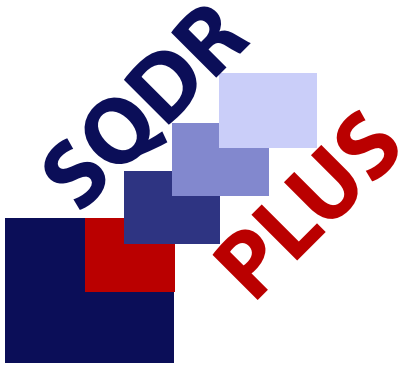

SQDR Plus for DB2 Universal Database

STARQUEST

Statement of Limitations on Warranty & Liability

StarQuest Ventures makes no representations or warranties about the suitability of the software and documentation, either expressed or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. StarQuest Ventures shall not be liable for any damages suffered by licensee as a result of using, modifying, or distributing this software or its derivatives.

StarSQL™, StarQuest Data Replicator™, SQDR™, and SQDR Plus™ are trademarks of StarQuest Ventures, Inc. Java and all Java-related products are registered trademarks of Oracle. All trademarks or registered trademarks are the property of their respective owners.

v3.72_07.28.2011

Contacting StarQuest Ventures

StarQuest Ventures provides technical support subject to the prices, terms, and conditions specified in your organization's maintenance contract with StarQuest.

If you have suggestions for product enhancements or need more information about how StarQuest products can provide solutions for connecting Windows and UNIX applications to IBM host resources, please contact us via any of the following methods.

Address	StarQuest Ventures, Inc. P.O. Box 1076 Point Reyes Station, CA 94956
Telephone	415-669-9619 Option 1: Sales Option 2: Technical Support
Fax	415-669-9639
Email	contact@starquest.com
World Wide Web	www.starquest.com

Contents

Introduction

StarQuest Data Replicator (SQDR) is a software product that replicates data between IBM DB2, Oracle, and Microsoft SQL Server database environments. You can replicate any table, portion of a table, or view. You can schedule replications to take place at specified intervals or on demand. Replicating data from one database management system to another allows you to:

- make the same data available to users of different database systems
- make the same data available to multiple sites
- balance network and database server loads by making the data available from two or more database servers
- create backup copies of data

StarQuest Data Replicator and SQDR Plus

The base StarQuest Data Replicator software provides full refresh replication, which copies all the source data that is selected for replication, regardless of when that data was last replicated. This provides a “snapshot” of the specified source data at the time of replication.

StarQuest also offers the separately licensed component, SQDR Plus, to provide incremental replication support for SQDR users. Incremental replication copies only the data that has changed, which reduces the amount of data that is transferred and allows the replication to occur more frequently. SQDR Plus for DB2 UDB is part of the SQDR Plus product family and provides incremental replication for IBM DB2 UDB for Linux, UNIX, and Windows platforms.

Incremental replication requires installing the SQDR Plus software on the host database system to monitor changes and communicate with the Data Replicator Windows-based service. Note that no host software is required to use only the base SQDR software to perform full refresh (snapshot) replication. As shown in the following illustration, snapshot replication operations require only the Data Replicator service running on a Windows-based computer and the appropriate ODBC drivers for access to the source and target databases, whereas incremental replication operations are allowed from hosts that have SQDR Plus installed.

Figure 1: Overview of SQDR and SQDR Plus

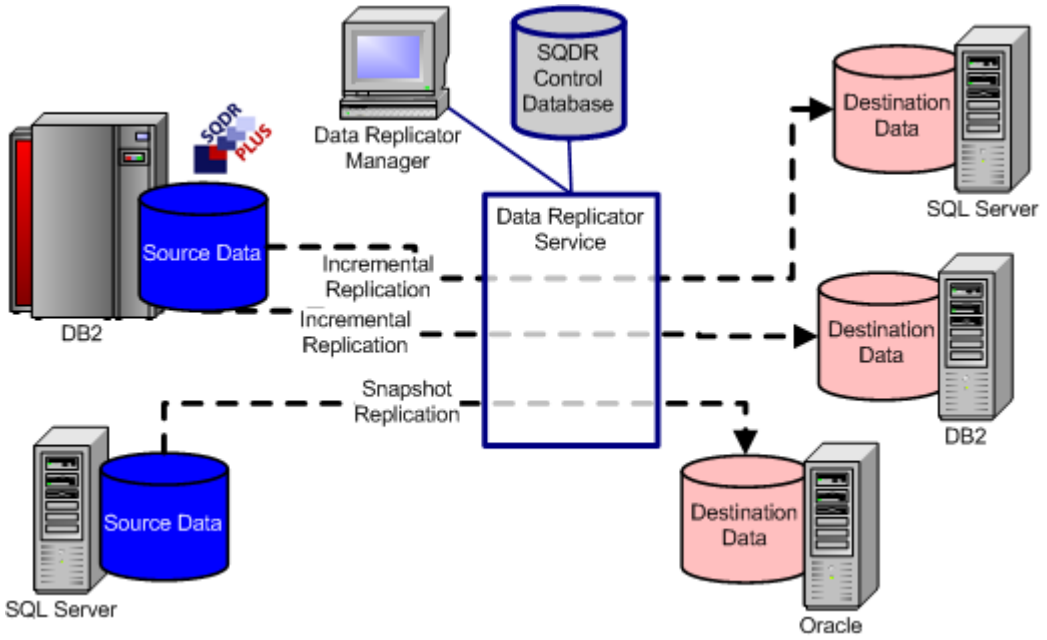

SQDR Plus Components

SQDR Plus for DB2 UDB includes the following components:

- Capture Agent and Java Service Wrapper
- Java and SQL Stored Procedures
- Replication Worker
- Log Reader
- Control Tables and Change Data Tables
- Capture Agent Maintenance Utility

The Java Service Wrapper allows the Capture Agent to be installed as a service on Windows-based computers and as a daemon on Linux- and UNIX-based computers. The Log Reader monitors the DB2 logs for changes and the Replication Worker stages the changes to the control tables and change data tables. The Capture Agent Maintenance program is provided to help manage the Capture Agent and is described in “Using the Capture Agent Maintenance Utility” on page 50.

The following illustration shows how the SQDR Plus components work within the DB2 UDB environment and with the base SQDR product to create a baseline snapshot replication from which point changes are tracked so they can be incrementally applied to the destination database.

Figure 2: SQDR Plus Components Enable Incremental Replication of Data

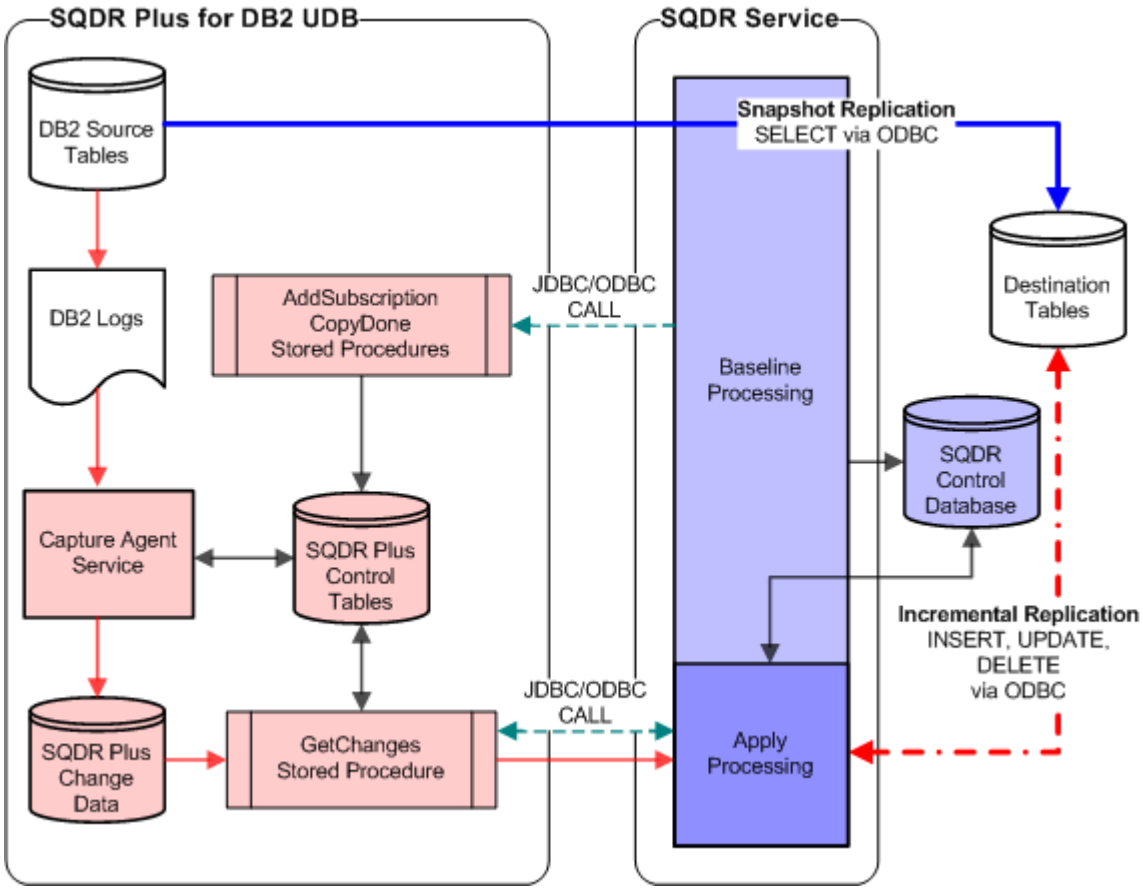

From the Data Replicator Manager, which runs on the same Windows computer as the Replicator Service, a subscription is created that defines what source data to replicate to what destination. If the source database is enabled with SQDR Plus, the subscription can perform incremental replication operations.

When a subscription that performs incremental replication is added to the SQDR Control Database, the SQDR service calls the AddSubscription stored procedure and a baseline snapshot replication is begun. The Capture Agent Service sets the table position and begins tracking changes. When the baseline snapshot replication is complete, the SQDR service calls the

CopyDone stored procedure to indicate that it is ready to receive change data. The Capture Agent sends any changes that accumulated while the baseline snapshot was being created, and continues monitoring the DB2 database log files to stage changes that are committed to the database.

The Capture Agent Service uses the Remote Method Invocation (RMI) to communicate with the Stored Procedures, which in turn communicate via JDBC/ODBC calls to the SQDR Service and Apply Manager. The Change Data Tables provide the primary input to the Stored Procedures, with secondary input from the Control Tables, which are managed by the Capture Agent service. The SQDR Apply Manager periodically polls for the changes to the source tables and applies any inserts, updates, and deletes the destination data as appropriate to keep it synchronized with the source data. You also can configure the Capture Agent service to send notification to the SQDR client when change data is available for applying to the target database.

SQDR Plus Features and Benefits

SQDR Plus offers an efficient replication solution that provides the following advanced features.

- Comprehensive statistics and easy-to-use graphical interface and Wizards.
- Incremental replication that copies only the data that has changed, providing significant reduction in the amount of data that is transferred.
- Enterprise-level security, allowing only operations that are permitted by the operating system, database system, and network.
- Ability to replicate what you want, when you want, with horizontal and vertical partitioning and flexible subscription scheduling.
- Ability to “mirror” the source database. In addition to configuring how frequently SQDR Plus polls for changes to the source database you can configure subscriptions to be notified when change data is available and reduce the latency such that updates occur at the target database almost instantaneously after they are committed on the source database.
- Ability to restrict which database tables can be subscribed to for replication. The SQDR Plus Capture Agent utility provides centralized control for publishing and unpublishing database tables. If you enable the publishing feature, SQDR users can subscribe only to tables that have been published for replication.

SQDR Plus Limitations

The SQDR Plus Capture Agent imposes a few limitations relative to the capabilities of the DB2 UDB server and the base SQDR software. For incremental replication operations, please note that:

- The maximum capacity of tables that can be replicated incrementally is slightly less than the database system allows to accommodate the mechanisms that track the change data. LOB data is replicated directly from the source instead of being stored in the change data (staging) tables. SQDR Plus needs a minimum of five columns (more if there are unique indexes) for managing the change data in the staged tables, therefore the tables cannot contain the maximum number of columns or the maximum bytes per row that are allowed by the system.

The system limits vary according to which release of DB2 UDB you are running, so refer to your system documentation if you work with very large tables and need to determine the maximums allowed.

- Foreign constraints are not replicated for incremental operations. Unique constraints are replicated as non-unique indexes, and indexes can be optionally replicated.
- You can incrementally replicate a subset of columns to a destination, but the subset must include all the columns associated with at least one unique index that also is replicated.
- Not withstanding the foregoing limitation, SQDR Plus can replicate tables that have no unique indexes defined if the committed changes include only insert operations. A new baseline snapshot must be performed if any delete or update operations are committed to the source table. To avoid frequent baseline replications consider altering the source table to include an identity column, as described in “Creating an Identity Column” on page 68.

The remainder of this document covers the installation, setup, and operation of the SQDR Plus for DB2 UDB host software. Although SQDR Plus functions similarly on all the supported platforms, differences in the operating systems require different procedures for some tasks. Where information applies to only one platform this documentation shows a UNIX or Windows icon in the margin so you can more easily locate or skip particular information, as illustrated below.

This information applies only to users of a computer that is running the Linux or UNIX operating system.

This information applies only to users of a computer that is running the Windows operating system.

Installing SQDR Plus for DB2 UDB

This chapter describes how to prepare for and install the SQDR Plus for DB2 UDB software. SQDR Plus may be installed on any number of DB2 UDB platforms within an organization. Licensing is controlled by SQDR and StarLicense and is managed according to each DB2 UDB source database and DBMS destination pair. Incremental replication is supported only for the specific source and destination pairs that are licensed for SQDR Plus. The online help for the Data Replicator Manager and StarLicense products describe how to configure licenses for using StarQuest products.

System Requirements for the DB2 UDB Computer

Following are the requirements for installing SQDR Plus on a DB2 UDB computer:

- DB2 for Linux, UNIX and Windows v8.2 with FixPak10 or later, or v9.1 or later. (DB2 FixPak 10 also is known as DB2 Version 8.2 with FixPak 3 of the V8.2 maintenance stream, depending on how you acquired the product.)
- SQDR Plus for DB2 UDB currently supports the following platforms:
    - AIX v5.3 & later
    - HP-UX PA-RISC
    - Linux x86 (2.6 kernel) and Linux on x86-64 (AMD64 and Intel EM64T)
    - Solaris 10 & later SPARC
    - Windows 32-bit and 64-bit

  See the IBM website, http://www-306.ibm.com/software/data/db2/udb/sysreqs.html, for detailed operating system requirements for DB2 UDB and to obtain the latest DB2 UDB fixes.
- Java Virtual Machine (JVM) installed and specified in the PATH. If you install DB2 UDB on a Windows computer using the DB2 Setup Wizard and the JVM is not already installed in the path, the Setup Wizard will install it. JVM 1.4.2 or later is required during installation, and recommended for use by the Capture Agent and Java stored procedures. If you need to download a JVM, refer to the appropriate operating system vendor website, such as the Oracle Java website for Solaris, Linux, and Windows systems, the Hewlett-Packard website for HP-UX, or the IBM developerWorks or Fix Central website if you are running AIX.

UNIX

- There are two options for installing SQDR Plus on a UNIX-based computer: a GUI installer or a text-based installer. The text-based installer involves editing a configuration file with a text editor such as vi or gedit, and requires fewer system resources than the GUI installer. In addition, support for some platforms is available only in the text-based installer, and future installer development will be oriented to enhancing the text-based installer. The GUI installer is recommended only for users unfamiliar with the UNIX command line environment; it requires an X-Windows GUI environment. You can run it from a console, use a VNC connection from a remote computer, or set the DISPLAY environment variable to point to another UNIX computer that has an X-Windows environment.

Preparing the DB2 UDB Server Computer

This section describes how to ensure that the server computer has the appropriate software and is configured to allow SQDR Plus to perform incremental operations.

Verifying the DB2 UDB Version

Make sure the DB2 UDB server software has been installed on the host computer and is v8 with FixPak 10 or later, or v9.1 or later. You can enter the **db2level** command at a command prompt to display the version and service level (build level and FixPak number) of the DB2 instance.

