. Also, ensure that it appears before any other Shadow Mainframe Adapter Client jar file on the CLASSPATH.



Doc Reference:

For more information about setting the CLASSPATH environment variable, see “Step 3: Verify the Search Paths” on page 3-20 within Chapter 3, “Installing Shadow Mainframe Adapter Client,” of this guide.

2. Set an environment variable called NEONTRACE as follows:

```
NEONTRACE=options JDBCLOG=:
```



Note:

When setting the NEONTRACE environment variable via the command line, you must enclose the settings in quotes.

Where `options` consists of the options you initially select (see “Shadow Mainframe Adapter Client Trace Facility Options” on page 6-11), delimited with spaces, and the colon (“:”) in the JDBCLOG keyword specifies the use of dynamic tracing.

Optionally, you can also specify an IP address and a listening port, as follows:

```
NEONTRACE=options JDBCLOG=ip-address:port-number
```

**Note:**

If you do not want to capture a base trace, set the LOG keyword of the NEONTRACE environment variable to NULL as follows:

```
NEONTRACE=options JDBCLOG=: LOG=NULL
```

3. From a telnet session, do the following to prepare for the dynamic trace:

- a. Log on to your UNIX system.
- b. Run scjdlog. In a default installation, this file is located as follows:

```
/home/IBM/Shadow/bin/
```

The scjdlog utility is now waiting for tracing information to be written to it. Leave it running while you continue the required steps.

**Note:**

This telnet session will become unusable because it is running the dynamic trace.

4. Open a second telnet session and do the following to verify that the trace information is being captured:

- a. Log on to your UNIX system.
- b. Go to the /home/IBM/Shadow/bin directory.
- c. Verify that the size of the neonj.txt file is increasing, indicating that the trace information is being written to this file.

5. From the second telnet session, you may issue commands to control the trace dynamically with the scjdlog utility. For a description of the commands supported by the scjdlog utility, see Table 6–2.

Examples:

- To generate a dynamic INFO level JDBC trace:

```
scjdlog -l INFO
```

- To stop writing the trace log:

```
scjdlog -l NONE
```

- To quit the scjdlog utility:

```
scjdlog -q
```


Shadow Mainframe Adapter Client Trace Facility Options

The Shadow Mainframe Adapter Client Tracing Facility offers the following options to control the trace:

- Severity level trigger.
- Tracing options.

Severity Level Trigger

You can specify the severity of the Mainframe Adapter Client messages to be traced. When a severity level is specified, all events at and above that severity level will be traced and recorded. The severity options are listed below from the highest to the lowest level:

- **NONE:** With this option, no messages are traced or recorded. No trace file is created or opened. This option effectively turns off the Shadow Mainframe Adapter Client Trace Facility.
- **FATAL:** This option, the highest severity level, causes only program termination messages to be traced and recorded.
- **SEVERE:** (Default) This option causes only events of a severe or higher level to be traced and recorded.
- **ERROR:** This option causes all error messages to be traced and recorded. These messages usually include a description of what went wrong, a record of where the error was detected, any relevant return code, and a detailed error text message.
- **WARNING:** This option allows any warning messages issued by Shadow Mainframe Adapter Client to be traced and recorded.
- **INFO:** This option allows all informational events to be traced, including SQL. Note that values passed to and returned from Shadow Mainframe Adapter Client are included here. This option will cause a large quantity of information to be traced and recorded and is not recommended for general use; however, this option can be useful for application debugging.
- **DETAIL:** This option, the lowest severity level, causes all events, including function call events, to be traced. This value will cause huge amounts of information to be traced and recorded, and it is not recommended for general use.

Tracing Options

Tracing options are available for both base traces and JDBC traces.

Base Tracing Options

You can specify additional tracing options for the base trace as follows:

- **BUFFER:** This option traces network buffers transmitted between the Mainframe Adapter Client and Mainframe Adapter Server.
- **COMMTRACE:** This option traces detail information on all TCP/IP and LU 6.2 calls.
- **DATA:** This option traces the EBCDIC to ASCII translation of data. This option causes all conversion operations to and from DB2, ODBC C, and ODBC SQL data types to be traced. In each case, the input and output data is displayed in hexadecimal and character format. In addition, information is provided that designates which row and column of the table is currently being processed.
- **ENVLIST:** This option will trace out all environment variable settings into the Mainframe Adapter Client trace after a connection request, regardless of whether the connection request was successful or unsuccessful.
- **FLUSH:** Use this option with STAY. It helps to ensure that messages are not lost if your machine fails.
- **LOCKTRACE:** This option enables tracing for locks (mutex) that Shadow Mainframe Adapter Client obtains and releases.
- **SQL:** This option causes all SQL statements passed to the host to be traced, even if the trace severity level is higher than INFO. It also causes the column information for SELECT statements to be displayed, including the column number, name, and data type. For SQL_NUMERIC and SQL_DECIMAL data types, the precision and scale values are also displayed. This option is automatically set when INFO or DETAIL severity levels are selected.
- **STAY:** This option keeps the trace file open while conducting tracing. This can improve performance if you are using a disk file to store trace messages. This is generally a good option when doing heavy tracing.
- **STORAGE:** This option causes all storage GET and FREE operations to be traced. The trace record shows the name of the file requesting or freeing the storage area and the line number within the file. The record also indicates the size of the data area being obtained or freed, the address of the data area, and other important information. The final storage report is generated in the Windows environment when the DLL containing the storage manager functions is unloaded.

- **THREADID:** (Only available as an environment variable setting) This option traces a task or thread ID associated with every trace row.

**Note:**

By default, the THREADID will always be included at the beginning of trace lines for traces in a multi-threaded environment.

- **UPDATE:** This option is for dynamic, live debugging.

JDBC Tracing Options

You can specify additional tracing options for the JDBC trace via the NEONTRACE environment variable as follows:

- **FLUSH:** This option helps to ensure that messages are not lost if your machine fails.
- **THREADID:** This option traces a task or thread ID associated with every trace row.

**Note:**

By default, the THREADID will always be included at the beginning of trace lines for traces in a multi-threaded environment.

CHAPTER 7: Accessing DB2

This chapter contains details about using DB2 stored procedures to access DB2 data using Shadow Mainframe Adapter Client for DB2, part of the Shadow Mainframe Adapter Client component of the Shadow Connect product.

This chapter will cover the following topics:

- Invoking a Stored Procedure
- Returning Results Sets
- Using SET CURRENT PATH
- Troubleshooting DB2 Stored Procedures
- Samples

Invoking a Stored Procedure

The syntax for invoking a DB2 stored procedure using Shadow Mainframe Adapter Client is as follows:

```
CALL SCHEMA.PROCEDURE-NAME( PARAMETER, PARAMETER, . . . )
```

Where:

SCHEMA

Specifies the schema name, as defined in the DB2 SYSIBM.SYSROUTINES table. The schema name specified depends on your version of DB2, as follows:

- **For DB2 v4.1 and v5.1:** This defaults to SYSPROC and must be specified.
- **For DB2 v6.1 or above:** Supports schema names other than SYSPROC, if desired.¹

PROCEDURE-NAME

Specifies the procedure name in the DB2 catalog.

PARAMETER

Specifies the parameters for the DB2 stored procedures.

This syntax matches the IBM SQL stored procedure naming conventions.

¹ The following schema names should not be used as they may cause conflicts: NEON*, DATA, VSAM, ADABAS, IMSDB, MAP, ACI, UTILITY, CICS, and IMS.

Returning Results Sets

Shadow Mainframe Adapter Client for DB2 supports returning zero or multiple result sets from an IBM DB2 stored procedure.

Coding Cursors in Return Result Sets

Certain issues apply when you are running a DB2 stored procedure that returns a result set. If a COMMIT is performed after a DB2 stored procedure is executed and before the result set is fetched, the result set will be lost. Since this is a DB2 restriction, there are several options available:

- Before calling the DB2 stored procedure, turn off auto-commit within the Mainframe Adapter Client application.
- Code WITH HOLD in your DB2 stored procedures cursor declarations. When coding DB2 stored procedures, the cursors used in stored procedures should be declared WITH HOLD in addition to WITH RETURN. Otherwise, any COMMIT (either issued by DB2 or by Shadow Mainframe Adapter Server) will destroy the result set(s). If WITH HOLD is not specified and if COMMIT_ON_RETURN is set to Y in SYSIBM.SYSPROCEDURES, each stored procedure will not appear to return any result sets, *even if it actually does*.

The following is an example of an EXEC SQL statement in a DB2 stored procedure:

```
EXEC SQL
  DECLARE C1 CURSOR WITH HOLD WITH RETURN
  FOR SELECT ID, NAME, DEPT, JOB,
           YEARS, SALARY, COMM
  FROM Q.STAFF
END-EXEC.
```

However, this has a drawback. Since the cursor is defined as WITH HOLD, a lock will remain on the table in order to preserve the cursor positioning until a ROLLBACK is issued. A ROLLBACK would need to be issued by the Mainframe Adapter Client application to free any locks. A COMMIT will *not* free all the locks.

- Set the Shadow Mainframe Adapter Server parameter AUTOCOMMITCALL to NO, as follows:

```
MODIFY PARM NAME(AUTOCOMMITCALL) VALUE(NO)
```

This will cause the Shadow Mainframe Adapter Server to *not* issue a COMMIT immediately after the DB2 stored procedure is executed. Instead, you will get a COMMIT (as long as you are running in auto-commit mode) at the point of the CLOSE-CURSOR after the result set has been fetched.

Retrieving Column Names

If you want to retrieve column names from any DB2 stored procedures that return result sets, the DESCSTAT DB2 installation parameter *must* be set to YES. (DESCSTAT is set in the DB2 DSNZPARM member at DB2 install time. DSNZPARM is created via the DB2 installation job DSNTIJUZ.)



Note:

If the DESCSTAT DB2 installation parameter is *not* set to YES, the result set column names will *always* be zero length strings. Some IBM documentation incorrectly implies that this is a BIND parameter. It is not a BIND parameter and must be set as part of the installation of DB2.

Using SET CURRENT PATH

Shadow Mainframe Adapter Client for DB2 supports the SET CURRENT PATH statement. In addition, the Shadow Mainframe Adapter Server RPCDEFAULTSCHEMA parameter allows users to set a default schema name.



Note:

Regardless of the value of RPCDEFAULTSCHEMA, the SET CURRENT PATH statement executed from Shadow Mainframe Adapter Client can, at any time, change the current path for all future SQL operations for that connection. The SET CURRENT PATH statement can be executed via Shadow Mainframe Adapter Client as any other dynamic SQL statement.

Default RPCDEFAULTSCHEMA Value

RPCDEFAULTSCHEMA defaults to NEON, so if any unqualified requests are issued, it is assumed that the request is calling a Shadow remote procedure call (RPC), and the Shadow Mainframe Adapter Server will attempt to execute it as such.

Thus, if RPCDEFAULTSCHEMA is left to the default value, to execute a DB2 stored procedure, you must invoke it fully qualified, as shown in the following example:

```
CALL SYSPROC.DB2PROC1('SALES')
```

Setting the *RPCDEFAULTSCHEMA* Value

If *RPCDEFAULTSCHEMA* is set to any value other than NEON, the following will be assumed:

- It will be assumed that the *RPCDEFAULTSCHEMA* value is the default schema name for all calls.
- If a *CALL* statement is issued, it is assumed the request is a DB2 stored procedure call.

Using the *SET CURRENT PATH* Statement

Each schema name specified on the *SET CURRENT PATH* statement can be any valid string up to eight bytes. The *SET CURRENT PATH* statement can be issued with or without an equal operator between the command and the schema name. In the following example, both statements are valid:

```
SET CURRENT PATH = SYSPROC
SET CURRENT PATH SYSPROC
```

The current path can be up to 254 bytes long, and specify multiple schema names all separated by a comma, as shown in the following example:

```
SET CURRENT PATH = SYSPROC, AI38AAS, AI38KPO
```

At any time, you can issue the *SET CURRENT PATH* statement to switch back to Shadow RPCs by setting the current path back to NEON. If the *SET CURRENT PATH* command is issued to specify NEON as the first schema, then the *SET CURRENT PATH* command is not even passed to DB2. Rather, Shadow Mainframe Adapter Client internally recognizes that all future unqualified calls are Shadow RPCs and will attempt to execute the call as if it were a Shadow RPC, regardless of what other schema names are in the current path. For example, if the following command is issued and then a call is executed, the call will be executed as a Shadow RPC without the call statement being sent to DB2:

```
SET CURRENT PATH = NEON, SYSPROC, AI38AAS
```

If you want to switch the default schema back to NEON and at the same time have fully qualified DB2 stored procedure calls using a specific path, then the following sequence of commands can be issued:

```
SET CURRENT PATH = SYSPROC, AI38AAS
SET CURRENT PATH = NEON
```

The above two commands will set the new SQL path in DB2 to SYSPROC, AI38AAS. It will then switch the Shadow default schema back to NEON. If a request is executed without a qualifier, it is assumed to be a Shadow RPC. However, if the call is a fully qualified DB2 stored procedure call such as the following, this stored procedure call is sent to DB2 and DB2 will search both SYSPROC and AI38AAS to satisfy this stored procedure request:

```
CALL SYSPROC.DB2PROC1 ('SALES')
```


**Doc Reference:**

For more information on the DB2 SET CURRENT PATH command or information on the SQL path, please reference the *DB2 UDB for OS/390 SQL Reference* manual.

Troubleshooting DB2 Stored Procedures

Be aware of the following situations that can arise when executing DB2 stored procedures:

- **0 SQLCODE:** It is common to get a 0 (zero) SQLCODE from a stored procedure that actually failed because it could not access the associated package (timestamp errors, etc.). The reason is that all of the DB2 operations attempted by the stored procedure failed and there is no indication why.

In addition, a 0 (zero) SQLCODE can also result from a DB2 stored procedure that is designed to return one or more results sets for at least two reasons:

- All of the EXEC SQL statements in the stored procedure are failing because the required package is not available.
- The cursors used to return the result sets were not declared with WITH HOLD and COMMIT_ON_RETURN is set to Y.
- **-204 SQLCODE:** A -204 SQLCODE indicates that the DB2 stored procedure name is unknown to DB2. Unlike Shadow's stored procedures, all IBM DB2 stored procedures must be defined in the DB2 catalog.
- **-440 SQLCODE:** A -440 SQLCODE indicates the parameters passed by the application do not match the parameter definitions in the DB2 catalog.
- **-480 SQLCODE:** A -480 SQLCODE indicates that the stored procedure created one or more results without using WITH HOLD cursors and a subsequent COMMIT issued by Shadow Mainframe Adapter Server destroyed those result sets before the DESCRIBE PROCEDURE was completed.
- **+466 SQLCODE:** All stored procedures that actually create result sets exit with a +466 SQLCODE unless the result set is destroyed by a COMMIT (as described in "Coding Cursors in Return Result Sets" on page 7-2). This is neither an error nor a warning. It just means the stored procedure successfully created at least one result set.

Samples

Various samples are available with Shadow Mainframe Adapter Client for DB2. These samples can be found at the following locations, assuming a default installation:

- **Windows:** C:\Program Files\IBM\Shadow\samples\JDBC
- **UNIX:** /home/IBM/Shadow/samples/JDBC

The following samples are offered:

- **db2proc1.java:** This is sample Java program to call the sample DB2 stored procedure DB2PROC1, returning a result set.
- **db2proc2.java:** This is sample Java program to call the sample DB2 stored procedure DB2PROC2, returning parameters.
- **db2proc4.java:** This is a sample Java program to call the sample DB2 stored procedure DB2PROC4, returning parameters and a result set.

Using the Sample Programs

To use the Java program samples, do the following:

1. If applicable, configure the “String url” with a valid JDBC connection string, either specifying the required keywords in the connection string or referencing an existing data source.
2. Change the CALL statement to a valid statement.
3. Compile the Java program to create a Java class.

Appendices

APPENDIX A: *Shadow Mainframe Adapter Client Keywords*

Shadow Mainframe Adapter Client, part of the Mainframe Adapter Client component of the Shadow Connect product, is controlled using certain keywords. These keywords are automatically given default values which can be changed depending on the Shadow function you are using.

This chapter covers the following topics:

- Setting a Keyword
- Keyword Descriptions
- Setting Keywords to Benefit Performance

Setting a Keyword

For details on how to set Shadow Mainframe Adapter Client keywords, see “Step 1: Configure the Data Source” on page 4-1 within Chapter 4, “Shadow Mainframe Adapter Client: Configuration,” of this guide.



Note:

In addition, certain application Mainframe Adapter Servers allow Shadow Mainframe Adapter Client keywords to be set via the application Mainframe Adapter Server graphical user interface.

Keyword Descriptions

Table A-1 on page A-2 lists the Shadow Mainframe Adapter Client keywords that can be set, along with a brief explanation of each and default settings.

Table A-1. Shadow Mainframe Adapter Client Keywords

Keyword	Description	Explanation	Default
ABCN	ADABAS Correlation Name Support	<p>This keyword specifies whether the Shadow Mainframe Adapter Client for ADABAS will support column name correlation IDs for this particular connection.</p> <p>Possible values are as follows:</p> <ul style="list-style-type: none"> • YES: Column name correlation IDs will be supported. • NO: Column name correlation IDs will not be supported. <p>Support for column name correlation IDs can also be enabled for all connections to a particular Shadow Mainframe Adapter Server by setting the Shadow Mainframe Adapter Server parameter, ADABASCORRELATIONIDS. The Shadow Mainframe Adapter Client keyword adds flexibility, allowing users to enable this support on an application by application basis.</p> <p>Note: Setting this keyword to YES may cause a conflict with earlier versions of the Shadow Mainframe Adapter Client for ADABAS, which accepted SQL syntax statements of the following form (note that there are no commas to separate the operands):</p> <pre>SELECT AA AB FROM EMPQA1</pre> <p>With this keyword set to NO (or on earlier versions of the Shadow Mainframe Adapter Client for ADABAS), this will select two columns, AA and AB. With this keyword set to YES, AB will be considered a correlation name for AA.</p> <p>To select two columns when this keyword is set to YES, commas must be used to separate the two column names as follows:</p> <pre>SELECT AA, AB FROM EMPQA1</pre>	NO
ACMT	Number of Active Statements	<p>This keyword is only applicable to the SQLGetInfo function called with an InfoType of SQL_ACTIVE_STATEMENTS, in which case the value will be as follows if this keyword is not explicitly set:</p> <ul style="list-style-type: none"> • 0: If either the AF (MS Access Compatibility) keyword (see “AF” on page A-2) is set or the ZEAS (Zero Active Statements) keyword (see “ZEAS” on page A-32) is set. A value of 0 (zero) indicates an unknown or unlimited value. • 50: If the normal cursor pool is being used. • 400: If the extended cursor pool is being used. See “EXCU” on page A-12. <p>However, this keyword can be set to a specific value.</p>	0
AF	MS Access Compatibility	<p>This keyword offers a solution to Microsoft Access compatibility issues. If this keyword is set to YES, all host decimal data values are returned to the Mainframe Adapter Client application without any trailing zeroes to the right of the decimal point. This keyword must be set to YES in order for Microsoft Access to work with DB2 decimal data. This keyword may be needed with any other product that uses the Microsoft Jet database engine, such as Visual Basic.</p>	NO
ALCD	Always Convert Dynamic SQL	This keyword should not be used!	YES

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
ALPO	Always Get Password from ODBC.INI/SHADOW.INI	This keyword controls whether the password will always be obtained from the SHADOW . INI file. If set to YES, the password will always be obtained from the SHADOW . INI file, and the password passed in the connection string will then be ignored. <i>Note:</i> See also “NOPW” on page A-19.	NO
ALUO	Always Get Userid from ODBC.INI/SHADOW.INI	This keyword controls whether the userid will always be obtained from the SHADOW . INI file. If set to YES, the userid will always be obtained from the SHADOW . INI file, and the userid passed in the connection string will then be ignored. <i>Note:</i> See also “NOUS” on page A-19.	NO
ALUS	Alternate (Secondary) Userid	This keyword is used to set the host secondary userid for Mainframe Adapter Client applications. This keyword must be eight or fewer characters long. If this keyword is set to a non-blank, non-null value, a SET CURRENT SQLID statement is issued after DSNALI OPEN processing is completed.	
APNA	Application Name	This keyword is the application name. The application name is sent to the host as part of the logon information. The application name is normally used to group SQL statements within a plan. If the application name is not set, all of the SQL associated with a plan will be considered to be part of one large group. If the SQL used with one plan must be divided into subgroups (for conversion to static SQL), the application name must be set.	
APPL	Application	This keyword must be set to ODBC (which is the default).	ODBC
APRO	Application Read-Only	(Read-only) This keyword is an internal API parameter set via the following: <ul style="list-style-type: none"> Below ODBC 3.0: SQLSetConnectOption function with an Attribute value of SQL_ACCESS_MODE. ODBC 3.0: SQLSetConnectionAttr function with an Attribute value of SQL_ATTR_ACCESS_MODE. 	
ASTM	Async Execution Timeout Value	This keyword controls the default wait time (in seconds) for an ODBC function that is executing asynchronously. This keyword must be an integer.	2

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
AT	Accessible Tables	<p>This keyword controls whether to restrict the list of tables for which the SQLTables function returns information. Many applications require information about the tables in the connected data source; however, in the case of DB2, this list of tables is unmanageably large unless you provide restrictions on the size. This keyword can be set as follows:</p> <ul style="list-style-type: none"> • YES: When the SQLTables function is called, only the following tables will be returned: <ul style="list-style-type: none"> – Tables matching the table filter (see “DP” on page A-11). – Tables for which the current authorization ID has SELECT authority. This includes those in PUBLIC and PUBLIC AT ALL LOCATIONS but not those where SELECT authority is granted by virtue of a secondary authorization ID. • NO: All tables matching the table filter (see “DP” on page A-11) will be returned. 	NO
AUST	Change Dynamic SQL to Auto-Static	<p>This keyword is used to control whether or not Shadow Connect should try to convert dynamic SQL to auto-static SQL.</p> <p>Auto-static SQL converts literals in SQL statements to parameter markers for better query matching in Shadow Connect. In addition, auto-static SQL offers dynamic SQL caching, which means that for SELECT queries, the cursor is kept open across commits to keep the prepared SELECT statement in order to run it statically. For INSERT, UPDATE, and DELETE queries, the statements are lost at commit; thus, in order to take advantage of auto-static SQL for these types of queries, auto-commit must be disabled within the Mainframe Adapter Client application code and commits must be less frequent to exploit the use of the static SQL statements.</p> <p>The following behavior will result from setting this keyword:</p> <ul style="list-style-type: none"> • YES: All dynamic SQL statements will be looked up in a lookaside buffer and converted to static SQL, if possible. • NO: Normal dynamic SQL processing will be performed. <p>Note: Dynamic SQL can only be converted to auto-static SQL if this keyword is set to YES and if the Shadow Mainframe Adapter Server supports auto-static SQL, as determined by the Shadow Mainframe Adapter Server parameter AUTOSTATICSQL.</p>	YES
AUSZ	Lookaside Buffer Size	<p>(Read-only) This keyword displays the value of the Shadow Mainframe Adapter Server LOOKASIDESIZE parameter on the host.</p>	
BIPA	Binary Passthrough	<p>This keyword controls whether or not host binary data should be returned to Mainframe Adapter Client applications without being converted to hexadecimal. This keyword can be set as follows:</p> <ul style="list-style-type: none"> • YES: Host binary data is passed through to Mainframe Adapter Client applications unchanged. • NO: Host binary data is converted to hexadecimal. 	NO
BYDB	Bypass Double Quote	<p>This keyword should be set to YES if double quotes should be bypassed. This keyword is provided to fix certain application bugs.</p>	NO