For example, a typical result of running the **db2level** command from a Windows command prompt on a computer that has DB2 UDB v9 installed would be:

W

```
DB21085I  Instance "DB2" uses "32" bits and DB2 code release "SQL09010"
with level identifier "01010107".
Information tokens are "DB2 v9.1.0.189", "n060119", "", and Fix Pack "0".
Product is installed at "c:SQLLIB" with DB2 Copy Name "db2build".
```

UNIX

On Linux and UNIX computers entering **db2level** from the command line shows the Release, Level, and the informational tokens, such as:.

```
DB21085I  Instance "db2inst2" uses "64" bits and DB2 code release
"SQL08027" with level identifier "03080106".
Informational tokens are "DB2 v8.1.3.128", "s061108", "MI00180", and
FixPak "14".
Product is installed at "/opt/IBM/db2/V8.1".
```

Verifying the DB2 Path

The SQDR Plus Installation Wizard invokes DB2 commands to install the software. To verify that you can access the DB2 command processor:

W UNIX

1. Enter **db2cmd**, and then enter the command **db2**.

   Run **db2profile** (sh/ksh) or **db2cshrc** (csh) in the sqllib directory of the instance owner, such as shown below, and then enter the command **db2**.

sh/ksh	$ ./home/db2inst1/sqllib/db2profile
csh	% source /home/db2inst1/sqllib/db2cshrc

2. If these commands fail, edit the PATH environment variable to include the directory in which DB2 is installed.

Setting Up the Java Environment

SQDR Plus uses Java in three different contexts, and the version of Java can be different for each use. The contexts in which Java must be available are:

- for SQDR Plus GUI installation, which locates Java from the PATH variable set for the installing user;
- for the SQDR Plus Capture Agent program, which locates Java according to the PATH set for the account that the service runs under (Local System account on Windows and Run As User on Linux and UNIX) or a parameter in conf/wrapper.conf; and
- for stored procedures, which locate Java according to the DBM configuration parameter JDK_PATH (32-bit instances) or JDK_64_PATH (64-bit instances).

The SQDR Plus installation uses the Java stored procedure, SQLJ.INSTALL_JAR, which is provided by IBM. Therefore, the Java stored procedure environment must be available before you can install SQDR Plus. In DB2 v8, DB2 UDB for Linux, Windows, and some versions of UNIX include the IBM Java SDK, whereas Solaris and HP-UX installations of DB2 UDB rely on the manufacturer's JVM to already be installed in the typical location. In DB2 v9.1 and later, IBM supplies a Java SDK for most platforms.

You must have a compatible JDK and JVM installed on your database server. Each operating system supports different levels of the JDK, and IBM maintains a compatibility chart at http://www-306.ibm.com/software/data/db2/ad/v8/java.html to help users determine which JDK/JVM levels are supported. Product requirements for DB2 UDB v8 and v9 are provided at http://www-306.ibm.com/software/data/db2/9/sysreqs.html.

Be sure that your Java environment is properly set up, as described in the following sections, before you run the SQDR Plus installation Wizard.

Verifying the JVM Version

The JVM is installed when you install DB2 UDB. If you have multiple versions of Java installed, make sure that you are using at least v1.4 with SQDR Plus. Specify which Java version to use by setting the PATH system variable, as described in the section "Setting the CLASSPATH for Java" on page 18.

1. Type **`java -version`** at a command prompt to display the active version of Java.

   If you are installing SQDR Plus on Solaris or HP-UX, issue the following command to verify that you have a 64-bit mode of Java available before you begin the installation.

   ```
   java -d64 -version
   ```

If the Java version is not returned, or is earlier than v1.4, download the JVM software from the IBM Fix Central or developerWorks website or, if you are running Solaris or HP-UX, from the Oracle or Hewlett-Packard website, respectively.

On most computers that support both 32-bit and 64-bit mode you have a choice during installation of whether you want to install the 32-bit or 64-bit version of the SQDR Plus Capture Agent. However, on a Linux x64 computer you must use the 64-bit Java if you want to install the 64-bit version of SQDR Plus.

Setting the CLASSPATH for Java

The path to the JDBC drivers is specified by the CLASSPATH environment variable.

On a UNIX computer, sourcing db2cshrc or running db2profile as described in “Verifying the DB2 Path” on page 16 also sets the CLASSPATH and the library path (LD_LIBRARY_PATH, LIBPATH, or SHLIB_PATH) to the appropriate jar files.

Use the System Control Panel to review and set the CLASSPATH variable on a Windows computer. Review the CLASSPATH variable to ensure it specifies the correct path to the IBM JDBC drivers `db2jcc.jar`, `db2jcc_license_cu.jar` and `db2java.zip`.

Verifying the JDK Configuration for the Database Instance

Each operating system supports different levels of the JDK and different types of database instances.The characteristics of the database instance determine the type of JDK required. For example, a 32-bit database instance requires a 32-bit JDK, and a 64-bit instance requires a 64-bit JDK.

Multiple versions of Java can be installed on the same system, in which case DB2 reads the JDK_PATH database manager configuration parameter to determine which to use. To run DB2 Java routines, such as stored procedures, the DB2 database manager configuration on the server must include the path that specifies where the appropriate JDK is installed.

1. To verify the value for the JDK_PATH field, enter the following command in a DB2 command window on the server.

   `db2 get dbm cfg`

2. In the Database Manager Configuration list, look for the value that appears for JDK_PATH (or JDK_64_PATH if you are running a 64-bit instance) to see the configured location of the JDK, such as in the following example for a Windows computer.

   ```
   Java Development Kit installation path    (JDK_PATH) = C:\Program
   Files\IBM\SQLLIB\java\jdk
   ```

3. If you need to set the path where the JDK is located, enter the following on the server command line, replacing the pathname as appropriate for where Java is installed on your computer.

   `db2 update dbm cfg using JDK_PATH` ***`/home/db2inst/jdk14`***

   If the server computer supports both 32-bit and 64-bit instances (AIX, HP-UX, Solaris, or Windows) use the **`jdk_64_path`** configuration parameter instead of the jdk_path parameter. The jdk_64_path parameter specifies the directory under which the 64-bit version of the JDK is installed.

4. If you change the JDK_PATH parameter, enter the following commands in a DB2 command window to stop and restart the DB2 instance so the change takes effect.

```
db2stop
db2start
```

If you are running the HP-UX or Linux operating systems you need to perform a few additional steps to successfully run stored procedures on the server computer. These steps are described in the following sections for each platform.

Setting Up the HP-UX Java Environment

To run Java stored procedures on a HP-UX computer with a 64-bit database instance you need to use the db2hpjv tool to enable Java routine support, which is disabled by default, and create symbolic links to the Java shared libraries, as described in the following steps.

1. Enter the following commands on the command line to enable the db2hpjv tool.

```
db2hpjv -e
db2stop
db2start
```

The HP-UX run-time linker must be able to access Java shared libraries and the DB2 system must be able to load the shared libraries and the JVM. The loading program runs with setuid privileges and therefore looks for the dependent libraries only in /usr/lib/pa20_64, so you need to create symbolic links to the Java shared libraries. You can create symbolic links in /usr/lib/pa20_64, or you can add the name of the directory that stores the Java shared libraries to the `/etc/dld.sl.conf` file.

2. Log in as `root` user and use one of the following methods to create symbolic links to the Java shared libraries.

   - Add a directory to the configuration file:

     Add the `/opt/java1.4/jre/lib/PA_RISC2.-W` and `opt/java1.4/jre/lib/PA_RISC2.0W/hotspot` directories to the /etc/dld.sl.conf file.

   OR

   - Create symbolic links in /usr/lib/pa20_64:

```
ln -a /opt/java1.4/jre/lib/PA_RISC2.0W/*.al /usr/lib/pa20_64
ln -s /opt/java1.4/jre/lib/PA_RISC2.0W/hotspot/*.al
/usr/lib/pa20_64
```

Setting Up the Linux Java Environment

To run Java stored procedures, the Linux run-time linker must be able to access certain Java shared libraries, and DB2 must be able to load both these libraries and the Java virtual machine. After you ensure that the JDK and JVM are available, you may need to create symbolic links in `/usr/lib` (for 32-bit version) or `/usr/lib64` (for 64-bit version) to point to the Java shared libraries, or you can add the name of the directory that stores the Java shared libraries to the `/etc/ld.so.conf` file.

Specifying the Java Shared Libraries Directory. If you add the Java shared libraries to `/etc/ld.so.conf`, you must execute the `ldconfig` command as `root` user to activate the changes, such as shown below for 32- and 64-bit instances.

32-bit Instance:

```
/opt/IBMJava2-142/jre/lib
/opt/IBMJava2-142/lib
/opt/IBMJava2-142/jre/bin
/opt/IBMJava2-142/jre/bin/classic

ldconfig
```

64-bit Instance:

```
/opt/IBMJava2-amd64-142/jre/lib
/opt/IBMJava2-amd64-142/lib
/opt/IBMJava2-amd64-142/jre/bin
opt/IBMJava2-amd64-142/jre/bin/j9vm

ldconfig
```

Restart DB2 for the changes to take effect.

Creating Symbolic Links in the User Library. The number of required libraries that you need to link to depends on the version of JDK that you are using. Examples of the symbolic links for both 32-bit and 64-bit instances of DB2 with JDK v1.4.2 are illustrated below.

32-bit Instance:

```
JAVAHOME=/opt/IBMJava2-amd-142
cd /usr/lib
     ln -fs $JAVAHOME/jre/bin/libjava.so .
     ln -fs $JAVAHOME/jre/bin/classic/libjvm.so .
     ln -fs $JAVAHOME/jre/bin/libhpi.so .
     ln -fs $JAVAHOME/jre/bin/libjsig.so .
     ln -fs $JAVAHOME/jre/bin/libjitc.so .
     ln -fs $JAVAHOME/jre/bin/libxhpi.so .
     ln -fs $JAVAHOME/jre/bin/libdbgmalloc.so .

ldconfig
```

64-bit Instance:

```
JAVAHOME=/opt/IBMJava2-amd64-142
cd /usr/lib64
     ln -fs $JAVAHOME/jre/bin/libjava.so .
     ln -fs $JAVAHOME/jre/bin/j9vm/libjvm.so .
     ln -fs $JAVAHOME/jre/bin/libjsig.so .

ldconfig
```

Setting Database Authority for SQDR Plus Installation

This section applies only if you plan to use the GUI installer to install SQDR Plus on a Linux- or UNIX-based computer and you want SQDR Plus to manage the log files as described in “DB2 Database Logging” on page 26. You must run the installation program as `root` user and the `root` user must be defined in a group that has authorization to run the DB2 UPDATE CFG statement. This requires the authorization of SYSADM, SYSCTRL, or SYSMAINT group.

1. Enter the following command in a DB2 command window to determine the database manager configuration value for the SYSADM_GROUP, SYSCTRL_GROUP, or SYSMAINT_GROUP.

   ```
   db2 get dbm cfg | grep [SYSADM_GROUP | SYSCTRL_GROUP |
   SYSMAINT_GROUP]
   ```

2. Edit the etc/group to add the `root` user to a group specified in SYSADM_GROUP (a group named DB2GRP1 typically is created during installation) or to a group in SYSCTRL or SYSMAINT.

3. Enter the following commands to restart the DB2 instance so the changes take effect.

   ```
   db2stop
   db2start
   ```

Database DSN and Logon Permissions for Clients

For a client computer to access the DB2 UDB database it must have an ODBC driver installed, such as the StarQuest StarSQL driver, and have a Data Source Name (DSN) to the source database defined. If you are using the StarSQL driver with DB2 UDB v8, make sure that you have installed FixPak 10 or later as FixPak 8 and 9 prevented proper binding of catalog packages due to APAR JR21670.

The user ID that the ODBC data source uses when connecting from the SQDR service that is running on a UNIX or Windows computer needs the following permissions:

- read access to tables that are being replicated
- access to run the stored procedures of the Capture Agent
- read access to the Capture Agent schema and its contents

Preparing the Target Database

Incremental replications rely on unique fields to identify change data. If the collation sequence of the database is Case Insensitive, character fields that are uniquely identified on the host may no longer be unique when they are replicated to the target. For example, fields named “First” and “first” that are unique on the source would not be unique if the target database is configured to be insensitive to capitalization. Certain data types that have different levels of precision in different database systems also can create non-unique fields. For example, DB2 on an OS/400 host supports micro-second timestamps whereas the SQL Server DBMS only supports milli-second precision.

To help ensure unique fields, configure the collation sequence of the target database, especially if it is a SQL Server database, to be Case Sensitive. If you have SQL Server Enterprise Manager (SQL Server 2000) or the SQL Server Management Studio (SQL Server 2005 & later), you can verify the collation sequence by viewing the database properties.

A case-sensitive collation has the characters “CS” appended to the Collation Designator. For example, the collation SQL_Latin1_General_CP1_CS_AS is a case-sensitive collation for U.S. English systems. You can set the collation for a new database using the SQL Server Enterprise Manager or the SQL Server Management Studio. To change the collation for an existing database, use the ALTER DATABASE (Transact-SQL) command with the COLLATE clause. Refer to your database system documentation for details on specifying the collation sequence for databases it manages.

Planning the Installation

The following sections describe the prompts that appear when you are installing SQDR Plus. Read these sections and complete Table 4 on page 29 with the values appropriate for your environment. This preparation will help you install and use SQDR Plus in the most efficient manner.

DB2 Installation Directory

When you install SQDR Plus you specify the directory where DB2 UDB is installed. The default location depends on the operating system environment, as indicated in the following table.

Table 1: Default Installation Directory for DB2

Operating Environment	Default Installation Directory	
	DB2 v8.2	**DB2 v9.n (where n=1, 5, or 7 for DB2 v9.1, v9.5, or 9.7**
AIX	/usr/opt/db2_08_01	/opt/IBM/db2/V9.n
HP-UX and Solaris	/opt/IBM/db2/V8.1	/opt/IBM/db2/V9.n
Linux	/opt/IBM/db2/V8.1	/opt/ibm/db2/V9.n
Windows	C:\Program Files\IBM\SQLLIB\	C:\Program Files\IBM\SQLLIB\

With DB2 v9 you can have multiple DB2 copies on the same computer. Use the Default DB2 Selection Wizard to select which DB2 copy will be used by default.

Run the **`db2swtch.exe`** command, which is located in the sqllib\bin directory of the DB2 copy, to launch the Default DB2 Selection Wizard on a Windows computer.

Use the **`db2ls`** command to list the paths where DB2 is installed on Linux and UNIX computers.

Database Name

The SQDR Plus installation is associated with a particular database, and you specify the name of the database during installation. The SQDR Plus installer will show a list of all the databases that are defined for the current DB2 instance. If you have multiple instances, or multiple DB2 copies with multiple instances, specify the instance that you want to use with SQDR Plus as the current instance before you run the installation.

Issue the following DB2 command to show the current database manager instance:

```
db2 get instance
```

Use the Control Center or issue the **db2ilist** command to view a list of instances. If you are using DB2 UDB v9 or later, only the instances from the DB2 copy where the **db2ilist** command is invoked are included in the list.

```
db2ilist
```

To set the DB2INSTANCE environment variable for Windows enter the command:

```
set db2instance=<instance_name>
```

To set the DB2INSTANCE environment variable for Linux or UNIX, execute the db2profile or db2cshrc script in the sqllib directory of the instance owner's home directory.

Host and Port

The default values for the host IP and port are 127.0.0.1 (localhost) and 50000. You can specify a different IP address or the computer host name, and any port number on which DB2 is listening for connections. You can review or set the port that DB2 is using from the DB2 Control Center.

1. Open the DB2 Control Center. (For DB2 v9 the Control Center can be run only from a Linux or Windows computer.)

   

   In Windows click Start–>Programs–>IBM DB2–>General Administration Tools –>Control Center.

   

   In Linux and UNIX open the IBM DB2 folder on the desktop and click Control Center, or enter the command **db2cc** to start the Control Center.

2. Select the database instance and select Setup Communications from the Object Details pane or the context menu.

3. In the Setup Communications dialog, click on the Properties(s) button to review or set properties for the TCP/IP protocol.

The Configure TCP/IP - DB2 dialog shows the port number that DB2 uses for TCP/IP communications. As shown in the following example the db2c_DB2 service on the computer GrayGhost is listening for connections on port 50000.