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
CALK	CALL Lock Value	This keyword controls the type of lock associated with CALL statements on the host, as follows: <ul style="list-style-type: none"> • EXCLUSIVE • SHARE • UPDATE • NONE: The host code will assume that CALL statements do not obtain any host database locks. 	EXCLUSIVE
CD	Convert Dynamic SQL	This keyword should <i>not</i> be used!	NO
CGFX	Cognos Impromptu Compatibility	This keyword can be set to YES to resolve certain problems with Cognos Impromptu. This keyword should not be set for any other reason.	YES
CLCT	Mainframe Adapter Client Certificate File	This keyword specifies the SSL Mainframe Adapter Client certificate file name.	
CLOR	Column Order Option	This keyword specifies how column names should be ordered when they are returned by catalog functions. Some functions explicitly specify an order; for functions that do not specify an order, this keyword is used by the catalog functions (SQLColumns). Possible values are as follows: <ul style="list-style-type: none"> • NAME • NUMBER 	NAME
CMNP	Confirm New Password	If this keyword is set to YES, the user will be asked to confirm the new password string if the new password is set in the PWD keyword.	NO
CNDG	Connection Dialog Type	This keyword specifies what type of connection dialog should be displayed, as follows: <ul style="list-style-type: none"> • DYNAMIC: Either the SIMPLE or the DETAIL dialog is displayed, depending on how much information is needed. • SIMPLE • DETAIL • NEVER This keyword can be set using the data source settings or using a connection string.	DYNAMIC

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
CNMD	Connection Mode	<p>This keyword controls the connection mode used by applications. The connection mode determines how long each physical connection (session or conversation) lasts and if SQL operations are blocked together. The possible connection modes are as follows:</p> <ul style="list-style-type: none"> • PERMANENT: Use a permanent connection and send each SQL operation standalone to the Mainframe Adapter Server. • BLOCK: The session is permanent; however, SQL operations are blocked and sent together. • TRANSACTION: Each SQL operation is sent standalone but the session is terminated at the end of each Logical Unit Of Work (LUOW). • TRANSBLOCK: The session is terminated at the end of each LUOW, and SQL operations are blocked and sent together. • MESSAGE: SQL operations are blocked and sent together using messages. No session is ever maintained. 	PERMANENT
CNNA	Connection Name	This keyword is used to specify the connection name. The use of the connection name is application specific. The name can be up to eight (8) bytes long; however, the application may or may not use all of the bytes. The connection name is padded on the right with blanks.	
CNPL		This keyword should <i>not</i> be used!	OFF
CNTM	Connection Timeout Value	<p>This keyword controls how many seconds the Mainframe Adapter Client will wait for a connection to the host Mainframe Adapter Server to complete. If this keyword is set to 0 (zero), then the system default value will be used. Otherwise, the specified value will be used. This value should never be negative but can be 0 (zero). This value has no effect in UDP messaging.</p> <p><i>Note:</i> The use of this keyword may cause some unpredictable results in many environments. This option is only available in the TCP/IP and the MQ link types (see “LINK” on page A-16), but it is applicable to both the permanent and the messaging modes (see “CNMD” on page A-6).</p>	0

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
COFX	COUNT() Fix Type	<p>This keyword controls how COUNT(column name) SQL functions are modified. Microsoft Access uses the COUNT(column name) function two ways, both of which are wrong:</p> <ol style="list-style-type: none"> 1. In some cases, COUNT(column name) really means COUNT (DISTINCT column name). 2. In other cases, COUNT(column name) means COUNT(*). <p>This keyword can be set as follows:</p> <ul style="list-style-type: none"> • DISTINCT: The DISTINCT keyword will be inserted (as in case 1, above). • ASTERISK: The column name will be replaced with a "*" (as in case 2, above). • NONE: No changes will be made. <p>Note: COUNT(column name) functions are never changed unless the MS Access Compatibility keyword is set to YES. See "AF" on page A-2.</p>	DISTINCT
COID	Correlation ID Support	<p>This keyword controls whether or not SQLGetInfo should return information about correlation IDs. When this keyword is set, SQLGetInfo will behave as follows:</p> <ul style="list-style-type: none"> • YES: SQLGetInfo will return information about correlation IDs. • NO: No information about correlation IDs will be returned. This keyword must be set to NO to enable some tools to work. 	YES
CPFX	Catalog Prefix	<p>This keyword specifies the PROCEDURE QUALIFIER (in ODBC 2.0 terminology) or the PROCEDURE CAT (in ODBC 3.0 terminology). Possible values are as follows:</p> <ul style="list-style-type: none"> • SYSIBM: Accept the default value to use the standard system catalogs (SYSIBM). • SHADOW: Specify SHADOW only if you are using the Shadow-optimized catalog feature, which requires DB2. SHADOW should represent the table owner (in DB2 terminology, the authorization ID) for a set of DB2 tables that are optimized to support catalog queries. <p>Note: Shadow-optimized catalogs are considered essential for smooth operation of pre-packaged applications with the Shadow product.</p> <ul style="list-style-type: none"> • SDBMAP: Specify SDBMAP to return metadata related to data maps in the Shadow Data Mapping Facility (DMF) for CICS and IMS pseudo-stored procedures. 	SYSIBM

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
CRBH	Cursor Commit/ Rollback Behavior	This keyword controls what value to return for the cursor commit and rollback behavior SQLGetInfo request, as follows: <ul style="list-style-type: none"> DELETE: This will close cursors and delete prepared statements. To use the cursor again, the application must re-prepare and re-execute the statement. CLOSE: This will only close cursors. For prepared statements, the application can call SQLExecute on the statement without calling SQLPrepare again. PRESERVE 	CLOSE
CRIN	Create Table Index Automatically	This keyword controls whether an index is automatically created when you create a table with a primary key or unique constraint. This index enforces the uniqueness. To disable this option, set the value to NO.	YES
CVNL	Convert Nulls to Blanks	This keyword controls whether or not nulls (zero bytes) in character string data should be returned from the Mainframe Adapter Server to the Mainframe Adapter Client, or should be converted to blanks. Both fixed length and variable length character data are affected by this keyword. This keyword can be set as follows: <ul style="list-style-type: none"> YES: Nulls will be converted to blanks. NO: Nulls will not be converted. 	NO
CVTS	Convert Timestamps	This keyword controls whether or not timestamp values should be converted to other types, as follows: <ul style="list-style-type: none"> YES: Timestamps that appear to be dates will be converted to dates and timestamps that appear to be times will be converted to times. These conversions are required to circumvent bugs in several products including MS Access and Crystal Reports. NO: Timestamps will not be modified. 	NO
D2VR	DB2 Version String	This keyword can overwrite the host returned DB2 version string data. This string should be in the following format: x.y.z Where x, y and z are all numeric digits, as follows: <ul style="list-style-type: none"> x (required) is the DB2 version byte. y (required) is the DB2 modification level. z (optional) is the DB2 release level. If the release level is not set, then 0 is assumed. <p>Examples:</p> <ul style="list-style-type: none"> 2.3 is for DB2 version 2, modification level 3. 4.1.1 is for DB2 version 4, modification level 1, and release 1. 	
DACN	Disable SQLCancel Tracing	This keyword controls whether to disable SQLCancel tracing. SQLCancel creates a separate connection to the host Mainframe Adapter Server, and when the SQLCancel operation is completed, this connection is freed. Setting this keyword to YES prevents all tracing of the operations used in creating and terminating the new connection. This is done to avoid cluttering the trace file.	YES

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
DAOP	SQLDescribeParam Support	This keyword controls if Shadow Mainframe Adapter Client should support SQLDescribeParam in a limited form. This keyword can be set to the following values: <ul style="list-style-type: none"> • Mainframe Adapter Server: The Shadow Mainframe Adapter Server code will decide whether this support is enabled. • YES: SQLDescribeParam is supported to a limited extent. • NO: SQLDescribeParam is not supported at all. 	Mainframe Adapter Server
DBD	Default Database	This keyword specifies the database where the table will be created when issuing CREATE statements.	
DBMD	DBCS Mode	(Only for DBCS or GRAPHIC data) This keyword is used to specify how DBCS (double byte character set) data is handled: <ul style="list-style-type: none"> • DEFAULT: (Default) When processing DBCS or GRAPHIC data, leaving the default value of DEFAULT will cause an error. • BINARY: All DBCS data will be treated as binary data. • CHAR: All DBCS data will be treated as character data, even if it is stored in GRAPHIC columns. • GRAPHIC: It is assumed that the DBCS data will be treated as SQL_GRAPHIC, SQL_VARGRAPHIC, or SQL_LONGVARGRAPHIC SQL types. <p>DBCS data will still be converted from host formats to Mainframe Adapter Client formats; however, in BINARY mode, data will be described as binary even though it is not stored in binary columns on the host.</p> <p><i>Note:</i> This keyword only applies if you are using DBCS or GRAPHIC data, in which case you must choose how to handle the DBCS data, as attempting to use DEFAULT will result in an error. The three choices are to treat it as binary data (BINARY), treat it as character data (CHAR), or treat it as graphic data (GRAPHIC).</p>	DEFAULT
DBQO	DBCS Remove Blanks	This keyword controls whether or not blanks are removed from within quoted strings for double byte languages (such as Chinese, Japanese, and Korean). Possible values are as follows: <ul style="list-style-type: none"> • YES: Blanks that are within quoted strings in SQL being sent to the host will be removed. • NO: Blanks that are inside quoted strings in SQL being sent to the host will not be changed. 	YES
DBTY	DBMS Type	This keyword provides SQL access to databases by specifying the means by which users will be accessing data. Supported values are as follows: <ul style="list-style-type: none"> • DB2 • Oracle • ADABAS • VSAM • VSAMCICS • DATA 	DB2

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
DCCM	Decimal Point is Comma	This keyword controls if the decimal point is a comma or not, as follows: <ul style="list-style-type: none"> • YES: A comma is considered to be the decimal point. This keyword should only be set to YES if the host uses a comma as the decimal point. This keyword will cause commas to be converted to periods before data is passed to the auto-static SQL conversion facility. • NO: A period is considered to be the decimal point. 	NO
DENU	Decimal as Numeric	This keyword determines whether DECIMAL and NUMERIC columns should be reported as DECIMAL columns or NUMERIC columns (because both types of columns are treated as the same by the Shadow Mainframe Adapter Client). Possible values are as follows: <ul style="list-style-type: none"> • YES: All DECIMAL and NUMERIC columns will be reported as NUMERIC. • NO: All DECIMAL and NUMERIC columns will be reported as DECIMAL. 	NO
Description	Description	This optional keyword can contain a string value describing this data source (user-defined).	
DFCL	Default Column Names	This keyword controls whether or not Shadow Mainframe Adapter Client will assign default column names (of the form COLnnn, where nnn is the column number) when the DBMS does not return names for the columns. If set to YES, default column names will be assigned.	NO
DFPR	Defer Prepare for CALLs	This keyword controls whether or not to defer prepare of CALL statements. Possible values are as follows: <ul style="list-style-type: none"> • YES: Prepare is always deferred for CALL statements. If this keyword is set to YES, prepare is deferred for CALLs that do not have parameter markers. • NO: Prepare is only deferred for CALLs that have one or more parameter markers. In other words, prepare is always deferred for CALLs that have parameter markers. 	NO
DFSC	Default Schema	This keyword controls the default schema for stored procedures with an implicit schema name. Possible values are as follows: <ul style="list-style-type: none"> • 0: The procedures are executed as Shadow RPCs. • 1: Shadow Mainframe Adapter Server uses the RPCDEFAULTSCHEMA parameter value as the schema name. • Specific schema name: If the value is a specific schema name, it is used as the default schema via the SET CURRENT PATH setting each time a connection is established. <p>Note: This keyword is only supported at a Shadow Mainframe Adapter Server functional level of 4 or higher (see “FULV” on page A-13).</p>	0

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
DIPO	Disable Prepare/ Open Optimization	<p>This keyword controls whether to disable the prepare/open optimization. Possible values are as follows:</p> <ul style="list-style-type: none"> • YES: Disables the prepare/open optimization feature, consequently adding a network round-trip (for the SQLExecute). The RDO layer on top of VB issues two SQLPrepares and only one SQLExecute for each select statement. This causes two result sets to be returned for every query. <p>For those Mainframe Adapter Client applications that expect the SQLPrepare to be executed as soon as it requested (for example, applications that require the metadata in the SQLCA), this keyword should be set to YES.</p> <ul style="list-style-type: none"> • NO: Enables prepare/open optimization, offering support for deferred prepare, meaning that the SQLPrepare and SQLExecute can be combined into single call that is sent at SQLExecute time. Prepare/open optimization minimizes network flow and improves performance. 	NO
DP	Table Owner Filter	<p>This keyword specifies the table filter to use for the SQLTables function, which returns a list of available tables in the remote database. This field, along with the accessible tables feature (see “AT” on page A-4), is used to restrict the information returned by the SQLTables function. When the application requests a list of tables, Shadow Mainframe Adapter Client constructs this list based on the value of the table filter. Only tables that fall under the first-level name (userid) are returned.</p> <p>This filter can include the two SQL wildcard characters “%” and “_”, as follows:</p> <ul style="list-style-type: none"> • “%” substitutes for zero or more characters. • “_” substitutes for exactly one character. <p>The table filter defaults to the userid value, which means Shadow Mainframe Adapter Client only returns tables that you own or created under your userid.</p> <p>Note: More than one table filter can be specified by comma delimiting the entries.</p>	<USERID>
DPSO	Duplicate Socket Descriptor	<p>(UNIX only) This keyword is a non-negative integer that is used to return a new and unused file descriptor for the connecting socket. The new socket number will be a descriptor number that is either equal to or greater than this number. The old descriptor will be closed after the replacement. This keyword is only enabled for the UNIX platforms and is designed to allow C run-time applications to get around the 255 maximum open file handle limitation in studio library. For example, a good value to set this keyword to is 256.</p>	0
DRIVER	Driver	This keyword is used for DSN-less connections to specify the name of the Shadow Mainframe Adapter Client driver being used.	

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
DRPM		This keyword controls whether or not Shadow Mainframe Adapter Client should prompt for missing logon information. This keyword can have one of the following values: <ul style="list-style-type: none"> • NOPROMPT: (Default) SQL_DRIVER_NOPROMPT. • COMPLETE: SQL_DRIVER_COMPLETE. • PROMPT: SQL_DRIVER_PROMPT. 	NOPROMPT
DSFL	DTS Plan File	This keyword should <i>not</i> be used!	
DSN	Data Source Name	This keyword specifies the name of the data source.	
DTCH	Convert DB2 Date to CHAR	When this keyword is set to YES, the DB2 date keyword will be converted to SQL_CHAR.	NO
DTFM	Date Format	This keyword is used to specify how ODBC dates are converted to character strings, as follows: <ul style="list-style-type: none"> • ODBC: The standard ODBC/ISO format yyyy-mm-dd will be used. • UK: The dd-mm-yyyy format will be used. • EUR: The dd.mm.yyyy format will be used. • USA: The mm/dd/yyyy format will be used. • USA2: The mm/dd/yy format will be used. 	ODBC
EXCC	SQLExecute Connection Check	This keyword is used to force a connection check during SQLExecute processing. Possible values are as follows: <ul style="list-style-type: none"> • YES: SQLExecute will always return an error condition if the connection has been reset or shutdown by the Mainframe Adapter Server. • NO: (Default) If prepare/open optimization is in effect (see “DIPO” on page A-11), SQLExecute will bypass the connection check and return SQL_SUCCESS. This keyword can be used by Shadow Mainframe Adapter Client users when prepare/open optimization is in effect to detect a stale connection during <code>statement.executeQuery()</code> processing instead of result set processing.	NO
EXCU	Extended Cursor Pool	This keyword controls whether or not the extended cursor pool should be used. Possible values are as follows: <ul style="list-style-type: none"> • YES: The extended cursor pool will be used. The extended cursor pool is much larger than the standard cursor pool; however, the extended cursor pool cannot be used with auto-static SQL (see “AUST” on page A-4). • NO: The normal cursor pool will be used. 	NO

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
FALG	Fast Logon	This keyword controls whether or not fast logon should be used. Possible values are as follows: <ul style="list-style-type: none"> • YES: Logon processing will be done with a minimum of network I/O. This is only possible if the host Mainframe Adapter Server supports fast logon. Failures will result if fast logon is requested and the host Mainframe Adapter Server does not support it. The ability of the Mainframe Adapter Server to handle fast logon must be verified before this option is used. • NO: Normal logon processing will be used 	NO
FEDL	Unload Communication DLLs	This keyword is used to determine whether the communication DLL (i.e., WINSOCK.DLL or CPIC.DLL) should be unloaded after an application disconnects. This keyword can be specified as follows: <ul style="list-style-type: none"> • YES: Shadow Connect always unloads the communication DLL after the application disconnects. • NO: Unloading the communication DLL while the communication service is still active will cause errors. Setting this keyword to NO will prevent the DLL from being unloaded by Shadow Mainframe Adapter Client. 	YES
FULV	Function Level Value	(Read-only) This keyword is set at connection time based on the maintenance level of the Shadow Sever code.	
FXFL	Fix Floating Point Rounding	This keyword is available to fix certain floating point rounding errors. These errors are typically seen when a floating point number is converted to a decimal or integer value. Possible values are as follows: <ul style="list-style-type: none"> • YES: Floating point values will be adjusted before they are converted. • NO: No adjustments will be made to floating point values. 	NO
FXIS	Fix INSERT Statements	This keyword controls whether or not the VALUES clause of each INSERT statement should be scanned for dates, times, and timestamps, and whether or not quotes should be added. Possible values are as follows: <ul style="list-style-type: none"> • YES: Quotes will be added to dates, times, and timestamps as needed. • NO: INSERT statements will be not modified. 	NO
FXOR	Fix Oracle 8.0.5 Bug	This field is used to fix the Oracle 8.0.5 bug.	NO
GAJI	GAIJI Extension Support	When set to YES, this keyword enables GAIJI extension support in a codepage.	YES
GATB	GAIJI Extension Table Name	This keyword sets the GAIJI extension table name. There is no default name. If this keyword is left blank, a Mainframe Adapter Server-side default table name will be used. This keyword has a maximum size of 4 characters.	

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
GHDS	GRAPHIC Data Processing	This keyword should be set to YES when inserting DBCS data into DB2 GRAPHIC columns. When set to YES, Shadow will determine whether a CHAR parameter is actually a GRAPHIC parameter so that Shadow can handle it correctly.	Conditional
GXSR	GRAPHIC Literal Processing	This keyword controls if some string literals should be converted to parameters. This conversion is needed for some applications operating in DBCS modes. If this keyword is set to YES, then all string literals not preceded by “G” or “g” will be converted to parameter markers. <i>Note:</i> This keyword is ignored if Shadow Mainframe Adapter Client is not in a DBCS environment.	NO
HD	Retain Cursor	This keyword controls how DB2 COMMIT operations affect cursors in the data source. Possible values are as follows: <ul style="list-style-type: none"> • YES: Cursors will be preserved in the same position as before the COMMIT operation took place. This enables the application to execute or fetch without preparing the statement again. <i>Note:</i> Setting this keyword to YES also causes Shadow Mainframe Adapter Server to issue COMMITs after every SELECT statement as long as auto-commit is enabled by the application. <ul style="list-style-type: none"> • NO: The application will close and delete cursors after a COMMIT operation, meaning you must prepare and execute the next statement from scratch. 	YES
HODB	Host Debugging Values	This keyword is used to control the debugging of host programs (generally stored procedures). The host program language must be specified correctly so that the program can be invoked with the correct debugging options. The supported values are as follows: <ul style="list-style-type: none"> • NONE • COBOL • PLI • C • C++ 	NONE
HOST	Host Name	This keyword specifies the host name or IP address for the connection. The value can contain a hostname string or a TCP/IP address in dot-notation format. Failover Shadow Mainframe Adapter Servers can be specified by listing host/port combinations as follows: HOST=Host1,Host2,Host3;PORT=1200,1210,1500 Based on the host/port combinations specified, Shadow Mainframe Adapter Client would try to connect to Host1:1200, Host2:1210, and Host3:1500, in that order. Specified host names can be the same host. Note that this syntax should <i>not</i> be used to specify Shadow Mainframe Adapter Server group directors.	

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
IDQO	Identifier Quote Option	This keyword is used to show what identifier quote character should be returned to the application using SQLGetInfo with SQL_IDENTIFIER_QUOTE_CHAR. This keyword is needed to fix certain application bugs. Possible values are as follows: <ul style="list-style-type: none"> • BLANK • SINGLE • DOUBLE 	BLANK
IGHI	Ignore High Bound Column Errors	This keyword controls whether or not SQLFetch should ignore high bound column errors. Possible values are as follows: <ul style="list-style-type: none"> • YES: SQLFetch will ignore column binding errors. This keyword must be set to YES to enable some tools (such as Microsoft Excel) to work. • NO: Column binding errors will be handled normally. 	NO
IGUN	Ignore Underscore Characters	This keyword controls whether or not the underscore character (“_”) should be considered as a wildcard or a regular character. This keyword is used to solve problems in which a table or procedure name actually has an underscore in the name and the underscore is misinterpreted as a wildcard. This keyword can be set as follows: <ul style="list-style-type: none"> • YES: Underscore will be handled as a normal byte. • NO: Underscore will be treated as a wildcard character that matches any other single byte. 	NO
KEQU	Keep Literal Quotes	This keyword is used to keep quotes around and embedded in string literals, as follows: <ul style="list-style-type: none"> • YES: Quotes around and embedded in string literals will be retained. • NO: Quotes will be removed from string literals. <p><i>Note:</i> This only applies to string literals passed to stored procedures, including MDI remote stored procedures (RSPs).</p>	NO
LAFX	Lotus Approach Compatibility	This keyword can be set to YES to resolve certain problems with Lotus Approach. This keyword should not be set for any other reason.	NO
LBCB	Convert LONG VAR Data to LOB	This keyword allows any SQL_LONGVARCHAR or SQL_LONGVARBINARY data to be treated as LOB data (either CLOB or BLOB).	YES
LGFX	Long Data Fix	This keyword controls whether or not a large number should be returned for the length and precision of LONG VARCHAR keywords, rather than the actual length and precision. This fix is needed to resolve certain problems in GetChunk processing. This keyword can be set as follows: <ul style="list-style-type: none"> • YES: A large number will be returned for the length and precision of LONG VARCHAR keywords. • NO: The actual values will be returned for the length and precision of LONG VARCHAR keywords. 	NO