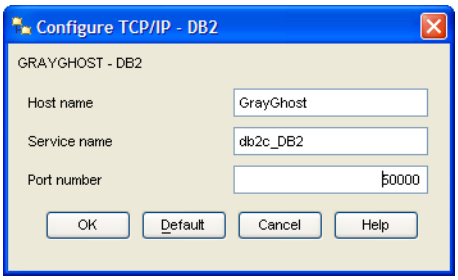

DB2 User ID and Password

DB2 obtains information about an authenticated DB2 user to determine what database operations the user may perform, and what data objects may be accessed. With each request there may be more than one authorization check, depending on the objects and operations that are involved. Authorization is performed using DB2 facilities, DB2 tables, and configuration files to record the permissions associated with each authorized username. When an authenticated user tries to access data, the authorization name of the user, and those of the groups to which the user belongs, are compared with the recorded permissions to determine whether access is allowed.

To install SQDR Plus you must be logged on to the computer as a user with Administrative rights (Windows) or with `root` authority (UNIX). As you install SQDR Plus you also must specify a DB2 user name and password that has access to the DB2 UDB database that SQDR Plus will monitor. Create these accounts before you start the installation if they do not already exist.

The DB2 user account needs authority to:

- Connect to the database
- Create tables
- Create packages
- Create and register routines
- Create schemas
- Create external routines

The installation process grants PUBLIC access to the SQDR Plus stored procedures and schema.

Database Configuration Directory

SQDR Plus stores configuration information and operational log files in a directory that you specify during the installation. The default location for this directory depends on the operating system environment, as shown in the following table.

Table 2: Default Database Configuration Directory

Operating Environment	Default Database Configuration Directory
AIX, HP-UX, Linux, and Solaris	/var/sqdrplus/<*database_name*>
Windows XP and 2003	C:\Documents and Settings\All Users\Application Data\StarQuest\sqdrplus\<*database_name*>\
Windows Vista and later	C:\ProgramData\StarQuest\sqdrplus\<*database_name*>

The Application Data folder is hidden on a Windows computer if the folder options are configured to hide system files and folders. Change the folder options to “Show hidden files and folders” if the Application Data folder is not visible.

SQDR Schema and Tablespace

A ***schema*** is an identifier that provides a logical grouping for database objects. The schema stores the components owned by the application, and it can store the database objects on which the components are based. A schema also qualifies names to prevent ambiguous references to objects and avoid namespace collisions.

During installation you specify a schema name and tablespace for SQDR Plus to store the components it needs to manage incremental change data for the database. The default value for the schema is SQDR and the default tablespace name is *DEFAULT.

SQDR Plus Capture Agent Port and Logging Level

During the installation of SQDR Plus you specify the port number that the SQDR Plus Capture Agent components use to communicate with the DB2 UDB system. This port does not need to be visible from any remote computer. The default port number is 50005.

The logging level that you configure determines the amount and type of operational information that is saved to SQDR Plus log files. There are eight levels of logging, and you also can turn off the logging activity by specifying OFF. As reflected in the following illustration, the broad levels of

logging can significantly degrade performance of all computers involved in the communications. The default logging value is INFO, which logs information about when the SQDR Plus Capture Agent is started and stopped and when it encounters warning or severe error messages.

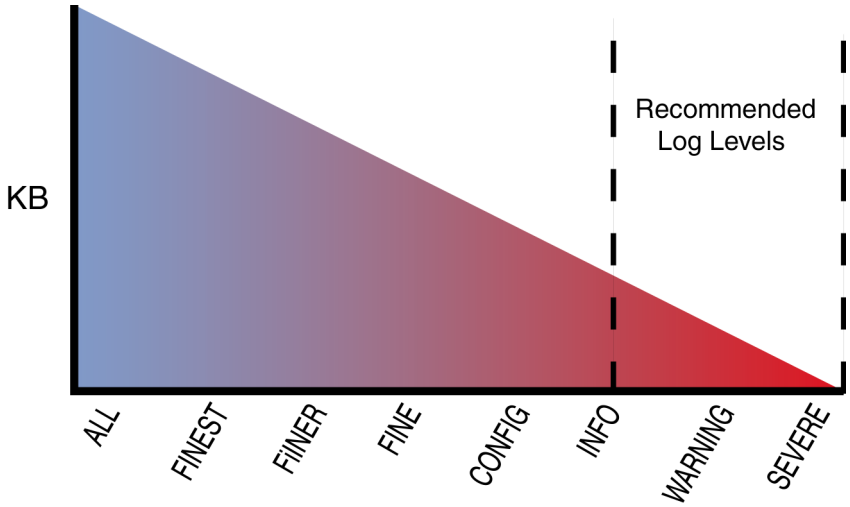

Restricting Replication of Tables

The Capture Agent Maintenance utility allows you to specify which tables in the database are available for replication. During installation you specify whether you want to use this publishing feature. Enable the checkbox for the Restrict Subscriptions to Published Tables option during the installation process to use the table publishing feature. Leave the checkbox disabled if you want all the database tables to be available for replication operations.

DB2 Database Logging

All databases have logs associated with them to record database changes. Logs keep a record of changes made to database objects and data. DB2 provides for ***circular*** and ***archive*** logging. Circular logging reuses the log files, which can overwrite transactions and prevent rollforward recovery. Archive logging does not overwrite log files, which allows rollforward recovery that uses both archived logs and active logs to recover transactions and restore a database.

The database must be configured to use archive logging for SQDR Plus to access the log files it needs to ensure incremental replication operations are applied properly. Circular logging is the default when a DB2 UDB database is created.

During the installation you can choose to have SQDR Plus manage the log files it needs for tracking changes to the database. If SQDR Plus manages the log files it installs a user exit program that archives and retrieves the log files for that database manager instance. If the database is configured for circular logging, the installer prompts you to approve changing the logging to archive, in which case it also installs the exit program and performs an offline backup of the database. If the database already is configured for archive logging and you choose to have SQDR Plus manage the log files the installer installs only the SQDR Plus exit program. If you have

SQDR Plus manage the log files you also can choose whether you want to delete archived logs when they are no longer needed, or save the logs, in which case you also specify a directory in which to store them.

The SQDR Plus installation proceeds more quickly if the database is already configured for archive logging. You can use the DB2 Control Center to enable archive logging for the database before you install SQDR Plus. From the DB2 Control Center select the database, right-click, and select the Configure Database Logging command to change the type of logging to Archive and specify other logging parameters.

Note that there can be only one user exit program, which is named `db2uext2`, per database manager instance. The exit program also must match the bit-width of the DB2 instance. On a Windows computer the exit program is installed in C:\Program Files\ibmj\sqllib\bin. On a Linux- or UNIX-based computer, the exit program is installed into `~<instance_owner>/sqllib/adm` directory if you choose to have SQDR Plus manage the log files. The DB2 `logarchmeth1` configuration parameter must be set to USEREXIT for the user exit program to archive and retrieve log files; the SQDR Plus installation Wizard automatically sets the `logarchmeth1` parameter if you choose to have SQDR Plus manage the log files.

If your database instance requires a custom user exit program, it must ensure that the log files remain accessible until they are no longer needed by SQDR Plus.

Capture Agent Service Parameters

The service parameters include an option for running the Capture Agent in 64-bit mode on some computers. If this option is not enabled during installation, SQDR Plus installs the 32-bit Capture Agent.

If you install SQDR Plus on a Windows-based computer you can specify the name of the Windows service and how it starts. The default name for the Windows service is SQDRPlus along with the database name, in the format of SQDRPlus$*database_name* for the service name and a display name of SQDRPlus(*database_name*). You also can enable an Automatic Startup option to have the SQDRPlus service start each time the operating system is started, and an option to start the service when the installation of SQDR Plus completes.

On a Linux or UNIX computer the SQDR Plus service daemon name is sqdr-<*database_name*>. During installation you can choose to enable the daemon for Automatic Startup, which creates the appropriate symbolic links to the service. Note that DB2 must be started before the SQDR Plus service daemon can start. Use the Fault Monitor Coordinator (FMC) to monitor each DB2 instance for availability. If the FMC detects that the DB2 instance exits prematurely, it restarts the instance. Refer to your DB2 UDB documentation for more information about starting the FMC using the `init` daemon.

The installer for Linux and UNIX also includes a prompt for the user ID that the Capture Agent will run as. The default value for the Run As User prompt is the installing user. The installing user must have `root` authority, whereas the Capture Agent service can run as a user with only enough privilege to run Java and write the PID file and `wrapper.log` files. The Run As User account must exist before you begin the SQDR Plus installation. If you want the Capture Agent service to

run with the minimal privileges, create a user account, such as *sqdr*, that you can specify for the Run As User prompt during the SQDR Plus installation. The installation Wizard automatically grants the privileges required to run Java and write the PID file and wrapper.log files.

 **NOTE for AIX USERS**	There is a known issue using the GUI-based installer on an AIX computer which requires that you install SQDR Plus using 32-bit Java and then use the 64-bit Java to run the Capture Agent. We recommend using the text-based installer to avoid this complication. If you do use the GUI-based installer, be sure the PATH variable is set of the installing user is set to use the 32-bit Java, and that the `.profile` file of the Run As User account is set to the location of the 64-bit Java.

Configuring Email Notifications

During the installation of SQDR Plus you can specify that notifications be sent to an email address to inform an administrative user of an error condition. The logging level determines how severe the error condition must be to warrant sending an email notification. If the Capture Agent encounters a problem within the level of logging that you select (CONFIG, INFO, WARNING, SEVERE), it sends an email to the specified email address using the specified SMTP server. The recommended notification level is for SEVERE error conditions, as the more general levels of errors (WARNING, INFO, and CONFIG) can generate a large quantity of email.

You can specify the SMTP server name or IP address of any SMTP mail server that does not require authentication. The SMTP mail server must be able to direct email to the user account specified in the "To" parameter and accept *any value* as the sender. You may want to use the "From" field to identify the Capture Agent on a particular computer, such as SQDRPlus@*computer_name*.

SQDR Plus Program Directory

During the installation you select the directory into which you want to install SQDR Plus program files. The default location for this directory depends on the operating system environment, as shown in the following table.

Table 3: Default SQDR Plus Program Directory

Operating Environment	Default SQDR Plus Directory
Linux or UNIX	/opt/StarQuest/sqdrplus
Windows	C:\Program Files\StarQuest\sqdrplus

Gathering Configuration Information

Table 4 lists all of the configuration parameters that you need to specify during installation. Make a note of the values appropriate for your environment before you start the installation program.

Table 4: Properties Configured During Installation

Configuration Parameter	Default Value	Your Value
Location of DB2 Installation	see Table 1 on page 22	
Database Name	none	
Host	127.0.0.1	
Port	50000	
DB2 User ID	none	
DB2 User Password	none	
Location for Database Configuration	**Linux** or **UNIX**: /var/sqdrplus/<*database_name*> **Windows XP/2003:** C:\Documents and Settings\All Users\Application Data\StarQuest\sqdrplus\<*database_name*> **Windows Vista/2008/7:** C:\ProgramData\StarQuest\sqdrplus\<*database_name*>	
SQDR Schema	SQDR	
Schema Tablespace	*DEFAULT	
Capture Agent Port	50005	
Logging Level	INFO (see “Configuring the Capture Agent” on page 37 for other valid values)	
Restrict Subscriptions to Published Tables	disabled (no checkmark)	
Allow SQDR to Manage Logs	Yes if database configured for Circular logging; No if database configured for Archive logging.	
Delete Archived Logs without Saving	Yes (no effect unless Allow SQDR to Manage Logs is enabled)	
Archived Logs Directory	none	
Windows Service Name	**Linux** or **UNIX**: sqdr-*database_name* **Windows:** SQDRPlus$*database_name*	

Configuration Parameter	Default Value	Your Value
Windows Service Display Name	**Linux** or **UNIX**: system-assigned PID **Windows:** SQDRPlus(*database_name*)	
Run as User	**Linux** or **UNIX**: root **Windows:** not applicable	
Run Service in 64 bit mode	enabled if supported on the target computer	
Automatically Start Windows Service	Yes	
Start Capture Agent Service (when installation is complete)	Yes	
Email Notification Information	OFF (see page 28 for description of the email notification parameters)	SMTP Server: From: To: Notification Level:
Directory for SQDR Plus Java Programs	**Linux** or **UNIX:** /opt/StarQuest/sqdrplus **Windows:** C:\Program Files\StarQuest\sqdrplus	

To change any of these properties after you install SQDR Plus you must edit the sqagent.properties file, as described in “Configuring the Capture Agent” on page 37.

Running the SQDR Plus GUI Installation Program

You run the SQDR Plus installer directly from the Linux, UNIX, or Windows computer on which you want to run SQDR Plus. The computer must have Java installed and configured as described in “Setting Up the Java Environment” on page 17.

Refer to the Quick Start Guide to Using SQDR Plus for DB2 UDB on Windows for additional information on installing SQDR Plus on Windows.

UNIX

To run the GUI SQDR Plus installer on a UNIX-based computer you need to have an X-Windows GUI environment. You can run the installer from a console, use a VNC connection from a remote computer, or set the DISPLAY environment variable to point to another UNIX computer that has an X-Windows environment.

See the next section for instructions on running the UNIX text-based installer.

The SQDR Plus for DB2 UDB installer consists of the files `DATA1.JAR`, `media.inf`, and a batch/shell script named `setup`.

1. Copy the `DATA1.JAR`, `media.inf`, and `setup` files to the computer on which you want to run the installation. On Windows, we recommend copying these files to a temporary directory on the same drive as your Windows directory.
2. Log on to the computer as a user with Administrator rights (Windows) or as a user with `root` authority (Linux and UNIX).

UNIX

3. If you are installing SQDR Plus onto a Linux- or UNIX-based computer, be sure the Run As User account has been created and configured as described in “Capture Agent Service Parameters” on page 27.
4. Change to the directory where you copied the SQDR Plus installation files.

W

5. Execute the `setup.bat` file if you are running the installation from a Windows computer.

UNIX

Execute the `setup` shell script if you are running the installation from a Linux or UNIX computer.

NOTE for Windows Vista, Windows Server 2008, and Windows 7 users:

setup.bat should be run in an elevated command window.

1. Locate the Command Prompt under Accessories, right-click and select "Run as Administrator" from the drop down list.
2. You will be prompted with an elevation prompt; respond OK.
3. In the Command Prompt window, navigate to the SQDR Plus installer and type "setup.bat".

The setup script initializes the InstallShield Wizard and displays a Welcome dialog. If the computer already has a VPD Registry for InstallShield, an “Initializing wizard” status dialog appears while the installer reads the VPD Registry. The Wizard checks for pre-requisites such as the presence of an appropriate version of Java and determines whether SQDR Plus is already installed. It also checks that the installing user has appropriate authority, and displays a warning dialog if the requisite software and authority are not adequate.

4. Click Next to proceed with the installation process after the Welcome dialog appears.
5. After you review the License Agreement for using the software, click Next to accept the terms and proceed with the installation.

The InstallShield Wizard then begins prompting for values that are required to configure the Capture Agent.

6. Enter the configuration values that you noted in Table 4, “Properties Configured During Installation,” on page 29, and click Next to proceed through the Wizard panes.

After you supply the configuration parameters, a Summary pane appears before the installation phase begins. Click the Back button if you need to review or change any configuration values. If you need to change these properties after you install SQDR Plus, refer to “Configuring the Capture Agent” on page 37.

7. After you review the installation summary information, click Install to begin the installation phase, and click Next as the Wizard displays panes that show the task results. If an error is encountered during the installation, back up to the prior installation pane and correct the error.

The installation phase installs the SQDR Plus service program and other objects to the specified schema, installs the Java .jar files and property files to the target directory, and creates an uninstaller and registers with the VPD.

After installing the necessary files and updating the VPD Registry, the Wizard displays a summary pane to indicate whether SQDR Plus was successfully installed.

8. Click Finish after you review the summary information.

If you enabled the option for starting the SQDR Plus Capture Agent service after installation, it is started after you click Finish.

Running the SQDR Plus Text-based Installation Program

The text-based installer involves editing a configuration file with a text editor such as vi or gedit, and requires fewer system resources than the GUI installer. In addition, support for some platforms is available only in the text-based installer, and future installer development will be oriented to enhancing the text-based installer.

You can use the text-based installer to upgrade from an existing installation performed with the GUI installer, after which future upgrades should also be done with the text-based installer.

Refer to the Quickstart Guide to Using SQDR Plus for DB2 UDB on Linux/UNIX for step-by-step instructions for installing SQDR Plus on UNIX with the text-based installer.

Before you install:

The text-based installer expects that the database is already configured for the method of logging you plan to use. For instance, if you select `managedLogs=true` (the default), the database should already be configured for `LOGARCHMETH1=USEREXIT`. If the database is configured for circular logging (`LOGARCHMETH1=NONE`) or if you select `managedLogs=true` and the database is configured for anything other than `USEREXIT`, an error message is displayed and the installation is aborted.