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
LGID	Language ID	<p>This keyword is used to specify the language to be used. The possible values are as follows:</p> <ul style="list-style-type: none"> • ARB: Arabic • CHS: Simplified Chinese • CHT: Traditional Chinese • DAN: Danish • DFT: Default • DEU: German • ENC: U.S. English Compat • ENG: U.K. English • ENU: U.S. English • ESN: Modern Spanish • ESP: Castilian Spanish • FIN: Finish • FRA: French • FRC: Canadian French • ISL: Icelandic • ITA: Italian • JPL: Japanese • JPY: Japanese Latin Ext • KOR: Korean • MDI: Micro Decisionware • NLD: Dutch • NOR: Norwegian • PPS: PeopleSoft English • PTG: Portuguese • SVE: Swedish • TUR: Turkish 	ENU
LGMG	Return Logon Messages	<p>This keyword controls whether or not non-error logon messages should be returned to Mainframe Adapter Client applications. Possible values are as follows:</p> <ul style="list-style-type: none"> • YES: If a message is returned to the application, the return code will be set to SQL_SUCCESS_WITH_INFO and the error message text will be available using the Solderer function (or some higher level API). • NO: If a message is returned to the application, it will be displayed using a dialog box. 	NO
LGPA	Convert Strings to Params	<p>This keyword controls whether or not long strings should be converted to LONG VARCHAR parameters. This conversion is needed because some applications produce literals that are longer than can be handled by the host database. Possible values are as follows:</p> <ul style="list-style-type: none"> • YES: All long strings will be converted to parameter markers. • NO: SQL strings will not be scanned for long strings. 	NO
LINK	Link Type	<p>This keyword specifies the connection information. Supported values are as follows:</p> <ul style="list-style-type: none"> • TCPIP • LU62 • MQSeries 	TCPIP
LKES	Add ESCAPE Option	<p>This keyword causes an ESCAPE clause to be added after each LIKE clause if the following conditions are met:</p> <ul style="list-style-type: none"> • This keyword is set to YES. • The application did not provide its own ESCAPE clause. • The LIKE clause contains one or more characters that need to be escaped. 	YES
LKTH	Use Thread Locks	(Read-only) This keyword is set depending on whether thread-safe mode is being utilized.	
LNDN	LAN Domain Name	(Read-only) This is an internal, read-only keyword that may not be set by the user.	

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
LNID	LAN Userid	(Read-only) This is an internal, read-only keyword that may not be set by the user.	
LVCT	Catalog Level Value	(Read-only) This keyword value is negotiated internally between Shadow Mainframe Adapter Client and the Shadow Mainframe Adapter Server and represents the level of support between the two components.	
LZCM	LZ Compression	If this keyword is set to YES and the Shadow Mainframe Adapter Server supports this feature, a variant of Lempel-Ziv compression will be used to compress all data flow between Shadow Mainframe Adapter Client and the Shadow Mainframe Adapter Server.	YES
MDBO	MDI Text/ Keywords	This keyword controls if variable text and keywords can be used together with MDI RSPs, as follows: <ul style="list-style-type: none"> • YES: Variable text can be used with MDI keywords. • NO: Variable text can not be used with keywords. 	NO
MDQO	MDI Keep Quotes	This keyword controls if quotes should be retained around MDI keyword values. Possible values are as follows: <ul style="list-style-type: none"> • YES: Quotes will be included in keyword values and keyword value lengths. • NO: Quotes will be stripped from MDI keyword values. 	NO
MGDG	Message Digest	This keyword is not used.	
MGTY	Message Type	This keyword controls the type of messages used for message oriented connections. It is only used if the messaging connection mode is employed (see “CNMD” on page A-6). Supported values are as follows: <ul style="list-style-type: none"> • NETWORK: A TCP/IP or LU 6.2 session or conversation is created for each message. • UDP: When the link type is set to TCP/IP (see “LINK” on page A-16), users can also choose the UDP message type as the messaging protocol. Using UDP datagrams for messaging usually incurs lower network overhead; however, UDP communication is not reliable and is discouraged when connecting over a WAN. • HTTP: HTTP and HTTPS methods will tunnel the communication session inside the HTTP protocol stream to allow the session to be established over firewalls and HTTP proxies. • HTTPS: HTTPS is the Netscape secured HTTP protocol which implements SSL and can be used when a secured channel between the two endpoints is required. 	NETWORK
MR	Maximum Rows	This keyword can be used to specify an integer value to which to limit the number of rows returned from a single query. If no value is entered, no restriction is placed on the number of rows returned.	0

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
MSMT	MTS Security SID Type	This keyword specifies the security SID option when using the MS Transaction Mainframe Adapter Server based network authentication/single signon model. Possible values are as follows: <ul style="list-style-type: none"> • Direct Caller SID • Direct Creator SID • Original Caller SID • Original Creator SID 	Direct Caller SID
MT	Enable Multitasking	This keyword can be set to YES to enable the multitasking capabilities of Shadow Connect. This feature will keep users' screens from locking up while calls are pending. <i>Note:</i> This feature should be used with extreme caution! Many applications will not operate correctly with this feature activated.	NO
MUNT	MultiNet TCP/IP Compatibility	This keyword can be set to YES to support applications using the Multi-Net TCP/IP product for host access.	NO
MXBU	Maximum Buffer Size	This keyword sets the Shadow Mainframe Adapter Client maximum communication buffer size for all data exchanges. The default for this keyword is 0 which currently sets the maximum buffer size to 100,000 bytes (32K bytes in 16-bit Windows). <i>Note:</i> This value will be used to negotiate with the Shadow Mainframe Adapter Server side limit. The actual runtime buffer size could be smaller but not larger.	0
MXDG	Floating to Character Digits	This keyword controls the maximum number of digits that should be generated for a floating point to character conversion. This keyword should only be specified if the default value is not acceptable.	6
MXRT	UDP Maximum Retry Attempts	This keyword controls the maximum number of UDP retries for an unacknowledged request. This value must be a non-negative number. If this keyword is set to 0 (zero), the UDP request retry feature is disabled.	10
NEONTRACE	NEONTRACE	This keyword can be used to control the Shadow Mainframe Adapter Client Trace Facility tracing options. For more information, see Chapter 6, "Shadow Mainframe Adapter Client Trace Facility."	
NBIO	Non-Blocking Network I/O	This keyword controls whether or not TCP/IP network I/O operations should be blocking. Possible values are as follows: <ul style="list-style-type: none"> • YES: Non-blocking network I/O operations will be used, and the Windows message queue will be used to wait for I/O operation completion. • NO: Blocking network I/O operations will be used. 	NO
NLGR	German NLS Support	This keyword can be set to YES to resolve certain problems handling SQL in the German language environment. This keyword should not be set for any other reason.	NO

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
NO3D	Disable CTL3D	This keyword is used to indicate whether or not CTL3D processing will occur. Setting this option to YES will prevent the use of the CTL3D DLL with the connection dialog. This option must be set to YES for some applications if an old copy of CTL3D is installed.	NO
NOAS	Disable Async Option	This keyword is used to indicate whether or not asynchronous execution will be allowed. Setting this option to YES will prevent all asynchronous processing.	YES
NODA	Return Code if No Rows Affected	This keyword controls the return code when an INSERT/UPDATE/DELETE does not affect any rows. The following values are supported: <ul style="list-style-type: none"> • INFO: The return code will be SQL_SUCESS_WITH_INFO. • NODATA: The return code will be SQL_NO_DATA_FOUND. • SUCCESS: The return code will be SQL_SUCCESS. • ERROR: The return code will be SQL_ERROR. 	INFO
NONL	No Nulls	This keyword controls whether or not zero length strings should be returned. Possible values are as follows: <ul style="list-style-type: none"> • YES: Zero length strings will be replaced by one blank. • NO: Zero length strings will not be modified. 	NO
NOPM	Disable All Prompts	This keyword controls whether to disable all interactive prompts or informational message boxes. By setting this keyword to YES, all interactive prompts informational message boxes will be disabled. This feature is required when Shadow Mainframe Adapter Client is being called from an NT service, a UNIX daemon process, or any Mainframe Adapter Server type application that cannot be interrupted.	NO
NOPW	Don't Get Passwords from ODBC.INI/SHADOW.INI	This keyword controls whether the password should not be obtained from the SHADOW . INI file. Possible values are as follows: <ul style="list-style-type: none"> • YES: Any password value in the SHADOW . INI will be ignored; the password must be provided by the connection string or the connection dialogs. • NO: The password will be obtained from the SHADOW . INI file in some cases. <p><i>Note:</i> See also “ALPO” on page A-3.</p>	NO
NOUS	Don't Get Userids from ODBC.INI/SHADOW.INI	This keyword controls whether the userid should not be obtained from the SHADOW . INI file. Possible values are as follows: <ul style="list-style-type: none"> • YES: Any userid value in the SHADOW . INI will be ignored; the userid must be provided by the connection string or the connection dialogs. • NO: The userid will be obtained from the SHADOW . INI file in some cases. <p><i>Note:</i> See also “ALUO” on page A-3.</p>	NO

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
O3MD	ODBC 3.0 Meta Data Support	<p>This keyword controls whether Shadow Mainframe Adapter Client will request catalog data conforming to the ODBC 3.0 specification. Possible values are as follows:</p> <ul style="list-style-type: none"> • YES: Shadow Mainframe Adapter Client requests catalog data as defined in the ODBC 3.0 specification. When using Shadow Mainframe Adapter Server optimized catalogs, the proper Shadow Mainframe Adapter Server catalog level support must be in place to complete this support (the Shadow Mainframe Adapter Server parameter ODBCCATALOGLEVEL must be set to 4). If the Mainframe Adapter Client requests it, but the host installation does not have the correct catalog level turned on, the connection will fail. • NO: Shadow Mainframe Adapter Client requests catalog data as defined in the ODBC 1.0 specification. 	NO
ODBC.INI	ODBC.INI	(Read-only) This keyword is used to display the location and name of the ODBC . INI file found and used for the connection on UNIX platforms.	
ODPG	ODBC Mainframe Adapter Client Code Page	<p>This keyword is used to specify the ODBC Mainframe Adapter Client code page. The following values are supported:</p> <ul style="list-style-type: none"> • DEFAULT: The default is a set of UNIX code pages for UNIX and Windows Latin 1 (ANSI) for Windows 95/NT. • LATIN1: Forces the use of the Windows Latin 1 code page in any environment. • UNIX: Forces the use of the UNIX code pages in any environment. <p><i>Note:</i> Windows Latin 1 is also known as ISO 8859.</p>	DEFAULT
OJFX	Fix DB2 Outer Joins	If this keyword is set to YES, parentheses around ON clauses of DB2 outer joins will be removed. While this syntax is valid ANSI SQL, DB2 V5 cannot handle it. This kind of SQL is generated by tools such as Brio.	NO
OLEDB	Caller is an OLE DB Provider	<p>This keyword indicates whether Shadow Mainframe Adapter Client is operating as an ODE DB provider.</p> <p><i>Note:</i> This keyword can only be set through the connection string.</p>	NO
ONWY	One-Way Messages	<p>This keyword controls whether or not responses to messages should be sent to the Mainframe Adapter Client. Possible values are as follows:</p> <ul style="list-style-type: none"> • YES: One-way message processing will be used. This means that no responses will be returned to Mainframe Adapter Client applications. • NO: Normal responses will be returned for messages. <p><i>Note:</i> This keyword can only be used with messages.</p>	NO