If you are using sh/ksh/bash, modify `/home/<db2inst_owner>/sqllib/userprofile`; see the Quickstart Guide to Using SQDR Plus for DB2 UDB on Linux/UNIX.

Additional pre-requisites and pre-installation tasks are described in the Quickstart Guide to Using SQDR Plus for DB2 UDB on Linux/UNIX.

Editing the configuration file:

If this is a new install, copy `setup.conf.template` to `setup.conf` and edit it. Refer to the comments in setup.conf and to the Quickstart Guide to Using SQDR Plus for DB2 UDB on Linux/UNIX for detailed information.

If you are upgrading from an existing installation performed with the text-based installation, copy setup.conf.template to `setup.conf` and confirm that the value of `INSTALLDIR` (default `/opt/StarQuest/sqdrplus`) in `setup.conf` is correct. All other values will be obtained from the configuration file used during the original installation.

If you are upgrading from an installation performed with the GUI installer, copy `setup.conf.template` to `setup.conf` and edit it using the values used during the previous installation. Refer to the comments in `setup.conf` for detailed information.

Running the installation:

1. Log on to the UNIX/Linux computer as a user with root authority.
2. Set up access to DB2 by sourcing db2profile (sh/ksh) or db2cshrc (csh) in /home/<*db2inst_owner*>/sqllib.
3. Change to the directory where you extracted the installer and edited setup.conf.
4. Run the installation:

```
# ./setup [-v] [-c]

-v = Verbose (optional)

-c = Recreate control tables during an update (optional;
ignored during a new installation)
```

The installer will attempt to verify whether this is a new installation, an update from a previous text-based installation, or an update from a previous GUI installation. It will ask for confirmation of the type of installation before proceeding.

You will be prompted for the password of the DB2 user configured in `setup.conf`; this user and password are used during installation for creating control tables and registering stored procedures, and are used during operation by SQDR Plus. The same DB2 user will be used during uninstallation, at which time you will be prompted for its password.

After installation, examine the logs in `/var/sqdrplus/<database_name>/logs` for unusual errors. Some errors are expected; for instance, after a new installation, you can ignore errors in `AgentUtilR.log`.

Changing the DB2 user password:

If the password of the DB2 user should change at any time after installation of SQDR Plus, run the following command to notify SQDR Plus of the change:

1. `# cd /opt/StarQuest/sqdrplus` (or wherever you installed to)
2. `# ../change_password`

You will be prompted for the new password.

Troubleshooting Installation Problems

This section addresses problems that can occur during installation that prevent the installation from running or prevent the Capture Agent from starting.

Setup Program Will Not Run

The SQDR Plus setup program invokes **db2cmd** during the installation. If the installation does not run and no errors are reported in the command window or the log.txt, the problem could be due to the character limitations of **db2cmd**. The setup program invokes **db2cmd** with the classpath option (-cp), and a long classpath may exceed the number of characters that **db2cmd** allows. Try running the **setup.bat** file from a command line and review the error messages to determine if the length of the classpath is preventing the program from running.

On a Windows 2000 workstation the **db2cmd** line can contain a maximum of 2047 characters; on Windows XP or 2003 workstations the **db2cmd** line can contain a maximum of 8191 characters. When using **db2cmd** in interactive mode you can enter up to 4095 characters in a command line. If the path configured for the computer is exceeding the **db2cmd** limitations, create a new batch file that specifies only the pathnames needed for Java, DB2, and SQDR Plus, such as shown below for a Windows installation.

```
db2cmd /c /i java -Dis.debug=1 -cp "DATA1.JAR;c:\java;
C:\PROGRA~1\IBM\SQLLIB\java\db2java.zip;C:\PROGRA~1\IBM\SQLLIB\java\db2j
cc.jar;C:\PROGRA~1\IBM\SQLLIB\java\sqlj.zip;C:\PROGRA~1\IBM\SQLLIB\java\
db2jcc_license_cisuz.jar;C:\PROGRA~1\IBM\SQLLIB\java\db2jcc_license_cu.j
ar;C:\PROGRA~1\IBM\SQLLIB\bin;C:\PROGERA~1\IBM\SQLLIB\java\db2jcc_licens
e_cu.jar;C:\PROGRA~1\IBM\SQLLIB\bin;C:\P{ROGRA~1\IBM\SQLLIB\java\common.
jar" run 1>log.txt
```

SQDR Plus Service Does Not Start

If the Capture Agent service fails to start after you install SQDR Plus for DB2 UDB, review the capture_agent0.log and wrapper.log files for errors.

For Windows Server 2003 and earlier, the default directory for these log files is C:\Documents and Settings\All Users\Application Data\StarQuest\sqdrplus\<*database_name*>\conf for the capture_agent0.log, and C:\Documents and Settings\All Users\Application Data\StarQuest\sqdrplus\<*database_name*>\logs for the wrapper.log file. For Windows Vista and later, the conf and logs directories can be found in C:\ProgramData\StarQuest\sqdrplus\<*database_name*>.

Also review the Windows Events for application errors. The Windows Event Viewer is typically located in the Administrative Tools program group.

For a computer running Linux or UNIX the log files are located by default in /var/sqdrplus/<*database_name*>/conf (capture_agent0.log) and /var/sqdrplus/<*database_name*>/log (wrapper.log).

If you need to capture more detail about the problem you can increase the level of information that is logged by the Capture Agent service and/or its Java wrapper. The `wrapper.conf` file, which is located in the `conf` directory, specifies the logging properties for the wrapper. Change the `wrapper.logfile.loglevel` property from INFO to DEBUG to increase the amount of information that is logged about the operation of the wrapper.

The Capture Agent service logs information specified by the `logLevel` parameter of the `sqagent.properties` file. Edit the logging level specified in the `sqagent.properties` file (see page 47) and try restarting the Capture Agent service. The most recent Capture Agent logging information is found in the `capture_agent0.log` file. You can view the log files using Notepad or Wordpad (Windows), vi or gedit (UNIX), or any other application that can read text files.

DB2 SQL Exception

The SQDR Plus installation program accesses DB2 UDB to create the control tables that the Capture Agent needs for tracking change data. An error similar to the following will appear if you are running DB2 UDB v8.2 and upgraded to FixPak 10 or later without rebinding packages.

```
SEVERE: [sqv][main][Sep 26, 2006 1:26:19 PM] CaptureAgentLog:
CaptureAgent.main:
com.ibm.db2.jcc.a.SqlException: DB2 SQL error: SQLCODE: -443, SQLSTATE:
38553,
SQLERRMC: SYSIBM.SQLPRIMARYKEYS;PRIMARYKEYS;SYSIBM:CLI:-805
```

After you apply FixPak 10 or later to a DB2 UDB v8.2 system, you must rebind packages to your existing databases as described in the FixPak README file. If you have not already rebound packages on your DB2 server, perform the following steps.

1. From a Windows computer enter **`db2cmd`** in a Command Prompt window, or select Command Window from Start–>Programs–>IBM DB2–>Command Line Tools, to open a DB2 Command window.

   From a Linux or UNIX computer make sure ~<*instance_owner*>/sqllib/db2profile or db2cshrc has been sourced.

2. Enter the following commands, replacing <*database_name*> with the name of the database to which the utilities should be bound, and <*path*> with the pathname where the bind files are located (typically `sqllib/bnd`).

```
db2 terminate
db2 connect to <database_name>
db2 bind <path>/db2schema.bnd BLOCKING ALL GRANT PUBLIC sqlerror
    continue
db2 terminate
```

Error Registering Stored Procedures During Installation

Stored procedures are registered with DB2 UDB when you install SQDR Plus. If the DB2 UDB system was updated to v8 but the database was not, you will receive the following error when the SQDR Plus installer tries to register the stored procedures:

```
DB2 SQL error: SQLCODE: -440, SQLSTATE: 42884,
SQLERRMC:SQLJ.DB2_INSTALL_JAR;PROCEDURE
```

SQL0440N No authorized routine named routine-name of type routine-type having compatible arguments was found.

To correct the problem run the db2updv8 utility to update the database to v8 format. Open a DB2 Command Window and enter the following command to update the database that you want to use with SQDR Plus.

db2updv8 -d <*database_name*>

Managing and Using SQDR Plus

SQDR Plus is designed to operate with minimal administration after the product is installed and configured. This chapter describes how to change configuration parameters and monitor the product, and provides considerations for optimizing the environment and replication operations. It describes:

- the configuration properties that you can set for SQDR Plus for DB2 UDB
- the logs that can provide insight to the SQDR Plus operations
- how to use the Capture Agent Maintenance utility
- how to remove the SQDR Plus software
- considerations for troubleshooting and tuning the performance of replication operations
- adding an identity column to a source table
- Support for partitioned tables
- Format of the SQ_STATISTICS TABLE

Configuring the Capture Agent

The Capture Agent is controlled by configuration parameters that are specified in a file named `sqagent.properties`, which is located in the Database Configuration directory (the directory specified for database configuration files during installation, as described on page 25).

The values that you specify during the installation of SQDR Plus (shown in Table 4, “Properties Configured During Installation,” on page 29) are saved to the `sqagent.properties` file. There are additional properties in the `sqagent.properties` file that are not presented during the installation of SQDR Plus.

To review or change the configuration values in the `sqagent.properties` file from a Windows computer, select Capture Agent Configuration from the SQDR Plus for DB2 UDB program group. The `sqagent.properties` file opens in Notepad.

On a Linux or UNIX computer you can use vi or any text editor to modify the `sqagent.properties` file, which typically is located in `/var/sqdrplus/<database_name>/conf`.

You must stop and restart the Capture Agent for changes to the `sqagent.properties` file to take effect.Table 5 describes all the properties in the `sqagent.properties` file for SQDR Plus for DB2 UDB.

Table 5: SQDR Plus sqagent.properties File for DB2 UDB

Configuration Parameter	Default Value	Description
archivedLogPath	none	If managedLogs is set to True, the archivedLogPath value specifies the directory in which archive log files are saved when they are no longer needed by SQDR Plus if the deleteLogsWithoutSave parameter is set to false.
broadcastAddress	none	Specifies the default IP address that Capture Agent will send notification to when there is change data available for a subscription that has notifications enabled but does not specify a different IP address to send them to.
clientDeleteGrace	1 (day)	When SQDR Plus detects an inactive subscription, it sends a notification email (type SEVERE) after seven days of inactivity, then automatically deletes the subscription and subscriber after a grace period of one additional day expires, allowing purging of staged data. Setting this value extends the grace period. A value of 0 disables the automatic deletion of subscriptions and subscribers. **Note**: You can recover deleted subscriptions by using the *Reset I/R Group* function from the SQDR Data Replication Manager; a new baseline will be performed.
clientMonitoringInterval	1440 (minutes)	The default behavior of SQDR Plus is to check for inactive clients every twenty four hours (1440 minutes). You can change this interval by setting the clientMonitoringInterval in sqagent.properties; the value is specified in minutes.
controlDbDriver	com.ibm.db2.jcc.DB2Driver	The Java class name of the JDBC driver that is used by the Capture Agent. Do not change the default value unless you are directed to do so by a StarQuest Technical Support engineer.

Configuration Parameter	Default Value	Description
controlDbSchema	SQDR	Name of the schema (collection) that contains the control tables that are used by a particular instance of the SQDR Plus Capture Agent.
controlDbUrl	jdbc:db2://127.0.0.1:50000/SAMPLE:driverType=4;deferPrepares=false;	Java driver information that the SQDR Plus installation Wizard obtains during setup. Do not change the default value unless you are directed to do so by a StarQuest Technical Support engineer.
deleteLogsWithoutSave	true	If managedLogs is set to True the deleteLogsWithoutSave value determines whether archive log files are deleted when SQDR Plus no longer needs them (true), or the archive log files are saved to a directory specified by the archivedLogPath parameter (false).
filterUserId	none	Specifies that the Capture Agent should bypass staging transactions associated with the specified userID. The property can have a value of a single userID. With the property in place, the Capture Agent examine each commit record; if the user-id in the commit matches the "filterUserId", the commit is handled like a rollback, and any staged data is removed
from	none	If the email notification feature is enabled (see smtpServer), the “from” parameter appears in the From field of the email. The “from” parameter can contain any string, such as the Capture Agent schema name and the computer on which it is located, to help identify the source of the email message.

Configuration Parameter	Default Value	Description
grantees	none	Grants the specified users access to the Capture Agent control tables. By default the Capture Agent grants SELECT privileges to PUBLIC. Specify other users, separated by a comma, if you want to limit who can view the Capture Agent control tables and see what subscriptions are active. Note that changing the grantees does not change grants for existing control tables, nor does it grant access to change data that is being staged.
logLevel	INFO (if the default during installation is accepted)	Determines the amount and type of information that is recorded to log files. Valid values are OFF, SEVERE, WARNING, INFO, CONFIG, FINE, FINER, FINEST, and ALL. See the discussion below regarding the level of detail that is logged.
managedLogs	true if the database instance is set to use circular logging	Specifies whether SQDR Plus manages the log files (true), or the database instance manages the log files (false).
numControlDbConnectionsPerVm	10	Specifies the maximum number of database connections that are allowed from a Capture Agent process. Do not change the default value unless you are directed to do so by a StarQuest Technical Support engineer.
maxLogFileSize	512000 bytes	Specifies the maximum number of bytes of changes that the log file can contain.
maxNumLogFiles	5	Specifies the maximum number of log files that can be created to track changes.
minimumNotificationDelay	5000 milliseconds	Determines the minimal time interval that must pass before the Capture Agent sends notification that change data is available for subscriptions that have been enabled to receive notifications.

Configuration Parameter	Default Value	Description
notificationLevel	OFF	Specifies the level of logging messages that are sent via email to the specified user account. Valid values are OFF, CONFIG, INFO, WARNING, and SEVERE.
password	none	The password for the User ID specified for the userId parameter. This value is encrypted and can be changed only by reinstalling SQDR Plus
pollSleepTime	5000 milliseconds	Specifies the interval at which the Capture Agent Log Reader should poll for changes.
port	50005	Specifies the number of a port that the Capture Agent components can use to communicate within the DB2 system. This port does not need to be visible to any remote computer, but it cannot be used by any other application.
pruneInterval	30 minutes	Specifies the interval for examining the Capture Agent control tables to determine what change data can safely be deleted.
requirePublication	false	Specifies whether all the database tables can be subscribed to for replication (false), or only tables that have been published can be subscribed to (true).
retryPublishLock	false	Do not publish a table until an exclusive lock can be obtained. See “Publishing Database Tables” on page 57
smInterval	0 seconds	Specifies whether to use the Storage Monitor; 0 means disabled. See “Monitoring Backlogs with the Storage Monitor” on page 67
smSQL	a SQL statement	See “Monitoring Backlogs with the Storage Monitor” on page 67
smThrottleLevel	500000	See “Monitoring Backlogs with the Storage Monitor” on page 67
smWarningLevel	450000	See “Monitoring Backlogs with the Storage Monitor” on page 67

Configuration Parameter	Default Value	Description
smtpServer	none	To enable email notification. specify a server name or IP address of an unauthenticated SMTP server that can direct the email to the user account specified in the “to” parameter, and specify a notificationLevel value other than OFF.
sourceType	UDB	Must be set to UDB when running SQDR Plus for DB2 UDB.
starquestNotification	SEVERE	Automatically send problem reports via email to StarQuest Support. To disable it, set starquestNotification=OFF in sqagent.properties or uncheck the checkbox during installation or custom update.
TABLEINFO	false	A stored procedure that SQDR uses to find indexes. The default value for SQDR Plus for DB2 UDB is false. Do not change the default value unless you are directed to do so by a StarQuest Technical Support engineer.
to	none	If the email notification feature is enabled (see smtpServer), the “to” parameter must contain a valid email user account in a format such as user@domain.com.
udpPort	2728	The port that Capture Agent uses to send notification to a subscribing SQDR computer that change data is available from a particular source table.
useAuthentication	true	Specifies whether the JDBC driver uses authentication (true) or adopted authority (false). The IBM type 2 driver for UDB can use adopted authority.
userId	none	The user ID of the DB2 account that SQDR Plus uses to access the DB2 control tables.
useTxSequence	false	Specifies that the Capture Agent should maintain the order of transactions. See “Maintaining the Order of Transactions” on page 46

Configuring Database Logging

The SQDR Plus configuration parameters that affect database logging are shown below: Some of the parameters have an effect only if SQDR Plus is managing the log files.

Table 6: Database Logging Configuration Parameters

Database Logging Parameter	SQDR Plus Behavior	DB2 UDB Behavior
managedLogs	true=SQDR Plus manages logs	false=system manages logs
deleteLogsWithoutSave	applicable only if managedLogs=true	no effect
archivedLogPath	applicable only if deleteLogsWithoutSave=false	no effect
maxLogFileSize	applicable only if managedLogs=true	no effect
maxNumLogFiles	applicable only if managedLogs=true	no effect

The effects of these parameters are influenced by whether you choose to have SQDR Plus manage the log files it needs, or whether the DB2 system manages the log files for the database SQDR Plus is monitoring. This choice is indicated by the value of the `managedLogs` parameter, with `true` specifying that SQDR Plus manage the logs, and `false` specifying that DB2 manages the logs. If SQDR Plus manages the log files it installs a user exit program that archives and retrieves the log files for that database manager instance.

Set `managedLogs` to `true` to have SQDR Plus manage the database logs that are used for incremental replication of the source database. If you set `deleteLogsWithoutSave` to `false`, specify the directory where you want to save the archive logs for the `archivedLogPath` parameter. The `maxNumLogFiles` and `maxLogFileSize` parameters control how many log files can be created to track the change data, and how much data each can contain. The optimum values for these parameters depend on how much change data the database generates and in what period of time. If the log files reach the maximum size and number it could prevent replication or rollback activities from completing.

Configuring Database Mirroring

After changes are committed to a source database that Capture Agent is monitoring, the SQDR client retrieves the changes from the staging table on a periodic polling interval. You also can configure the SQDR Plus Capture Agent to send a notification to the client via the User Datagram Protocol (UDP) when change data is available, rather than waiting for the SQDR client to poll for changes. The UDP change data notifications are sent using the port number specified by the udpPort parameter at a frequency no greater than the number of milliseconds specified for the minimumNotificationDelay parameter.

To “mirror” the source database, achieving almost immediate update of the target database when changes are committed to the source, set a small interval for the minimumNotificationDelay and the pollSleepTime parameters in the sqagent.properties file. For example, using the default value of 5000 milliseconds for the minimumNotificationDelay and pollSleepTime parameters would typically provide less than a 5-second lag time in sending notification to the client that changes have been committed to the source, triggering the client to retrieve and apply the changes to the destination database.

Change data notifications also can be useful if a source database is updated infrequently and you want to minimize how often the Capture Agent polls the log files for changes. In this case you could rely primarily on sending UDP notifications when the source database changes, having the SQDR client poll for changes infrequently, such as only once every hour or every day to ensure that changes are applied to the destination database with at least that frequency.

From the SQDR Replication Manager you can configure a group of incremental replication subscriptions to receive notifications at a specific IP address. If a subscription that is enabled to receive notification of change data specifies an IP address of 255.255.255.255 for the notification, the Capture Agent sends a broadcast message to the IP address specified for the broadcastAddress parameter in the sqagent.properties file.

If the server on which SQDR Plus is running has a firewall, be sure the firewall is configured to allow outbound traffic through the UDP port number that you specify for SQDR Plus notifications. The SQDR Plus host computer must allow outbound traffic on the specified `udpPort`, and the computer that the subscription specifies for receiving the change notifications must allow inbound UDP traffic on that port.

For example, if you are using the Windows Firewall, follow the steps below to unblock port 2728,

1. Select Windows Firewall from the Control Panel and click the Exceptions tab of the Windows Firewall property window.
2. Click Add Port and enter a name, such as SQDR Notifications, and the port number you specified for the `udpPort` parameter,
3. Click the UDP protocol radio button and click OK to open the specified port, such as shown below for opening port 2728.

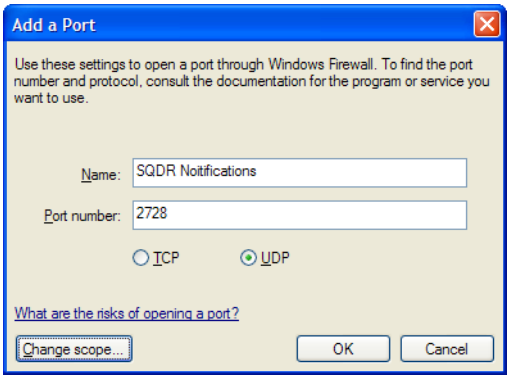

Other types of hardware and software firewalls provide different ways to specify which ports you want to block or whether you want to allow inbound or outbound traffic through them, such as shown below for Zone Alarm Pro.

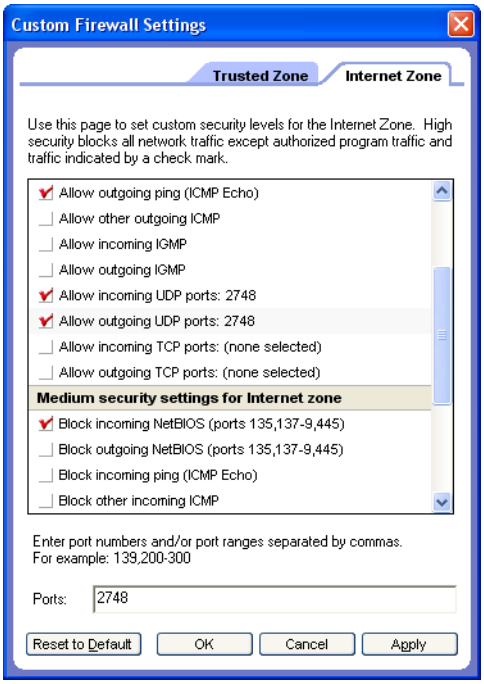

Refer to the information that came with your firewall hardware or software solution if you need further information about how to configure a port to allow UDP traffic.

Configuring Pruning of Control Tables

The `pruneInterval` controls how frequently the Capture Agent control tables are reviewed to delete change data that is no longer needed. The default behavior is for the Capture Agent to prune the control tables every 30 minutes. The optimum pruning interval depends on how much change data is generated and the activity and available storage space of your network. The pruning operations can affect the availability of the Capture Agent such that one or more subscriptions may report an error in polling for changes. This type of error is automatically resolved when the Capture Agent is available again. If there is adequate storage space you can increase the frequency of the polling interval (the `pollSleepTimer` parameter) to allow more change data to accumulate before it is pruned.

Maintaining the Order of Transactions

The default behavior of SQR Plus is to ensure that replications are consistent at transaction boundaries. This behavior can be changed by setting **useTxSequence=true** in sqagent.properties. This property is useful if the target table has any foreign constraints defined on it, such as the target of an incremental replication defined with "unique constraints." The setting preserves the order of

the source updates on the target side and is useful if you are using triggers on the target and need updates to arrive in the order they were performed at the source. The update order is always maintained within each table; this setting is useful if the update order between tables is important.

An example is a single transaction that updates table A, then table B, then table A, then table B. Without this setting, the changes would be applied as A,A,B,B, Commit; with this setting, the order would be A,B,A,B, Commit.

Configuring Operational Logging and Email Notifications

The logging level that you configure determines the amount and type of information about SQDR Plus operations that is saved to log files. There are eight levels of logging, and you also can turn off the logging activity by specifying OFF. For sending email notifications about SQDR Plus operations you can specify only a value of CONFIG to SEVERE to avoid flooding the email server with notifications that would be generated with the more detailed logging levels. SEVERE is the error level that is recommended for sending an email notification. The amount of information that is logged increases as the levels of logging become more broad, as illustrated below.

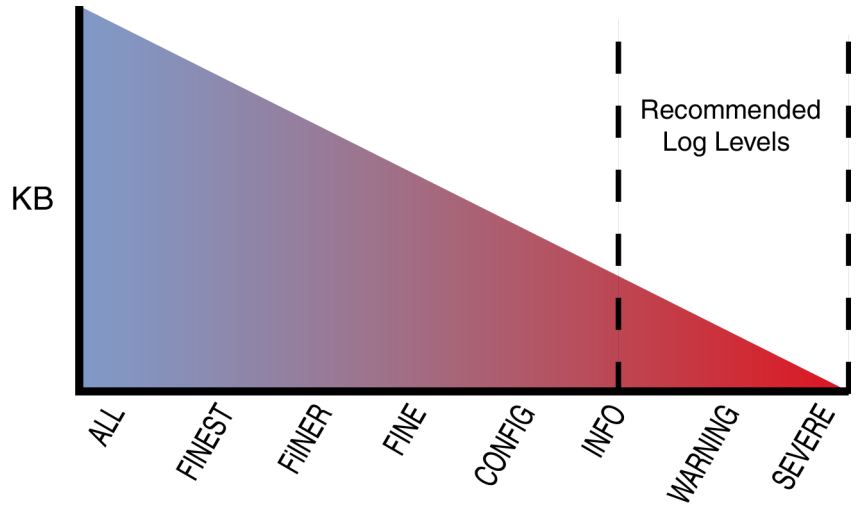

The broad levels of logging can significantly degrade performance of all computers involved in the communications. If you need to enable detailed logging of operations to troubleshoot a problem, remember to restore the prior logging level as soon as you have collected the required diagnostic information.

Restricting Replication of Database Tables

The default behavior of the SQDR Plus Capture Agent is to allow SQDR users to define subscriptions that replicate whatever database object they otherwise have appropriate authority to use. When you install SQDR Plus there is a prompt, Restrict Subscriptions to Published Tables, which sets the initial value (disabled by default) for the `requirePublication` parameter of the `sqagent.properties` configuration file.

After SQDR Plus is installed you can select which tables you want to publish so they can be subscribed to for replication. Publishing and unpublishing database tables provides a centralized method for a database administrator to control what data can be replicated.

After you set the `requirePublication` parameter of the sqagent.properties file to `true` (either by enabling the Restrict Subscriptions to Published Tables during installation or by editing the sqagent.properties file), use the Capture Agent Maintenance utility to specify which database tables are published. Refer to page 58 for details about using the Publish command.

Automatically Starting the Capture Agent Service

The SQDR Plus Capture Agent must remain running to process incremental changes to the source database.

 **NOTE**	After SQDR Plus is monitoring subscriptions for incremental changes there may be repetitive notification messages sent to the administrator if there is a long delay between shutting down and restarting the SQDR Plus Capture Agent service. You can avoid flooding the message queue while the database logs are unavailable by pausing the incremental subscriptions and deleting the connections to the host from the Data Replicator Manager. After the Capture Agent service has been restarted, resume the incremental subscriptions, which automatically creates connections to the host again.

On a Windows computer the SQDR Plus Capture Agent service is configured to start automatically whenever the Windows operating system starts, and it is defined with a dependency to the DB2 instance that is managing the database to help ensure that DB2 also is running. If you need to stop the Capture Agent service or change the logon account, use the Services utility, which typically is found in the Administrative Tools of the Control Panel.

On a Linux or UNIX computer you have an option during installation to configure the service to start automatically, which creates the necessary symbolic links to the Capture Agent service daemon. If these links are not created during installation and you later need to automate starting the Capture Agent, perform the following procedures to create the necessary entries based on the operating system you are running. The default name of the Capture Agent service on a Linux or UNIX system is sqdr-<*database_name*>.

Adding Capture Agent Daemon to AIX Startup

The runtime control (rc) script for the Capture Agent, `/etc/rc.sqdr-<database_name>`, is created when SQDR Plus is installed. Use the **`mkitab`** command to add an entry to the `/etc/inittab` file that instructs the operating system to start the Capture Agent daemon when the AIX operating system starts.

1. Log on as `root` user.
2. Execute the following command at the command line:

   `/usr/sbin/mkitab 'sqdr-<database_name>:2:once:/etc/rc.sqdr-<database_name> start'`

3. View the contents of the `/etc/inittab` file to ensure the entry was added correctly.

Adding Capture Agent Daemon to HP-UX Startup

If you are running the HP-UX operating system, create symbolic links in the `/sbin/rc2.d` and `/sbin/rc1.d` directories that point to the Capture Agent daemon startup and shutdown script.

1. Log on as `root` user.
2. Verify that a symbolic link exists in `/sbin/init.d/sqdr-<database_name>` and points to `/var/sqdrplus/<database_name>/bin/sqdrplus`. If the link does not exist, create it using the following command:

   `$ ln -s /var/sqdrplus/<database_name>/bin/sqdrplus /sbin/init.d/sqdr-<database_name>`
3. Create the following two additional symbolic links. The link in the `rc2.d` directory starts the Capture Agent daemon with a priority of 99, using the linked script, when entering level 2 of the startup process. The link in the `rc1.d` directory kills the Capture Agent daemon with a priority level of 99, using the linked script, when entering level 1 of the shutdown process.

   `$ ln -s /sbin/init.d/sqdrplus /sbin/rc2.d/S99sqdr-<database_name>`
   `$ ln -s /sbin/init.d/sqdrplus /sbin/rc1.d/K99sqdr-<database_name>`

Adding Capture Agent Daemon to Linux Startup

Use the `chkconfig` command to add the Capture Agent daemon to the Linux boot process.

1. Log on as `root` user.
2. Verify that the `/etc/init.d/sqdr-<database_name>` symbolic link exists and points to `/var/sqdrplus/<database_name>/bin/sqdrplus`.
3. Run the following command to create a link in `/etc/rc2.d`:

   `$ /sbin/chkconfig --add sqdr-<database_name>`

Adding Capture Agent Daemon to Solaris Startup

For the Solaris operating system you create symbolic links in the `/etc/rc2.d` and `/etc/rc1.d` directories that point to the Capture Agent daemon shell script.

1. Log on as `root` user.
2. Verify that a symbolic link named `/etc/init.d/sqdr-<database_name>` exists and points to `/var/sqdrplus/<database_name>/bin/sqdrplus`. If the link does not exist, create it using the following command:

   `$ ln -s /var/sqdrplus/<database_name>/bin/sqdrplus /etc/init.d/sqdr-<database_name>`

3. Create the following two additional symbolic links. The link in the `rc2.d` directory starts the Capture Agent daemon with a priority of 99, using the linked script, when entering level 2 of the startup process. The link in the `rc1.d` directory kills the Capture Agent daemon with a priority level of 99, using the linked script, when entering level 1 of the shutdown process.

```
$ ln -s /sbin/init.d/sqdrplus /sbin/rc2.d/S99sqdr-<database_name>
$ ln -s /sbin/init.d/sqdrplus /sbin/rc1.d/K99sqdr-<database_name>
```

Automatically Starting the DB2 Instance

The Capture Agent service requires that the instance of DB2 be running for it to start.

On a Windows computer services have a Startup Type property, which for DB2 UDB and the SQDR Plus Capture Agent, typically is set to Automatic so the services are started when the system is started. In addition, the SQDR Plus service, which appears as SQDRPlus<*database_name*> in the Services management console, is specified with a dependency on the database instance to help ensure they will be started and stopped at the same time.

On a UNIX-based system you can configure the DB2 Fault Monitor Coordinator (FMC) to ensure that the DB2 instance is running when you start the Capture Agent service. The FMC monitors an instance of DB2 and restarts the instance if it shuts down in any way other than with the **`db2stop`** command. Use the **`db2fm`** command to modify the fault monitor registry file, referring to your DB2 documentation if you need more information about configuring and starting the FMC.

You can determine whether a database instance is set to start automatically by running the command **`db2set`** as the instance owner. If the instance is set to start automatically the value for DB2AUTOSTART is set to YES. You can set the instance to start automatically by running the following command as the instance owner.

```
db2iauto -on <database_instance>
```

Reviewing Capture Agent Logs

If the logging property is enabled in the `sqagent.properties` configuration file (See “Configuring the Capture Agent” on page 37), the SQDR Plus service program logs operational information to the location specified for the Database Configuration Directory parameter during installation (see “Database Configuration Directory” on page 25).

Use a text editor to review the information in the `capture_agent<n>.log` files. The most recent logging information is always contained in capture_agent0.log When the capture_agent0.log file reaches the maximum size allowed it is renamed to capture_agent1.log (rippling the filename change for all the older log files by one) and a new capture_agent0.log file is started.

Using the Capture Agent Maintenance Utility

The SQDR Plus software includes a Capture Agent Maintenance Utility to help with management tasks. You can use the Maintenance utility to display and delete orphan subscriptions that can result if an SQDR client stops using SQDR Plus without cleaning up, or if the SQDR service abandons a local control database without first deleting the incremental subscriptions. You also

can use the Maintenance utility to start and stop the Capture Agent, to control which database tables are published for replication, and to send SQDR Plus log files to StarQuest Customer Support if you need help troubleshooting a problem.