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
OPRW	Optimal Row Count	This keyword limits the number of rows that will be returned from the host each time a request for rows is made. Since any number of requests for rows can be made for one query, this value will have no effect on the total number of rows returned by a query. This value can be used to control the size of the buffers used to return rows from the host.	0
OPTM	Operation Timeout Value	This keyword controls the timeout value (in seconds) for all Mainframe Adapter Client network operations (i.e. query preparation, execution, retrieving result set) <i>after</i> the connection has been established. This value should never be negative, but it can be zero. This value has no effect in UDP messaging. <i>Note:</i> The use of this keyword may cause some unpredictable results in many environments. This option is only available in the TCP/IP and the MQ link types (see “LINK” on page A-16), but it is applicable to both the permanent and the messaging modes (see “CNMD” on page A-6).	0
PAHN	Parameter Handling	This keyword is used to control how parameters are handled when they are passed by applications. In some cases, special processing is needed to fix values passed by some applications. Possible values are as follows: <ul style="list-style-type: none"> • NONE: Parameter values passed by applications are not changed. • INPUT: Required for Crystal Reports, because Crystal Reports passes parameters using obsolete values that must be modified before they can be used. • LENGTH: Fixes some application parameter length errors. • ADO: Fixes some ADO related errors. 	NONE
PAOP	SQLParamOptions Support	This keyword controls if Shadow Mainframe Adapter Client should support SQLParamOptions in a limited form. This keyword can be set to the following values: <ul style="list-style-type: none"> • Mainframe Adapter Server: The Shadow Mainframe Adapter Server code will decide whether this support is enabled. • YES: SQLParamOptions is supported to a limited extent. • NO: SQLParamOptions is not supported at all. 	Mainframe Adapter Server
PBFU	PowerBuilder Compatibility	This keyword can be set to YES to resolve certain problems with PowerBuilder. This keyword should not be set for any other reason.	NO

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
PLAN	Plan Name	<p>This keyword specifies the plan name that Shadow Connect uses on the host to make a connection to DB2. The format of this keyword is as follows:</p> <ul style="list-style-type: none"> The first two letters of the name are always “SD.” The third letter is always the last letter of the Shadow Mainframe Adapter Server subsystem name, which is “B” by default. The fourth letter designates the isolation level of the plan, as follows: <ul style="list-style-type: none"> C Cursor Stability SQL_TXN_READ_COMMITTED R Repeatable Read SQL_TXN_SERIALIZABLE S Read Stability SQL_TXN_REPEATABLE_READ U Uncommitted Read SQL_TXN_READ_UNCOMMITTED <p>Note: If the fourth character is any character other than “R,” “S,” or “U,” it is assumed the plan was bound with an isolation level of Cursor Stability.</p> <ul style="list-style-type: none"> The final four characters are always a version number (such as “1010” or “1040”). <p>For example, the default plan name value is SDBC1010. This default is release-dependent.</p> <p>Note: If the Mainframe Adapter Client sets this keyword to either DFLT or DEFAULT, then the DEFAULTDB2PLANNAME parameter that is set on the Shadow Mainframe Adapter Server will be used for that connection.</p>	SDBC1010
PORT	Port Number	<p>This keyword specifies the Shadow Mainframe Adapter Server listener port.</p> <p>Failover Shadow Mainframe Adapter Servers can be specified by listing host/port combinations as follows:</p> <p>HOST=Host1,Host2,Host3;PORT=1200,1210,1500</p> <p>Based on the host/port combinations specified, Shadow Mainframe Adapter Client would try to connect to Host1:1200, Host2:1210, and Host3:1500, in that order.</p> <p>Specified host names can be the same host. Note that this syntax should <i>not</i> be used to specify Shadow Mainframe Adapter Server group directors.</p>	1200
PRAD	Precision Adjustment	<p>This keyword controls the precision-adjustment Shadow Mainframe Adapter Client makes when passing a numeric value as a literal. When passing a literal numeric value, Shadow Mainframe Adapter Client will adjust the precision to be an odd number by adding one (1) to even precision values. For example, the value 12.34 will be sent with a precision of 5. This adjustment sometimes causes the rejection of acceptable values; thus, this precision-adjustment can be disabled by setting this keyword to NO.</p>	YES

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
PRES	Process Escapes	This keyword is used to turn off ESCAPE clause processing. Possible values are as follows: <ul style="list-style-type: none"> • YES: Perform ESCAPE clause processing. • NO: Turn off ESCAPE clause processing. This may result in improved performance by reducing CPU overhead. 	YES
PRNF	Procedure Name Filter	This keyword is used to filter the procedure names returned by the product. This keyword is used only if it is set to a non-blank value and if the host application does not provide a procedure name filter string. The procedure name filter string from either the host application or from this keyword restricts the information returned by procedure catalog inquiries (SQLProcedures and SQLProcedureColumns). Only rows that match the procedure name pattern provided will be returned. If this keyword is set to a single percent sign ("%"), then all procedures will be returned. There is no default value for this keyword.	
PROF	Procedure Owner Filter	This keyword is used to filter the procedure owners returned by the product. This keyword is used only if it is set to a non-blank value and if the host application does not provide a procedure owner filter string. The procedure owner filter string from either the host application or from this keyword restricts the information returned by procedure catalog inquiries (SQLProcedures and SQLProcedureColumns). Only rows with owner IDs that match the procedure owner pattern provided will be returned. If this keyword is set to a single percent sign ("%"), then all procedures will be returned. This keyword defaults to the current userid or alternate userid.	
PROW	Procedure Owner Handling	This keyword is used to control how procedure owner values are returned to applications, as follows: <ul style="list-style-type: none"> • NONE: The procedure owner value is not changed. • NULL: The actual owner value is used as a prefix to the procedure name, and the owner keyword is returned as a null value. • EMPTY: The actual owner value is used as a prefix to the procedure name, and the owner keyword is returned as an empty string. 	NONE
PWD	Password	This keyword specifies the mainframe password with which to connect to the host.	
PXPW	Proxy Mainframe Adapter Server Password	When HTTP tunneling messaging mode is enabled and the HTTP proxy Mainframe Adapter Server requires user authentication, this keyword allows the user to specify the proxy user password string. At this time, Basic HTTP Access authentication is the only method supported.	

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
PXSR	Proxy Mainframe Adapter Server Information	<p>When HTTP tunneling messaging mode is enabled and the HTTP stream must first connect to a proxy Mainframe Adapter Server, this keyword should be set to identify the proxy Mainframe Adapter Server's hostname (or IP address) and the port number. The accepted formats for this keyword are as follows:</p> <ul style="list-style-type: none"> • MY_PROXY:80 • 10.22.33.44:1200 <p>The first part of the value is the hostname or the IP address of the proxy Mainframe Adapter Server and the second part is the port number on which the proxy Mainframe Adapter Server is listening.</p> <p>Note: For some proxy Mainframe Adapter Server setups, user authentication might be required.</p>	
PXUS	Proxy Mainframe Adapter Server Userid	<p>When HTTP tunneling messaging mode is enabled and the HTTP proxy Mainframe Adapter Server requires user authentication, this keyword allows the user to specify the proxy userid (or name) string. At this time, Basic HTTP Access authentication is the only method supported</p>	
QUNA	Qualifier Name Separator	<p>This keyword should be set to YES if SQLGetInfo must return a period when obtaining SQL_QUALIFIER_NAME_SEPARATOR. This keyword is provided to fix certain application bugs.</p>	NO
RDFX	RDO Compatibility	<p>This keyword can be set to YES to resolve certain problems with Remote Data Objects. This keyword should not be set for any other reason.</p>	NO
RDON	Read Only	<p>This keyword is set to YES to make the data source read-only. Setting the data source to read-only will not actually prevent update operations; however, it will prevent any index information from being returned to the application. This will generally prevent any updates from being attempted.</p> <p>Note: Setting this keyword to YES will greatly improve the performance of some applications using the standard DB2 catalogs.</p>	NO
RMEQ	Remove Equals	<p>This keyword controls if the equals byte should be removed from MDI keyword names as part of MDI RSP (Remote Stored Procedure) invocation. Possible values are as follows:</p> <ul style="list-style-type: none"> • YES: The equals byte will be removed from each keyword name. • NO: The equals will not be removed from the keyword name. <p>Note: This keyword only applies to MDI RSPs invoked using TSQL 0/1/2 syntax. It does not apply to MDI RSPs invoked using the extended ODBC CALL syntax.</p>	NO
RO	Optimized Fetch	<p>This keyword controls whether or not to use the block fetch feature of Shadow Mainframe Adapter Client. By using the block fetch feature, the increase in performance is significant.</p> <p>Note: When the block fetch feature is enabled, cursor-based deletes (DELETE WHERE CURSOR OF) cannot be used.</p>	YES

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
RPGT	Allow Repeating SQLGetData	This keyword controls whether to allow repeating SQLGetData calls for the same unbound column while the cursor is still on that result set row. For variable length columns, applications can retrieve the column data in pieces. After the data is completely retrieved, the next SQLGetData call will retrieve the data from the beginning again. <i>Note:</i> LOB type columns cannot be retrieved repeatedly.	YES
RSBK	Result Set Cache Size	This keyword specifies the number of SQL statements contained by the result set cache (see “RSCH” on page A-25). <i>Note:</i> A higher value translates into more disk space.	256
RSCH	Result Set Caching	If this keyword is set to YES, Shadow Mainframe Adapter Client will cache result sets of SQL statements. When the same SQL statement is encountered, results are retrieved from a local cache if the underlying tables have not changed since the last cache update.	NO
RSIN	Result Set Caching Interval	This keyword specifies an interval (in seconds) within which stale results are tolerated. This keyword can be specified as follows: <ul style="list-style-type: none"> Any positive number: The result set will be returned without host processing if the following conditions are met: <ul style="list-style-type: none"> If result set caching is enabled (see “RSCH” on page A-25) and a previously cached SQL statement is encountered. The number of seconds since the result set was stored is within this interval. 0 (zero): All statements will be submitted to the host, where the underlying tables will be checked for changes; this will guarantee no stale results. -1: All result sets will always retrieved from a local cache (if one is available). <i>Note:</i> Tolerances greater than zero save network round-trips, but a value of zero is still useful since the result set is still retrieved from a local cache if the underlying tables haven’t changed.	0
RTGL	Return Global Tables	This keyword controls whether or not global temporary tables (GTTs) will be returned as normal tables. Possible values are as follows: <ul style="list-style-type: none"> YES: GTTs will be treated and described as normal tables. NO: GTTs will be handled as GTTs. This keyword is provided to allow standard tools to be used with GTTs. Many of these tools bypass GTTs that are described as GTTs; however, the same tools will work correctly with GTTs that are described as normal tables.	YES

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
RTSY	Return System Tables	<p>This keyword controls whether or not system tables (SYSIBM) will be returned as normal tables. Possible values are as follows:</p> <ul style="list-style-type: none"> • YES: System tables will be treated and described as normal tables. • NO: System tables will be handled as system tables. <p>This keyword is provided to allow standard tools to be used with system tables. Many of these tools bypass system tables that are described as system tables. However, the same tools will work correctly with system tables that are described as normal tables.</p>	NO
RTTM	UDP Datagram Retry Timeout	<p>This keyword controls the initial timeout period (in seconds) before a UDP request is retransmitted. This timeout value is doubled on successive retries; however, this value will not increase without bound—the maximum timeout allowable is 30 seconds.</p> <p>Note: If the UDP request retry feature is disabled (see “MXRT” on page A-18), this keyword will have no effect.</p>	30
RTPP	Return TABLE_TYPE Column for Meta Data Calls	<p>This keyword controls whether the TABLE_TYPE column is returned for the SQLPrimaryKeys and the SQLStatistics calls. If this keyword is set to YES, the TABLE_TYPE column will be returned as the last column of the result set.</p>	YES
SBMD	SBCS Mode	<p>(Only for DBCS or GRAPHIC data) This keyword is used to specify how SBCS (single byte character set) data is handled. SBCS strings are assumed to contain a mixture of single byte characters and double byte characters. Each set of double byte characters starts with an SO byte and ends with an SI byte. This keyword can be set as follows:</p> <ul style="list-style-type: none"> • DEFAULT: (Default) When processing DBCS or GRAPHIC data, leaving the default value of DEFAULT will cause an error. • BLANK: All SO/SI bytes will be converted to blanks. • DELETE: All SO/SI bytes will be deleted from each string. <p>Note: This keyword only applies if you are using DBCS or GRAPHIC data, in which case you must choose how to handle the SO/SI (shift out/shift in) bytes, as attempting to use DEFAULT will result in an error. The two choices are to convert them to blanks (BLANK) or remove them (DELETE).</p>	DEFAULT
SECN	Application Mainframe Adapter Server	<p>This keyword is used to select the initial application Mainframe Adapter Server value on the host. It must be sixteen or fewer characters long. If this keyword is set to a non-blank, non-null value, the current application will be connected to the specified application Mainframe Adapter Server after DSNALI OPEN processing is completed.</p>	