Starting the Capture Agent Maintenance Utility

To start the Capture Agent Maintenance Utility from a Windows computer, select Start–>Programs–>SQDR Plus for DB2 UDB–>Capture Agent Maintenance.

From a Linux or UNIX computer you can run the Capture Agent Maintenance utility from a command line as described in the following steps.

1. Change to the `/var/sqdrplus/<database_name>/bin` directory.
2. Enter the following command to run the Capture Agent Maintenance utility.

   ```
   # ./camaint
   ```

3. Enter a valid DB2 user ID and password. The default user ID is the DB2 User specified during installation of SQDR Plus (see "DB2 User ID and Password" on page 24).

If there is more than one Capture Agent installed, a menu appears so you can enter the number that corresponds to the agent you want to manage.

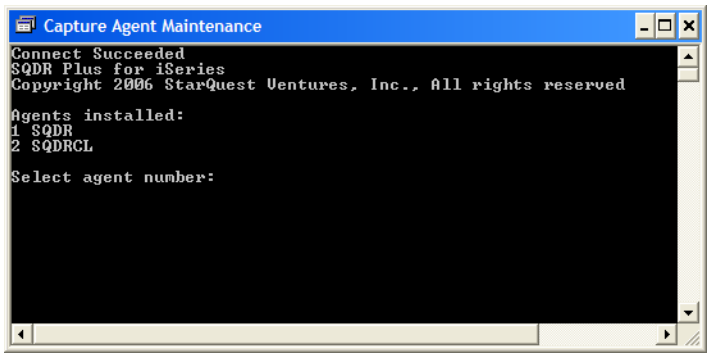

The Maintenance program starts, displaying a Main Menu from which you can select a task to perform.

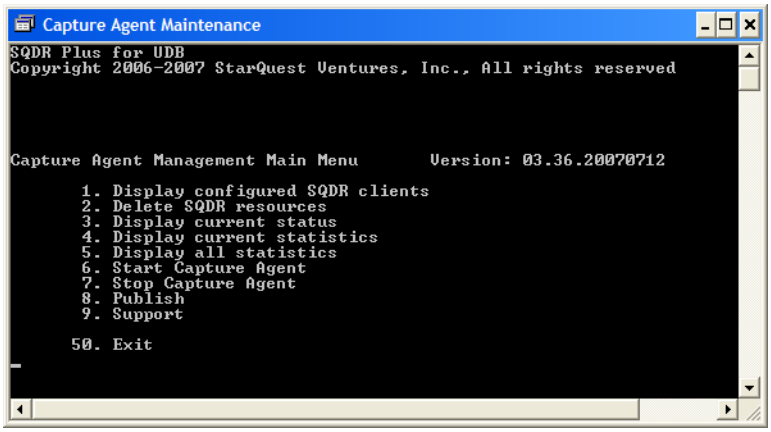

Displaying Configured SQDR Clients and Transaction Status

Select option 1 from the Capture Agent Management Main Menu to view the SQDR client computers that are configured to use the SQDR Plus Capture Agent. The utility displays a list similar to the following, which shows the computer name of the client, and the name and ID of the SQDR control database the client is using.

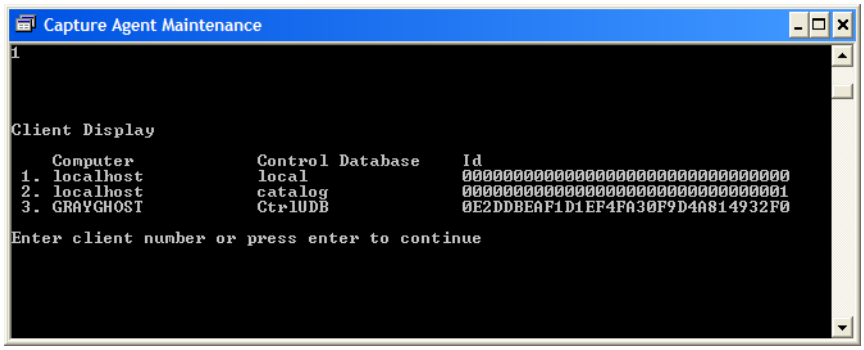

Enter the number that appears next to the SQDR client computer to display the list of incremental groups that Capture Agent is monitoring for the client. The information in the Transactions and Oldest Transaction columns can help you determine whether the client is retrieving change data from the Capture Agent. Note that, in the following illustration the status indicates that there are 2 transactions pending for the MAX_PING_PONG group, with a hexidecimal string that shows the oldest transaction.

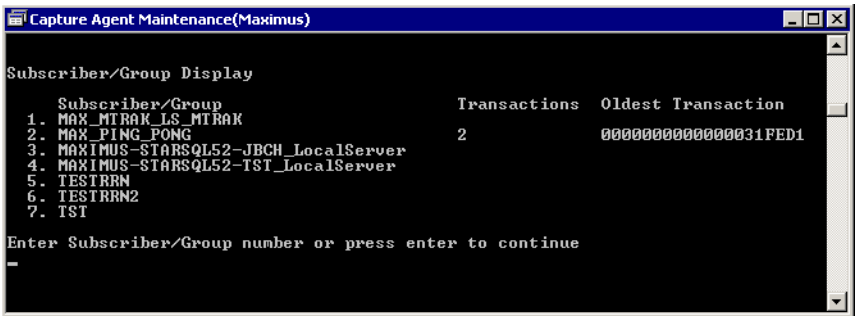

If you review the status again after a period of time and the value for the Oldest Transaction is different, as shown in the following illustration, the SQDR client is retrieving changes from the Capture Agent. If the value for Oldest Transaction does not change, check that the SQDR service

on the client computer is running and that the incremental group has not been paused. If the Capture Agent has changes pending for subscriptions that have been orphaned, you may need to delete the resources as described in “Deleting SQDR Resources” on page 53.

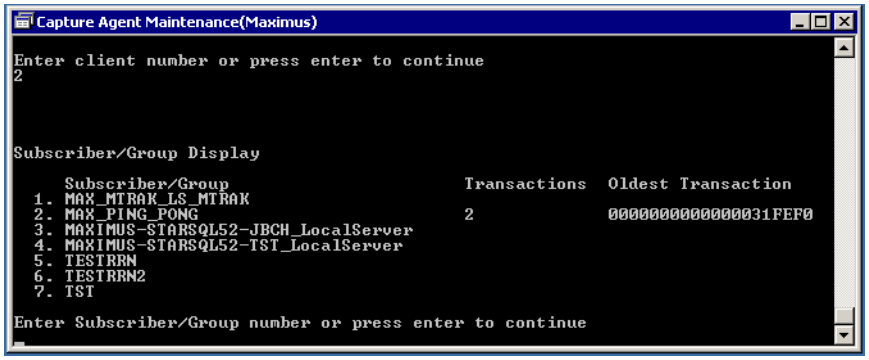

Enter the number that appears next to a particular group to list the specific tables that are subscribed to and how many transactions are pending for each.

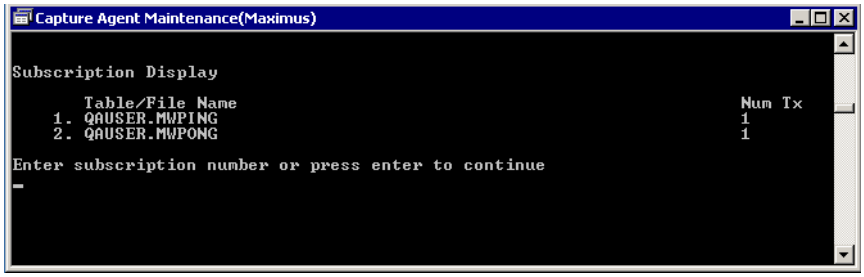

Deleting SQDR Resources

The Capture Agent automatically prunes the change data control tables and database archive logs that are no longer needed. However, incremental replication subscriptions can become “orphaned” if the subscriptions are not cleaned up correctly from the SQDR Replicator Manager on the Windows client. This can occur if the SQDR Replicator Manager is re-configured to use a different local control database, ceases to run the defined incremental replication subscriptions, or the client software is uninstalled prior to deleting the defined subscriptions.

Option 2, Delete SQDR Resources, of the Capture Agent Management Main Menu allows you to delete an incremental subscription, a group, or a client computer.

 **WARNING**	Be sure you select the correct client computer, subscription, or group when you are deleting SQDR resources. The resource that you select will be permanently deleted and there is no confirmation for the delete operation.

1. Select option 2, Delete SQDR Resources, from the Capture Agent Management Main Menu.
2. From the Delete SQDR Resource Menu, enter option 1, 2, or 3, depending on whether you need to delete a configured client computer, a group, or a subscription, respectively.

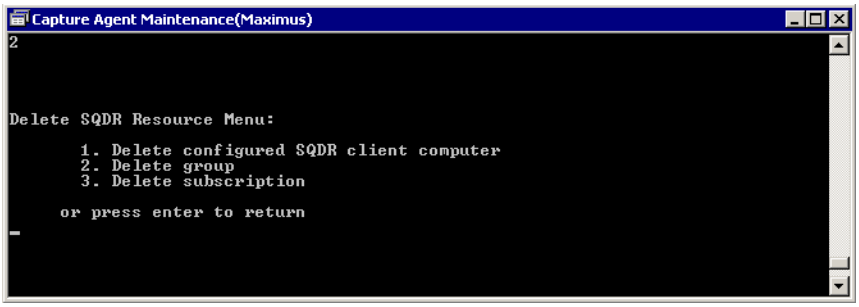

3. From the list of SQDR clients, enter the number that appears next to the client for which you want to delete the resource. Clients are identified by computer name, SQDR control database name, and a unique control database ID.

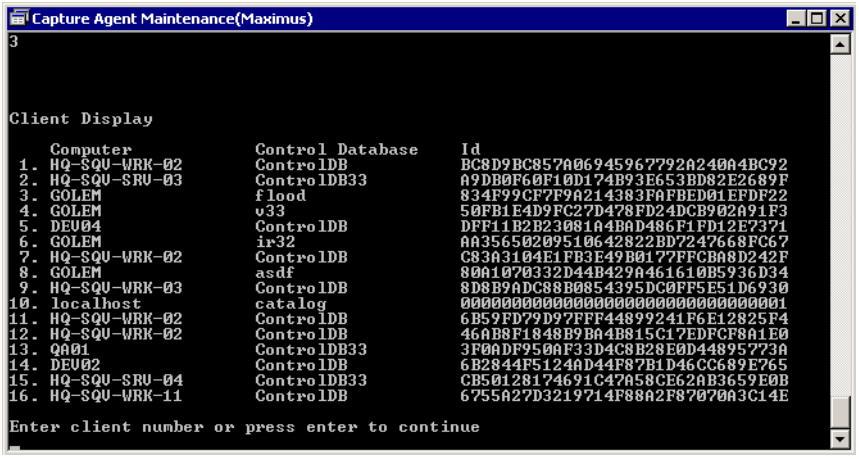

4. If you are deleting a group or a subscription, enter the number that appears next to the appropriate group.

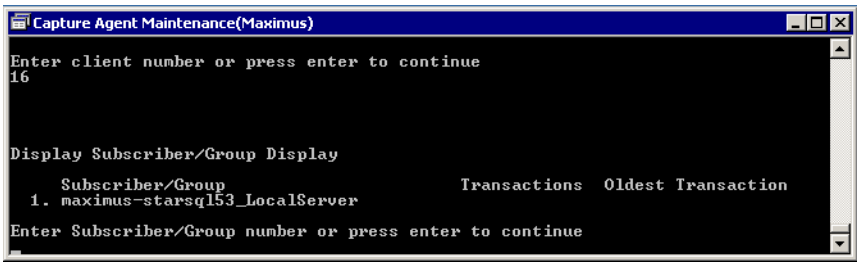

5. If you are deleting a subscription, enter the option number that appears next to the subscription you want to delete.

6. When you are finished deleting resources, press Enter several times to return to the Main Menu.

When you delete resource objects, SQDR Plus releases the system resources that were used by the client, group, or subscription.

Displaying Status Information

Option 3 of the Capture Agent Maintenance Main Menu displays the status of the SQDR Plus components for the selected Capture Agent. For example, the following screen shows that the Capture Agent and several Log Readers are running within the selected schema.

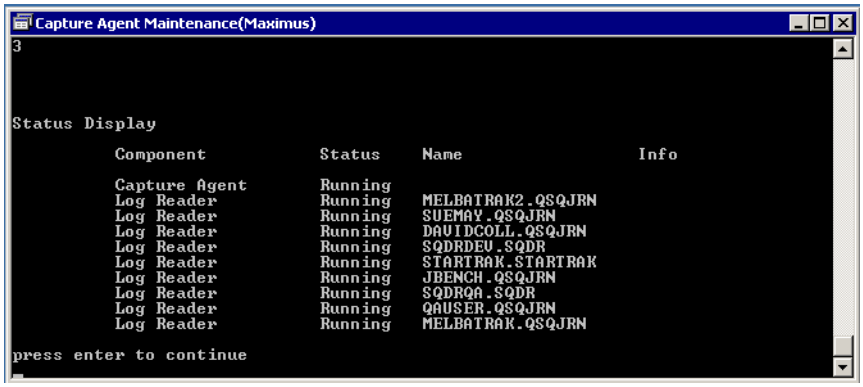

Displaying Statistics

Options 4 and 5 of the Capture Agent Maintenance Main Menu display current or all statistics, respectively. The current statistics reflect one hour of activity for the date and time shown. The time is reported according to 24-hour notation, such as 11 for statistics gathered between 1100 (11:00 a.m.) and 1200 (12:00 p.m.).

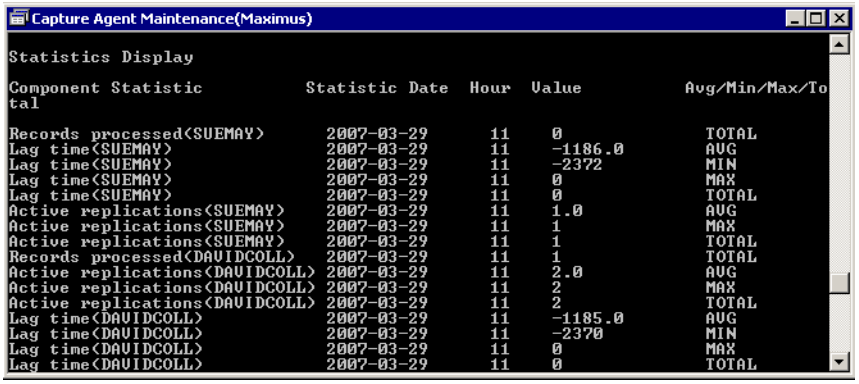

For active replications, the statistics show the total number of records processed, the average, minimum, maximum, and total lag time, and the average, maximum, and total number of active replications.

The minimum, maximum, average, and total lag times indicate the number of elapsed seconds between reading the database log and applying the change data. A total lag time of 0 indicates that, during the hour for which the statistics were sampled, the Log Reader was maintaining pace with the database changes, reading the log and recording the change data with no delays. The statistics also show the average number of seconds and the maximum number of seconds in which the Log Reader lagged behind in recording the change data.

You also can obtain the statistics by running a query on the `sq_statistics` table using another tool, such as Microsoft Excel. For example, you can issue a query such as `SELECT * from <schema>.sq_statistics`, and then save or print the results.

Starting and Stopping the Capture Agent

The Capture Agent Maintenance utility displays options for starting and stopping the Capture Agent; however these options (provided for compatibility with CAMAINT in an IBM iSeries environment) do not function properly in a DB2 UDB LUW environment.

If you are running SQDR Plus on a Windows computer, use the Windows Services utility to stop and start the Capture Agent. The Services utility typically is located in the Administrative Tools of the Control Panel. Start the Services program and review the services list to find the Capture Agent service that is associated with a particular database. The Capture Agent service appears in the list with the name of the database it is monitoring, in the format of SQDRPlus(*database_name*). As shown in the following screen, right-click the service to display a context menu that allows you to stop and start the service.

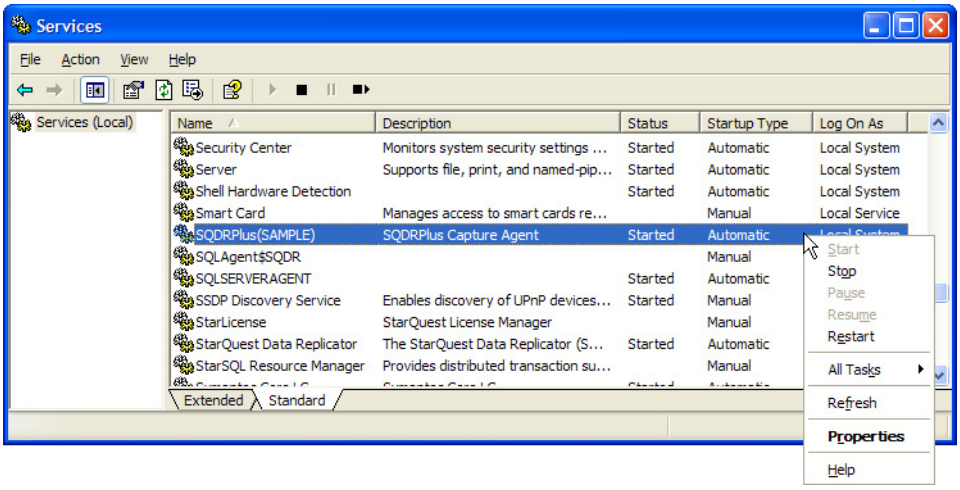

UNIX To start or stop the Capture Agent, or view the status, on a Linux or UNIX computer, execute the `sqdrplus` script with the Start, Stop, or Status option as shown below.

```
cd /var/sqdrplus/<database_name>/bin
./sqdrplus start | stop | status
```

If you run the `sqdrplus` script as a user other than `root` or the Run As User, you are prompted to enter the password for the Run As User to start or stop the Capture Agent.

Publishing Database Tables

Option 8, Publish, on the main Capture Agent Maintenance menu allows you to specify which database tables you want to make available to SQDR clients for replication. The Publish feature provides centralized control that a database administrator can use to restrict specific database tables from being replicated to other destinations. When the `requirePublication` configuration parameter is set to `true`, only the database tables that have been published from the Capture Agent can be successfully subscribed to for replication. Note that publishing and unpublishing tables affects only new subscriptions that are created from the SQDR Replicator Manager.

The Manage Publications Menu that appears after you select option 8 includes functions for displaying the published tables, and for selecting which tables to publish or unpublish.

Displaying Published Tables

Select option 1 to display a list of tables that have been published. The tables are listed in the format *schema.tablename*, with a timestamp and a state of Initial, Subscribed, Published, or Failed. *Initial* indicates that the request to publish the table has been initiated. *Subscribed* means that the

Capture Agent is setting up the change data table, and *Published* indicates that changes are being tracked and the source table can be subscribed to. If the table cannot be published the state indicates *Failed* with additional text to explain the problem.

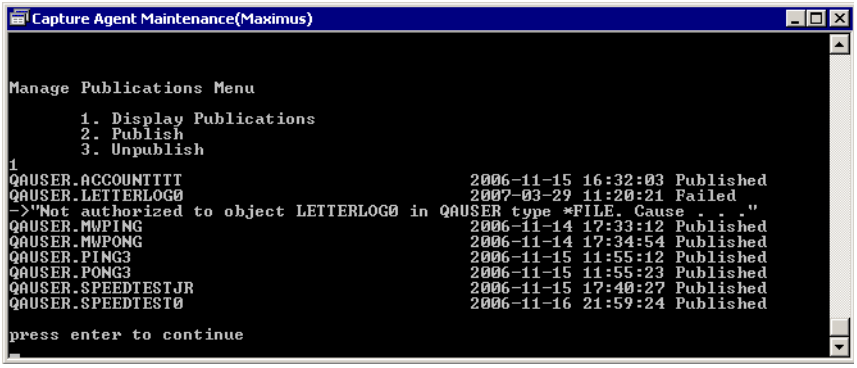

Publishing Tables

Select option 2, Publish, from the Manage Publications Menu to specify which database tables you want to make available for SQDR users to subscribe to. After you enter the option number, specify the schema and name of the table you want to publish. You can use wildcards to specify a pattern for the schema and the name to publish multiple tables at once, as shown below to publish all the tables in the QAUSER schema that begin with the letter P.

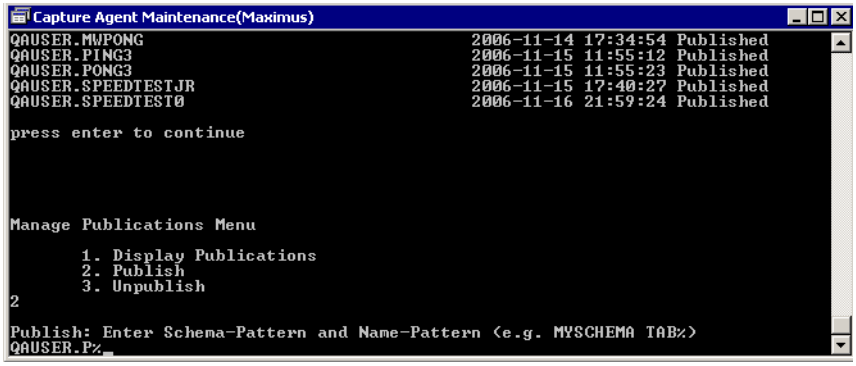

After the Capture Agent creates a staging table for the published table(s), new subscriptions for replicating the source table data can be created from the SQDR Replicator Manager. Note that the Capture Agent publishing functions affect only new, not existing, SQDR incremental replication subscriptions.

Prior to SQDR+ v3.60, the Publish function required that an exclusive lock be acquired on a table before being considered published; this is no longer necessary. If an attempt to lock the table for shared update fails, the subscription will be published with a "*" character next to the state of *Published* to indicate that a lock was not obtained prior to starting to stage changed data.

If the original more limiting behavior is desired, it can be configured by setting the keyword **retryPublishLock=true** in the sqagent.properties file; in this case the publication will be displayed as Subscribed, and the lock operation will be retried every 10 minutes. Once successfully locked (and immediately released), the state will be updated to *Published*.

Unpublishing Tables

After a table has been unpublished it cannot be subscribed to from the SQDR Replicator Manager. However, removing a table from the published list has no affect on existing subscriptions to the table.

If you need to remove one or more database tables from the published list, select option 3, Unpublish, from the Manage Publications Menu. After you enter the option number, specify the schema and name of the table you want to unpublish. You can use wildcards to specify a pattern for the schema and name to remove multiple tables at once from the published list.

Recovery Menu

Option 10 of the Capture Agent Maintenance menu brings up the Recovery menu, which allows the SQDR Plus Administrator to change the startup mode for the Capture Agent and take an individual incremental replication subscriptions offline or online.

The Recovery menu is shown below.

Agent Recovery Menu
1. Set Startup Mode (Current: "Normal")
2. Take Subscription Online
3. Take Subscription Offline
4. Quiesce Capture Agent

The Startup Mode affects the behavior of the Capture Agent, the next time it is started. The Online/Offline functions take effect immediately. The following sections detail each feature in detail.

Set Startup Mode

The default Startup Mode for the Capture Agent is Normal. Before changing the Startup Mode, the Capture Agent must be stopped.

When to change the Startup Mode

The Startup Modes are useful in situations where the source database is moved from a production system to a backup system ("fail-over") or where SQDR operations are restored on a primary server ("fall-back"). Both operations can be performed without the need for SQDR to run baseline replications. This option is termed a “Warm” startup mode. Another new startup option, “Cold” startup, will force all subscriptions to be flagged as requiring new baselines. This “Cold Start” option is useful when the host control tables are intact and in agreement with the SQDR client(s), but the User source and destination database tables are no longer be assumed to be synchronized. For example, this situation might occur in the event of unrecoverable log damage, necessitating a restore. Once the Capture Agent has completed the startup-time recovery, it resets the startup mode for the next invocation to “Normal”. The Set Startup Mode has three options:

1. Startup with fully recovered control and staging tables.

2. Force all subscriptions offline ("online" needed after restart).
3. Normal Startup.

Normal Startup

SQDR Plus assumes that the state of its log, staging tables and log positions are still current and accurate. However, at startup, the Capture Agent checks for source table structure changes. The Startup Mode should be set to Normal except when a "fail-over" or "fall-back" operation is being performed.

“Warm” Startup with fully synchronized control and staging tables

SQDR Plus assumes that logging tables and staging tables are in a working state and that it only needs to recover the current log position.

This option is useful when copying the entire SQDR Plus schema (e.g., library) to a new system, or other types of complete restore operations where the consistency of the database is guaranteed. SQDR Plus skips any journal entries showing that tables having been deleted and re-created by moving its log/journal read position to "now".

As in a Normal startup, the Capture Agent validates source table structures. If auto-journaling is enabled and SQDR Plus discovers that a source table is no longer journaled, it will automatically try to restore journaling. Also, grants to the views on the staging tables are made to match the source tables.

“Cold” Startup

SQDR Plus makes no assumption about the consistency of log and staging tables. It skips any log/journal changes and takes all subscriptions "offline". The user can then manually move subscriptions "online" in a controlled fashion once SQDR Plus is restarted. Taking a subscription online will cause a new baseline replication to be required. (See “Taking Subscriptions Offline/Online” on page 60 for more details).

As with the *Startup with fully recovered control and staging tables* mode, metadata verification occurs, journaling is restored, if necessary, and grants to the views on the staging tables are made to match the source tables.

In the latter two cases, SQDR Plus assumes that the source and destination tables are in sync at the time of the "fail-over" or "fall-back" operation, and that only its log position is to be considered invalidated.

Taking Subscriptions Offline/Online

The ability to take a subscription offline/online should be used when significant changes are made on the source tables. In this case, moving subscriptions to offline will suspend staging of changes and will require a new baseline in order to make new changes available. This option is particular useful if a set of tables is restored or loaded with a large number of new rows.

This functionality allows significant changes to be made to a source file or a set of source files (e.g., dropping/re-creating source tables or adding/removing members from tables) without requiring the user to re-create the subscriptions in SQDR.

Taking a subscription offline will flag a given source table as being unavailable for normal Capture Agent processing (e.g., staging changes), until further notice. A subscription in the offline state appears to the SQDR client as if there is no activity on the subscribed table. This allows a database administrator to restore the table to a desired version without the Capture Agent processing and staging all records, which may cause significant delay on large tables.

Once the restore operation is complete, the administrator uses the *Take Subscription Online* option to notify the Capture Agent to resume processing for the subscribed table. At this point the Capture Agent will compare its saved metadata for the table to its current state. If the structure of the subscribed table is unchanged, the SQDR client is notified that only a new baseline snapshot is required. If there have been structural changes or if the subscribed table no longer exists, a "table altered or deleted" notification is generated and the subscription must be deleted and re-created. If the subscription is already online when you issue the online request, no changes occur.

If the table is no longer journaled after the restore, SQDR Plus will attempt to restore journaling for the table. If the attempt to start journaling fails, the subscribed table will remain "offline".

The Online and Offline operations can be performed on multiple libraries (or schemas) and tables by using iSeries CL wildcards (when using the VRYSUB command) or SQL search patterns (when using CAMAINT or invoking Java directly)

The Online/Offline commands are available in the following ways:

Use the Agent Recovery Menu of CAMAINT

```
Agent Recovery Menu
    1. Set Startup Mode (Current: "Normal")
    2. Take Subscription Online
    3. Take Subscription Offline
Enter selection or press enter to continue
```

Start AgentUtil with the –online or –offline argument

```
java –jar AgentUtil.jar [-online | -offline ] <agentSchema> <schemaPattern> <namePattern>
```

"quiesce" Shutdown Option

When the Quiesce Capture Agent option of the Recovery Menu is selected, CAMAINT makes a Quiesce call to the Capture Agent. Upon receipt of the Quiesce request, the Capture Agent terminates the pruning, publishing, and client notification threads. It then notifies each replication worker (one per journal/log being processed) to quiesce its operation.

The replication worker will process any outstanding log/journal records until a call to the log/journaling API returns no records. The replication worker will then enter a quiesced state, where it will remain until the capture agent is stopped.

To monitor the state of the Capture Agent's replication workers, choose option 3 ("Display current status") from the CAMAINT menu. Once all replication workers have entered quiesced state, you can be assured that all log/journal entries have been processed and placed in the capture agent's staging tables, assuming that no new activity has been allowed on the subscribed tables after the replication worker entered quiesced state.

If a replication worker experiences a resource failure while processing in this mode, the Capture Agent will restart the worker and, barring other failures, will process any remaining log/journal entries, eventually entering quiesced state.

Sending SQDR Plus Log Files to StarQuest

If you need to provide the SQDR Plus log files to StarQuest Customer Support, select option 9, "Support," from the Capture Agent Management Main Menu to display the Support Menu.

Support Menu

1. Send SQDRPlus Logs (ftp)
2. Change ftp server (ftp.starquest.com)
3. Change ftp userid (sqdrdrop)
4. Change ftp password
5. Save SQDRPlus Logs to local file system (as zip file)
6. Set directory for saving logs locally (.)

Make a selection or hit RETURN to continue.

Select option 1 to transmit the SQDR Plus log files to the StarQuest FTP server using default values for the FTP server (ftp.starquest.com), user ID (sqdrdrop), and password. The Capture Agent Maintenance utility copies the SQDR Plus log and configuration files to three compressed files and sends those to the StarQuest FTP server, as illustrated in the following screen.

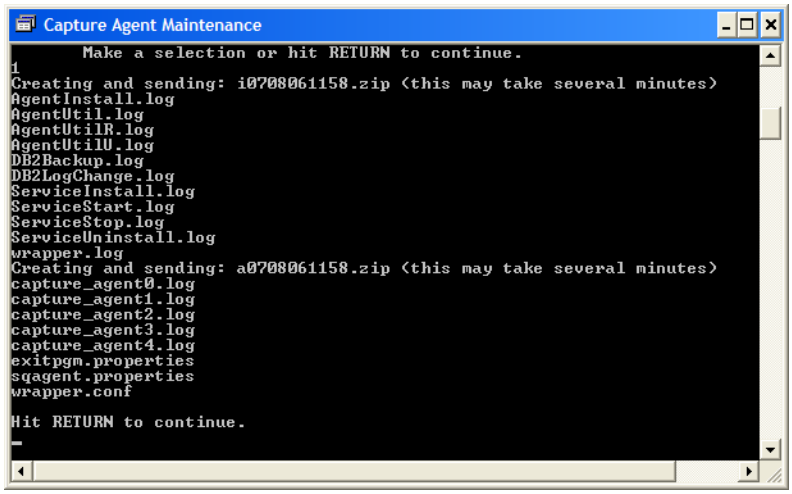

If you cannot connect directly to the StarQuest FTP server, execute option 2 from the Support Menu to specify another FTP server that you have access to. Option 3 allows you to change the user ID from the default value of sqdrdrop to another user ID, and option 4 allows you to change the password that is used for the FTP connection.

If you must send the log files to a different FTP server, make a note of the .zip filenames and send the compressed files to StarQuest (contact@starquest.com) as an email attachment.

If firewall policies or other restrictions prevent access to any ftp server, use option 5 to save SQDR Plus logs to the local file system so they may be forwarded using another method. You may also specify the target directory for the zip files with option 6. The files are saved to the DB2 UDB server's file system (when running SQDR Plus/UDB's CAMAINT), or the local workstation's file system (when running SQDR Remote Admin's CAMAINT).

Maintaining Log Files

If SQDR Plus is managing the database log files (managedLogs is set to True), a user exit program is installed to manage the archiving and retrieval of log files. When the Capture Agent is restarted it must be able to locate the last position in the log files in order to resume applying incremental changes.

Be careful not to manually delete logs that the Capture Agent needs when it is started again. If you override the protection of the exit program and delete log files from the chain you will need to re-run a baseline snapshot before the subscription can resume tracking and replicating incremental changes.

If SQDR Plus is not configured to manage the log files it needs it relies on the DB2 UDB database manager to track the log files. Since the database manager is archiving and retrieving database log files that would be needed for a roll forward recovery, it is unlikely that the log files would be deleted before the Capture Agent has applied the logged transactions. As long as the Capture Agent is running whenever DB2 is running, the Log Reader position should be close to the end of the DB2 log where the most recent log records are located. This minimizes the number of older log files that Capture Agent needs.

Determining the Log In Use by Capture Agent

The numbering scheme for archived logs begins with S0000000.LOG and continues through S9999999.LOG. To determine which log file Capture Agent currently is using you can find the log position and then use the DB2 Find Log Sequence Number command (**`db2flsn`**), as described in the following steps.