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
SECU	Network Authentication	<p>This keyword controls the network authentication model.</p> <p>Support of network authentication models allows users to connect without being prompted for their password, assuming they have already been authenticated via their NT domain and their userid is the same on the OS/390 or z/OS system. It is required that both the authenticated LANID and TSOID that are used during logon processing are the same and are defined in the OS/390 or z/OS security database for RACF, ACF2 or Top Secret.</p> <p>The following values are supported:</p> <ul style="list-style-type: none"> • NONE • Domain-based: When the domain-based authentication is selected, Shadow Mainframe Adapter Client will verify if the userid associated with the current process has been authenticated by a domain-based system. For Windows platforms, this requires the user to first log on to a NT domain. For UNIX platforms, the local machine must be a member of a NIS domain, and the password database used to authenticate the user must be NIS-mapped. • MS Transaction Mainframe Adapter Server: The Microsoft Transaction Mainframe Adapter Server (MTS) based authentication model allows an MTS Mainframe Adapter Server COM component to pass on the account name of the Mainframe Adapter Client applications invoking the object. • MyNet: This setting is used with the CKS MyNet single signon product support. • SNA: This setting works with the MS SNA Mainframe Adapter Server single signon facility. The Mainframe Adapter Client must be a MS/SNA CPI-C Mainframe Adapter Client. <p><i>(Continued on next page)</i></p>	NONE
SECU <i>(Continued)</i>	Network Authentication	<p><i>(Continued from previous page)</i></p> <p>Note: In order to support the functionality of this keyword, the Shadow Mainframe Adapter Server parameter Mainframe Adapter ClientLOGON must be set to YES.</p> <p>On the Shadow Mainframe Adapter Server, a RACINIT call is still performed to properly set up the ACEE information for the connection; however, the call will specify that the connection is trusted.</p>	NONE
SEDG	Current Degree	<p>This keyword is used to set the initial current degree value on the host. The only possible values for this keyword are as follows:</p> <ul style="list-style-type: none"> • ANY • 1 <p>No other values are supported at this time. If this keyword is set to a non-blank, non-null value, a SET CURRENT DEGREE statement is issued after DSNALI OPEN processing is completed.</p>	

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
SEPK	Current Packageset	This keyword is used to set the initial current packageset value on the host. This value must be eighteen or fewer characters long. If this keyword is set to a non-blank, non-null value, a SET CURRENT PACKAGESET statement is issued after DSNALI OPEN processing is completed. The SET statement assigns the specified value to the CURRENT PACKAGESET special register.	
SERL	Current Rules	This keyword is used to set the initial current rules value on the host. The only possible values for this keyword are as follows: <ul style="list-style-type: none"> • DB2 • STD No other values are supported at this time. If this keyword is set to a non-blank, non-null value, a SET CURRENT RULES statement is issued after DSNALI OPEN processing is completed.	
Mainframe Adapter Server	Mainframe Adapter Server	This keyword can be set to bypass a FILEDSN problem with Microsoft Access. <i>Note:</i> This keyword can be set through the connection string, but it is not available via jConfig tool dialog boxes.	
SEVL	System Engineering Value	This keyword can be set to various values to obtain diagnostic and debugging data. This keyword should only be used at the specific request of the Customer Support staff.	0
SKEY	Customer Secret Key	This keyword is not for general use.	
SRMD	Instance Message Digest	This keyword is not used.	
SQVA	SQL Mainframe Adapter Server Support	This keyword is used to control whether or not SQLTypeInfo calls should show that variable data is 255 bytes long. DB2 variable data is actually only 254 bytes long; however, some applications do not support that length. Possible values for this keyword are as follows: <ul style="list-style-type: none"> • YES: SQLTypeInfo will return 255 as the length of variable data. If this keyword is set to YES, the behavior of SQLCancel is modified to circumvent vendor implementation errors. • NO: SQLTypeInfo will return 254 as the length of variable data. 	NO
SSLX	Secured Channel Protocol	This keyword allows the user to select a secured channel protocol to encrypt and authenticate the Mainframe Adapter Client/Mainframe Adapter Server session, as follows: <ul style="list-style-type: none"> • Standard: The default is the standard Shadow Connect secured data stream, which provides optimal performance but does not support strong encryption and authentication. • SSL2: This is based on the SSL version 2 standard. • SSL3: This is based on the SSL version 3 standard; however, this value is not supported at this time. 	Standard

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
STFX	Fix String Length	This keyword can be set to YES to fix a specific bug that can occur in some products, like Microsoft Access, where the string length is incorrectly set to zero when the SQL_CHAR keyword only contains blanks. <i>Note:</i> If the keyword AF (MS Access Compatibility) is set to YES (see “AF” on page A-2), this keyword does not need to be set. The keyword AF includes the functionality of keyword STFX.	NO
STZO	Suppress Decimal Trailing Zeros	This keyword controls whether or not trailing zeros should be removed from decimal keywords. If this keyword is set to YES, trailing zeros will be removed. <i>Note:</i> If the keyword AF (MS Access Compatibility) is set to YES (see “AF” on page A-2), this keyword does not need to be set. The keyword AF includes the functionality of keyword STZO.	NO
SUBSYS	Subsystem	This keyword specifies the DB2 subsystem. This keyword should only be set if using DB2; otherwise this keyword must be set to NONE.	
SYFU	Sybase Compatibility	If this keyword is set to YES, a set of Sybase fixes will be used. These fixes bypass bugs in the Sybase ODBC support.	NO
SYOP	Table Owner - Synonym Option	This keyword is specifies how table owners should be handled for tables that are actually synonyms. Some of these choices are required to make specific desktop productivity tools work with Shadow Mainframe Adapter Client. Possible values are as follows: <ul style="list-style-type: none"> • ZERO • NORMAL • NULL • FORCEZERO 	ZERO
TBFL	Table Name Filter	This keyword is used to filter the table names returned by the product. This keyword is only used if it is set to a non-blank value and if the host application does not provide a table name filter string. The table name filter string from either the host application or from this keyword restricts the information returned by catalog inquiries (SQLTables, SQLColumns, SQLTablePrivileges). Only rows that match the table name pattern provided will be returned. If this keyword is set to a single percent sign (“%”), all tables will be returned. There is no default value for this keyword. <i>Examples:</i> <ul style="list-style-type: none"> • TBFL=STAFF will cause only information from the table labeled “STAFF” to be displayed. • TBFL=STA% will cause information from all tables that begin with “STA” to be displayed. • TBFL=STAFF BOOK% will cause information from the table labeled “STAFF” and all tables that begin with “BOOK” to be displayed. 	

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
TCLM	Table & Column Name Modification	This keyword controls adjustment of table and column names for DB2, as follows: <ul style="list-style-type: none"> NONE: The names are considered correct and are not processed. BEA1: The BEA formula converts names to 18 characters by using the first character and last 17 characters for names longer than 18 characters. 	NONE
TQOP	Table Qualifier Option	This keyword specifies how table qualifiers should be returned to application programs. Some databases use table qualifiers as part of the overall name of each table. In other words, two tables can have exactly the same name in all other respects, if their table qualifiers are different. There are three possible values: <ul style="list-style-type: none"> NORMAL NULL ZERO 	NORMAL
TRLT	Truncate Literal String	This keyword controls whether literal strings are truncated. Setting this keyword to NO will not truncate the literal string even if the string length is greater than 20.	YES
TRNA	Transaction Name	This keyword is used to specify the transaction name. The use of the transaction name is application specific. It can be up to eight (8) bytes long. The transaction name is padded on the right with blanks.	
TSCH	Convert DB2 Time to CHAR	When this keyword is set to YES, the DB2 timestamp and time keyword will be converted to SQL_CHAR. Note: If the keyword AF (MS Access Compatibility) is set to YES (see “AF” on page A-2), this keyword does not need to be set. The keyword AF includes the functionality of keyword TSCH.	NO
UPCH	Uppcase All Character Data	This keyword controls whether or not all character data sent to the host should be converted to uppercase. Possible values are as follows: <ul style="list-style-type: none"> YES: All character data will be converted to uppercase. NO: Character data will not be modified. 	NO
UPDB	Uppcase Double Quote String	This keyword controls whether or not strings in double quotes should be converted to uppercase. Strings must be converted to uppercase in some cases because DB2 treats table names in double quotes as literals. Possible values are as follows: <ul style="list-style-type: none"> YES: Strings in double quotes will be converted to uppercase. NO: Strings in double quotes will not be modified. 	NO
UPNL	Uppcase All Non-Literals	This keyword controls whether all SQL keywords that are non-literals will be converted to uppercase. Note: Although the default for this keyword is NO in most cases, if the following conditions are met, this keyword will default to YES: <ul style="list-style-type: none"> If the AUST keyword is set to NO (see “AUST” on page A-40). If the LGID keyword is set to JPL or JPX (see “LGID” on page A-16). 	NO

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
USERID	Userid	This keyword specifies the mainframe userid with which to connect to the host.	
USERPARM	Host User Parm	This keyword is sent to the host as part of the logon information. The host uses this keyword to complete logon to the host security system and/or host databases.	
VAFX	Visual Age Compatibility	This keyword can be set to YES to resolve certain problems with Visual Age. This keyword should not be set for any other reason.	NO
WACU	Display Wait Cursor	This keyword controls whether or not to update the cursor (with an hourglass). Users may observe a dramatic improvement in performance by setting this option to NO, which results in no updates to the cursor and no display of the hourglass.	NO
WGPH	Wollongong Pathworks 3.1 Support	This keyword must be set to YES when running in the Wollongong Pathworks 3.1 16-bit TCP/IP environment	NO
WIUR	Add WITH UR to Queries	This keyword controls whether or not a WITH UR clause should be added to queries. The WITH UR clause reduces the amount of CPU time needed for some queries. There is also the possibility it could change the results of these queries by allowing uncommitted data to be read. Possible values are as follows: <ul style="list-style-type: none"> • YES: A WITH UR clause will be added to queries. • NO: A WITH UR clause will not be added to queries. 	NO
WRPR	Enable SQLPrepare for ADABAS/VSAM	This keyword controls whether or not Shadow Mainframe Adapter Client should enable SQLPrepare support for ADABAS and VSAM statements that have parameter markers. Possible values are as follows: <ul style="list-style-type: none"> • YES: Enable SQLPrepare support for statements that have one or more parameter markers. • NO: Prepare will be deferred for statements that have parameter markers. <p>Note: For the statements that do not have parameter markers, regardless of the setting of this keyword, SQLPrepare is always supported and the result-set is returned at SQLPrepare.</p>	YES
XAEN	X/OPEN XA Support	This keyword should be set if the target Mainframe Adapter Server will participate in an X/OPEN XA transaction. Possible values are as follows: <ul style="list-style-type: none"> • None: This value implies that no part of the transaction will use the XA protocol. • TWO-PHASE: This value should be used in a situation where the transaction involves more than one resource manager and the target Mainframe Adapter Server supports the two-phase commit protocol. The value TWO-PHASE also allows for the one-phase commit scenario whenever the coordinator thinks it is appropriate. 	None

Table A–1. Shadow Mainframe Adapter Client Keywords (Continued)

Keyword	Description	Explanation	Default
XAOP	X/OPEN XA Transaction Manager	<p>This keyword is used to set the XA transaction manager type. This option must be set correctly when participating in a distributed transaction coordinated by a monitor. Supported values are as follows:</p> <ul style="list-style-type: none"> • None: This value implies that this is not a distributed transaction. • TUXEDO-TMS: Used for BEA Tuxedo coordinated transactions to identify that the connection will be owned by the Tuxedo TMS Mainframe Adapter Server. • TUXEDO-SQL: Used for BEA Tuxedo coordinated transactions to identify that the connection will be owned by the Tuxedo SQL application Mainframe Adapter Server. • MSDTC: MSDTC is the Microsoft Distributed Transaction Manager used mostly by the MTS applications. • JTS: Used for any transaction manager compatible with the Java Transaction API (JTA). 	None
ZEAS	Zero Active Statements	<p>This keyword controls whether or not zero should be returned for the number of active statements. Possible values are as follows:</p> <ul style="list-style-type: none"> • YES: Zero will be returned for the number of active statements supported by Shadow Mainframe Adapter Client. • NO: The actual number of supported active statements will be returned. 	NO
ZECL	Zero Column Names	This keyword must be set to YES if column names should be set to binary zeros. This is a performance optimization for production applications.	NO

Setting Keywords to Benefit Performance

There are several keywords that can be set to improve performance and to benefit users of certain tools. This section will highlight these keywords and their benefits for various applications.

3-Tier Applications

- **NOPM (Disable All Prompts)**

You must set this keyword to YES when Shadow Mainframe Adapter Client is being called from an NT service, a UNIX daemon process, or any other Mainframe Adapter Server-type application that cannot be interrupted. This will prevent interactive prompts or informational message boxes, such as those informing you about license or password expiration status, from interrupting an application.

General Use

- **LGPA (Convert Strings to Params)**

Setting this keyword to YES results in the conversion of strings greater than 254 bits to parameter markers. This is necessary for those applications that produce literals that are too long to be handled by the host database.

Without this setting, you will get the SQLCODE error -102, indicating that the host database could not handle the long string.

- **DIPO (Disable Prepare/Open Optimization)**

Setting this keyword to NO enables the prepare/open optimization feature. The prepare/open optimization feature saves a network round-trip for applications that use the ODBC functions SQLPrepare and SQLExecute instead of SQLExecDirect. Using SQLPrepare to prepare a SQL statement and SQLExecute to execute that SQL statement is a two trip operation. With prepare/open optimization, Shadow Connect has combined these two functions into a single network operation. If the application issues a SQLPrepare, Shadow Mainframe Adapter Client, in conjunction with Shadow Mainframe Adapter Server, prepares *and executes* the SQL statement. As a result, when the Mainframe Adapter Client issues the SQLExecute, Shadow Mainframe Adapter Client immediately returns SQL_SUCCESS without performing the additional network round-trip.

This works great for most applications; however, if the applications issue SQLPrepare calls “under-the-covers” to access meta data information, the SQL statements will never be executed. To disable the prepare/open optimization feature, set the DIPO keyword (Disable Prepare/Open Optimization) to YES.

- **FALG (Fast Logon)**

When set to YES, this keyword allows for logon processing with a minimum of network I/O. However, since this feature can only be used by host Mainframe Adapter Servers that support fast logon, it will fail with older versions of the Shadow Mainframe Adapter Server.

- **USERPARM (Host User Parm)**

The information stored by this keyword is sent to the host as part of the logon information and is also inserted into the SMF records. This keyword is used to complete logon to the host security system and/or host database.

- **OPTM (Operation Timeout Value)**

This keyword controls the timeout value (in seconds) for all Mainframe Adapter Client network operations after the connection has been established. When set to the default value of 0, it will not cancel a long query; however, when set to a number greater than 0, it will cancel the connection after the operation has exceeded the timeout value.

- **OPRW (Optimal Row Count)**

This keyword allows you to specify the number of rows returned per block fetch. This will limit the number of rows that will be returned from the host each time a request for rows is made. This value will have no effect on the total number of rows returned by a query, but it will control the number per block fetch.

- **RDON (Read Only)**

Setting this keyword to YES will greatly improve the performance of some applications using the standard DB2 catalogs. It will make the data source read-only, preventing any index information from being returned to the application. This will generally prevent any update attempts, although it will not actually prevent update operations.

- **LGMG (Return Logon Messages)**

When this keyword is set to YES, if a logon returns a message, a return code of SQL_SUCCESS_WITH_INFO will be set. This way, a dialog box is not displayed, and the message text will be available using the SQLError function.

- **TQOP (Table Qualifier Option)**

This keyword is used to specify how table qualifiers should be returned to application tables. Typically, an application will issue SELECT statements that contain the database name, the qualifier, and the table name. Possible values for this keyword are NORMAL, NULL, and ZERO. When this keyword is set to ZERO, Shadow Mainframe Adapter Client will return strings of zero length for the database name when calling the SQL tables.

APPENDIX B: National Language Support

This chapter describes the National Language Support (NLS) of Shadow Mainframe Adapter Client. Topics include the following:

- Languages Supported
- Setting the Language

Languages Supported

Table B–1 below lists the languages that Shadow Mainframe Adapter Client supports, along with the code pages used for each.

Table B–1. Shadow Mainframe Adapter Client NLS Support

Language	Windows Language Code	EBCDIC Code Page Number	Windows Code Page Number
Arabic	ARB	N/A	1256
Canadian French	FRC	037	1252
Danish	DAN	277	1252
Dutch	NLD	037	1252
English (UK)	ENG	285	1252
English (US)	ENU	037	1252
Finnish	FIN	278	1252
French	FRA	297	1252
German	DEU	273	1252
Icelandic	ISL	871	1252
Italian	ITA	280	1252
Japanese	JPL	290	932
Korean	KOR	037	1252
MDI	MDI	N/A	1252
Norwegian	NOR	277	1252
PeopleSoft	PPS	N/A	1252
Portuguese	PTG	037	1252
Portuguese (Brazilian)	PTB	037	1252

Table B–1. Shadow Mainframe Adapter Client NLS Support (Continued)

Language	Windows Language Code	EBCDIC Code Page Number	Windows Code Page Number
Spanish (Modern)	ESN	284	1252
Spanish (Castilian)	ESP	284	1252
Swedish	SWE	278	1252

Setting the Language

Shadow Mainframe Adapter Client accommodates your choice of language in a transparent manner. By default, Shadow Mainframe Adapter Client uses the language setting you have already specified for your machine; however, this language value can be overwritten using connection string or the data source settings.

Shadow Mainframe Adapter Client attempts to obtain the language information from three sources. The three sources and the priority order by which the language is set are as follows:

- Connection string.
- Data source settings.
- Locale settings.



Note:

If the language has not been set by any of the methods mentioned above, Shadow Mainframe Adapter Client will default to English (US).

1. **Connection String:** The connection string is passed by an application. To specify a language using a connection string, the LGID keyword must be coded. An example JDBC connection string follows:

```
jdbc:neon:<data-source-name>;
HOST=MKT.NEONSYS.COM;PORT=1200;UID=UID1;PWD=PWD1;
SUBSYS=DSN1;PLAN=SDBC1010;CPFX=SYSIBM;LGID=DAN
```

2. **JDBC Data Source Settings:** The language information may also be specified in the data source settings, which are configured by using the jConfig tool to configure the data source settings by selecting the data source and navigating to **Configure** → **Parameters**. Then, from the **Parameter Settings** list, select **Language ID** (keyword LGID) and select an appropriate value from the **Value** list. See Figure B–1.




Figure B–1. Optional Parameters Information -- Setting the Language ID



Doc Reference:

For more information about configuring a JDBC data source, see “Configuring a JDBC Data Source” on page 4-5 within Chapter 4, “Shadow Mainframe Adapter Client: Configuration,” of this guide.

3. **Locale Setting:** The locale setting can also be used to set the language as follows:
 - **Windows:** To set the language for the Shadow Mainframe Adapter Client, go into the **Control Panel**, select the **Regional Options** icon, and select the language using the drop-down menu. This causes the new value to be stored in the WIN.INI file.
 - **UNIX:** You will need to adjust your machine-specific locale environment variable (see your system administrator for details).

APPENDIX C: JDBC Compliance

Shadow Mainframe Adapter Client, part of the Mainframe Adapter Client component of the Shadow Connect product, offers JDBC compliance as detailed in this appendix.

JDBC Compliance

CallableStatement - *Partially Supported*

Methods *not* supported include:

- `getArray (int)`
- `getBlob (int)`
- `getClob (int)`
- `getObject (int, java.util.Map)`
- `getRef (int)`
- `registerOutParameter (int, int, String)`

Connection - *Partially Supported*

Methods *not* supported include:

- `getTypeMap ()`
- `setTransactionIsolation (int)`



Note:

The transaction isolation level is set by selecting the right plan at connection time. Once it's selected, there is no way of changing it.

- `setTypeMap (java.util.Map)`

ConnectionPoolDataSource - *Fully Supported*

DatabaseMetaData - *Fully Supported*

Methods that return hard-coded values include:

- `deletesAreDetected (int)` - Always returns false
- `getExtraNameCharacters ()` - Always returns null
- `getUDTs (String, String, String, int[])` - Always returns null
- `insertsAreDetected (int)` - Always returns false
- `othersDeletesAreVisible (int)` - Always returns false
- `othersInsertsAreVisible (int)` - Always returns false
- `othersUpdatesAreVisible (int)` - Always returns false

- `ownDeletesAreVisible (int)` - Always returns false
- `ownInsertsAreVisible (int)` - Always returns false
- `ownUpdatesAreVisible (int)` - Always returns false
- `supportsANSI92EntryLevelSQL ()` - Always returns true
- `supportsANSI92FullSQL ()` - Always returns false
- `supportsANSI92IntermediateSQL ()` - Always returns false
- `updatesAreDetected (int)` - Always returns false

DataSource - Fully Supported

Driver - Fully Supported

PooledConnection - Fully Supported

PreparedStatement - Partially Supported

Methods *not* supported include:

- `setArray (int, Array)`
- `setBlob (int, Blob)`
- `setClob (int, Clob)`
- `setNull (int, int, String)`
- `setRef (int, Ref)`

ResultSet - Partially Supported

Methods *not* supported include:

- `absolute (int)`
- `afterLast ()`
- `beforeFirst ()`
- `cancelRowUpdates ()`
- `deleteRow ()`
- `first ()`
- `getArray (int)`
- `getArray (String)`
- `getBlob (int)`
- `getBlob (String)`
- `getClob (int)`
- `getClob (String)`
- `getFetchSize ()`
- `getObject (int, java.util.Map)`
- `getObject (String, java.util.Map)`
- `getRef (int)`
- `getRef (String)`
- `getRow ()`
- `insertRow ()`
- `isAfterLast ()`
- `isBeforeFirst ()`
- `isFirst ()`

- `isLast ()`
- `last ()`
- `moveToCurrentRow ()`
- `moveToInsertRow ()`
- `previous ()`
- `refreshRow ()`
- `relative (int)`
- `rowDeleted ()`
- `rowInserted ()`
- `rowUpdated ()`
- `setFetchSize (int)`
- `updateXXX` - Not all of the updates methods are supported

ResultSetMetaData - *Fully Supported*

Statement - *Partially Supported*

Methods *not* supported include:

- `cancel ()` - Asynchronous execution is not supported
- `getFetchSize ()`
- `getQueryTimeout ()`
- `setEscapeProcessing (boolean)`
- `setFetchSize (int)`

XAConnection - *Fully Supported*

XADataSource - *Fully Supported*

XAResource - *Fully Supported*

Xid - *Fully Supported*