1. Issue the following SQL statement against the SQ_READERS table of the desired database.

   `SELECT logpos FROM <`*`schema`*`>.SQ_READERS`

The SELECT command returns the log position in a hexadecimal format such as 0x000000026afbdbc3. You provide this log position when issuing the **`db2flsn`** command to get the filename that contains the Log Sequence Number (LSN). The input LSN can be a 12 character string, so you need to trim the leading 0x0000 from the hexadecimal format. For example, if the value returned by the SELECT logpos statement is 0x000000031697B887, you would enter the log position variable as 00031697B887.

2. Open a DB2 command window. If you are using DB2 for UDB v8, you must change to the database directory that contains the SQLOGCTL.LFH file, or copy the SQLOGCTL.LFH file to the current directory, to execute the **`db2flsn`** command. The SQLOGCTL.LFH file is created when the database is created, and is typically found in a directory path such as C:\DB2\NODE0000\SQL0000*n* on Windows or ~<*instance_owner*>/<*instance_owner*>/NODE0000/SQL0000*n* on a Linux or UNIX computer.

3. Enter the following command, substituting the LSN that was returned from the SELECT statement, *without* the leading 0x0000, for the <*log_position*> variable.

   For DB2 for UDB v8:

   db2flsn <*log_position*>

   For DB2 for UDB v9:

   db2flsn -db <*database_name*> <*log_position*>

The **db2flsn** command returns the name of the log file that contains the LSN, which indicates which log file the Capture Agent is currently using, in a message such as "Given LSN is contained in log page 2 in log file S0000002.LOG." You must retain this log file and more recent log files to ensure that the Capture Agent can apply the transactions that are recorded in the log files.

Manually Pruning Log Files

If you do need to explicitly prune log files that are older than the file in use, use the **db2 prune** command, specifying the files that are older than the log file that the Capture Agent is currently using. Following is an example of the command to prune log files that are older than S000354.LOG.

```
db2 prune logfile prior to S000354.LOG
```

The log file that you specify in the prune command is not deleted.

Troubleshooting Operational Problems

Following are some general steps for troubleshooting operational problems:

1. Use the Capture Agent Maintenance Utility to view the status of the Capture Agent and its Log Reader. (See "Displaying Status Information" on page 55.)
2. If the status shows the Capture Agent is stopped or failed when it should be active, review the operational logs for SEVERE conditions in the capture_agent<*N*>.log files.
3. Use the Capture Agent Maintenance Utility to review the statistics to determine whether the Records Processed counter is updating at least every few minutes. (See "Displaying Statistics" on page 56.)
4. Use the Windows Task Manager to determine if any of the SQDR Java processes are using excessive CPU resources.

   From a Linux or UNIX computer use the status option to obtain the process identifier (PID) of the Capture Agent daemon and then filter the information to show the process statistics, as shown in the following example:

   ```
   cd /var/sqdrplus/<database_name>
   ./sqdr-<database_name> status
   ```
   (for AIX, **/etc/rc.sqdr-<*database_name*> status**)
   ```
   ps -eaf | grep <PID>
   ```

If you need to increase the logging level to capture more details, edit the logging level specified in the `sqagent.properties` file (see page 37) and restart the Capture Agent service. Log files have names such as capture_agent0.log. They can be viewed using Notepad (Windows), vi (UNIX), or any other application that can read text files.

StarQuest Technical Support engineers may need you to send the Capture Agent logs to help troubleshoot a problem. The Capture Agent Maintenance utility provides a Support option to make this task easier, as described in "Sending SQDR Plus Log Files to StarQuest" on page 63.

Using the DB2 Diagnostic Logs to Monitor SQDR Plus Operations

The DB2 diagnostic logs also can provide insight to operational problems. The db2diag.log typically is located in the home directory of the database instance owner (C:\Program Files\IBM\SQLLIB\DB2 on a Windows computer and ~<*instance_owner*>/sqllib/db2dump/ on a Linux or UNIX computer).

If the Capture Agent exit program is functioning normally, you should find a message in the db2diag.log similar to the following:

```
FUNCTION: DB2 UDB, data protection, sqlpgArchiveLogFile, probe:3180
MESSAGE : ADM1844I
Successfully archived log file "S0000011.LOG" to "USEREXIT" from
"/home/db2inst2/db2inst2/NODE0000/SQL00002/SQLOGDIR/".
```

If the Capture Agent is not running, the exit program may stop working, in which case a message similar to the following would appear in the db2diag.log:

```
FUNCTION: DB2 UDB, data protection, sqlpgArchiveLogFile, probe:3160
MESSAGE : Failed to archive log file S0000012.LOG to USEREXIT from
"/home/db2inst2/db2inst2/NODE0000/SQL00002/SQLOGDIR/ with rc = 32.
```

UNIX

If the DB2 instance owner has permission to write to the /var/sqdrplus directory (such as by issuing the command `chmod 777 /var/sqdrplus`) and SQDR Plus is managing logging, the SQDR Plus exit program will create additional files named ARCHIVE.LOG and USEREXIT.ERR that can provide insight to the SQDR Plus operations. Review the contents of the active logging directory (/home/db2inst2/db2inst2/NODE0000/SQL00002/SQLOGDIR/) and the exit program cache directory (/var/sqdrplus/<*database_name*>/cache) to help ensure that incremental changes are flowing properly.

Monitoring Backlogs with the Storage Monitor

SQDR Plus provides a mechanism for monitoring backlogs and will send a warning notice and suspend log readers if necessary to avoid excessive usage of the table space where the staging tables reside. The default mechanism uses a simple row-count based throttling mechanism, but is user-extensible.

By default, the Storage Monitor uses a simple SQL statement to count the number of rows in complete transactions available in the staging tables. You can define alternate criteria with the **smSQL** property in sqagent.properties; this should be specified as an SQL statement which returns an integer number.

When the number of rows exceed a certain value (defined as **smWarningLevel**), then a warning notification is issued.

When the number of rows exceeds the **smThrottleLevel** value, another warning notification is issued and all log readers (replication workers) are paused. The log readers resume when the number of rows drop below the **smThrottleLevel** value. When the "storage monitor" has throttled staging, pruning runs at a higher frequency than usual (using the "storage monitor" interval).

The functionality is controlled by the following sqagent.properies settings:

SmSQL - a SQL statement used to monitor staging activity; it should return a single integer or bigint value. Note that the default SQL statement (used when **smSQL** is not specified) is valid only when SQDR Plus is monitoring only one journal.

Examples:

```
smSQL=SELECT SUM(ROWCOUNT) FROM SQDR.SQ_TRANSACTIONS WHERE TXID IS NOT NULL
AND TXID <= (SELECT MIN(MAXTX) FROM SQDR.SQ_READERS)
```

smInterval (default 0 (off)) – Frequency in minutes of when to run the storage monitor. 0 means "off".

smWarningLevel (default 450000) - Any integer value above this value will cause a notification/log message.

smThrottleLevel (default 500000) - Any integer value above this value will pause all log readers (replication workers) and cause a notification/log message.

Note that in some cases this feature may have undesirable side effects. For example, if staging is suspended because of the "storage monitor" condition, the storage used by the archived transaction logs will grow instead; this may be less desirable, especially if SQDR Plus is staging only a small portion of the data in the receivers. You should analyze your system and compare the amount of available storage for the container holding the SQDR storage tables with the amount of available store for DB2 archived logs to determine whether or not this feature is suitable for your environment.

Performance Suggestions

On a DB2 UDB system that has limited disk space and a high rate of database transactions, you may want to decrease the interval in which SQDR Plus prunes change data from the control tables to free disk space more frequently. See “pruneInterval” on page 42 for details about changing the pruning interval.

Creating an Identity Column

If a source table does not have unique indexes defined SQDR Plus can support only the INSERT* operations when replicating data to the target database. If any DELETE* or UPDATE* operation is committed to the source database, a new baseline snapshot replication is required.

To avoid having to perform a new baseline replication frequently you can alter the source table to include an identity column. An identity column provides a way for DB2 to automatically generate a unique numeric value for each row in a table, providing unique primary key values. Each table can have a single column that is defined with the identity attribute.

It is difficult to add an Identity column to existing DB2 UDB tables, so the following procedure describes how to create a new table with the same layout as the original table, add an identity column, and then populate the new table with the data from the original table. If you are using the StarSQL driver to access the datasource, you must disable the Strict Parsing option in the DSN configuration for the following CREATE statement to succeed.

1. Create a new table that has the same layout as the original table, and add an identity column, as shown in the following example.

   ```
   create table NEW.TABLE (c1 varchar(22), ident integer not null
   generated always as identity (start with 0 increment by 1 no
   cache))
   ```

2. Populate the new table with the data from the original table, such as:

   ```
   INSERT INTO NEW.TABLE (C1) SELECT C1 FROM OLD.TABLE
   ```

3. Drop or rename the original table.

   ```
   RENAME TABLE OLD.TABLE TO NEW.NAME
   ```

4. Now you can give the new table with the identity column the same name as the original table.

   ```
   RENAME NEW.TABLE TO OLD.TABLE
   ```

5. Create a unique index or primary key for the identity column.

Applications that follow standard practices in the SQL they produce generally will not have a problem operating on a table that includes an identity column. For example, INSERT statements that specify a “column-list” in addition to the “values-clause” will execute successfully. Be sure that any application that operates on the table can work with the extra identity column.

Support for partitioned tables

SQDR Plus for UDB supports incremental replications for partitioned tables in DB2/UDB LUW 9.1 and later.

Adding empty partitions and subsequently populating them requires no special considerations. SQDR Plus will detect newly-added partitions, and automatically starts processing transaction log entries for the new partition.

Considerations for ATTACH/DETACH partition:

SQDR Plus will detect when partitions are DETACHED, and when partitions are ATTACHED (and INTEGRITY is set for the table). These events will cause SQDR Plus to notify the SQDR client that a new baseline is required. If the SQDR client (v3.62 and later) is configured to “ignore deletes”, the DETACH notification is ignored. There may be other considerations associated with “ignore deletes” that should be considered when using this option.

Removing the SQDR Plus Software

If you need to remove the SQDR Plus software and will not be re-installing it on the same computer, follow the procedures in this section.

There are some tasks that you need to perform prior to running the uninstall program to cleanly remove the software from the host computer.

1. From the SQDR Manager that is running on a Windows computer, delete any subscriptions that perform incremental replication operations from the DB2 UDB database before you remove the SQDR Plus software from the DB2 UDB computer.
2. Review the contents of the Database Configuration directory on the host computer (see page 25 for the default directory location) and copy any log files or other files that you want to retain to a different directory. The Database Configuration directory will be cleared and deleted when you uninstall SQDR Plus.

Uninstalling from a Windows Computer

After you complete the preparation steps in the previous section, use the Windows **Add or Remove Programs** (or **Programs and Features**) utility to completely remove the SQDR Plus software and environment.

1. Be sure you are logged on as a user with Administrator rights.
2. Select **Add or Remove Programs** from the Windows Control Panel. On Vista and later, select **Programs and Features**.
3. In the Add or Remove Programs window select SQDR Plus for DB2 UDB and click the Change/Remove button.
4. Click Next when the InstallShield Wizard displays the Welcome pane that informs you the SQDR Plus for DB2 UDB software will be uninstalled.
5. Click Next to indicate that you want to uninstall the SQDR Plus Java programs.
6. From the summary pane click Uninstall to begin the process of removing the software.
7. If the SQDR Plus service is running, click Next to stop it.
8. Click Next after the service has been stopped and removed.
9. On the next pane enter the password for the DB2 user. If you plan to re-install SQDR Plus for UDB on this computer, de-select the option Drop Control Database. Otherwise leave this option enabled to remove the SQDR Plus control tables. Click Next to continue.
10. Continue clicking Next and responding to the prompts as the Wizard removes the stored procedure registration and performs other tasks to uninstall the software.
11. Click Finish when the Wizard has successfully uninstalled SQDR Plus for DB2 UDB.

If the SQDR Plus installer changed the logging type of the database from circular to Exit Program (USEREXIT), change it to either archive or circular, or install another exit program, after removing SQDR Plus.

Uninstalling from a Linux or UNIX Computer (GUI installer)

Follow the steps below to remove the SQDR Plus software from a UNIX-based computer. The computer must have an X-Windows environment, or you can run the uninstall script from a console, use a VNC connection from a remote computer, or set the DISPLAY environment variable to point to another UNIX computer that has the X-Windows environment.

1. Log on to the computer as a privileged user.
2. Change to the directory where the SQDR Plus program files are installed and run the **`uninstall`** script, such as:

   ```
   cd /opt/StarQuest/sqdrplus
   ./uninstall
   ```

3. Click Next when the InstallShield Wizard displays the Welcome pane that informs you the SQDR Plus for DB2 UDB software will be uninstalled.
4. Click Next to indicate that you want to uninstall the SQDR Plus Java programs.
5. From the summary pane click Uninstall to begin the process of removing the software.
6. If the SQDR Plus service is running, click Next to stop it.
7. Click Next after the service has been stopped and removed.
8. On the next pane enter the password for the DB2 user. If you plan to re-install SQDR Plus for UDB on this computer, de-select the option Drop Control Database. Otherwise leave this option enabled to remove the SQDR Plus control tables. Click Next to continue.
9. Continue clicking Next and responding to the prompts as the Wizard removes the stored procedure registration and performs other tasks to uninstall the software.
10. Click Finish when the Wizard has successfully uninstalled SQDR Plus for DB2 UDB.

Symbolic links that were created when the SQDR Plus software was installed are deleted when the software is uninstalled.

If the SQDR Plus installer changed the logging type of the database from circular to Exit Program (USEREXIT), change it to either archive or circular, or install another exit program, after removing SQDR Plus.

Uninstalling from a Linux or UNIX Computer (Text-based installer)

1. Log on to the UNIX/Linux computer as a user with root authority.
2. Set up access to DB2 by sourcing db2profile (sh/ksh) or db2cshrc (csh) in /home/<*db2inst_owner*>/sqllib.
3. cd /opt/StarQuest/sqdrplus (or wherever you installed to)
4. Run the uninstallation:

   ./uninstall [-v]

-v – verbose (optional)

You will be prompted for the password of the DB2 user that was configured in setup.conf.

You will be prompted whether to remove the SQDR Plus control tables. After uninstalling the product, the bin, lib, conf, and logs subdirectories of `/var/sqdrplus/<DATABASE>` are renamed with a prev suffix – e.g. bin is renamed to bin.prev to bin.prev, etc.

You can examine the results of some of the actions of uninstalling by examining the logs in `/var/sqdrplus/<DATABASE>/logs.prev`. In addition, conf.prev/sqagent.properties and logs.prev/install_history.log may be useful as a reference of your previous configuration if you choose to reinstall SQDR Plus.

If you changed the type of database logging from circular to USEREXIT prior to installing SQDR Plus, you may want to change the database logging back to circular if you no longer need to retain archive log files for replication operations.

Appendix A: sqagent.properties sample

The following is the contents of the template sqagent.properties file, from which the installer creates sqagent.properties during the initial installation. Refer to `DBCONFDIR/conf/sqagent.properties` in the installation directory (default `/opt/StarQuest/sqdrplus/` on UNIX; `C:\Program Files\StarQuest\sqdrplus` on Windows) to view the latest version of the template.

Items that are not commented out are properties that are configured by the installer based on user input (or, in the case of the text-based installer, by setup.conf) during the installation; items that are commented (leading #) are advanced properties that can be configured by the user; the default value is displayed here.

```
# SQDR Plus for DB2 UDB
# CAPTURE AGENT PARAMETERS
#
controlDbSchema=SQDR
logLevel=INFO
port=50005
requirePublication=false
# pollSleepTime=5000
# pruneInterval=30
# grantees=
# numControlDbConnectionsPerVm=10
# singleTrans=true
# useTxSequence=false
# clientDeleteGrace=1
# clientMonitoringInterval=1440
# filterUserId=myuser
#
# Space Monitor Parameters
# smSQL=the sql used to monitor staging activity
# returns a single integer or bigint value that reflects usage.
# sminterval=frequency of monitor in minutes. 0 means "off".
# smInterval=0
# smwarning=issue a notification/log message if value returned by smSQL
exceeds this value
# smWarningLevel=450000
# smThrottleLevel pause all log readers (replication workers)
# and issue a notification/log message if value returned by smSQL exceeds
this value
# smThrottleLevel=500000
#
#
# DB2 UDB CONNECTIVITY PARAMETERS
#
```

```
sourceType=UDB
userId=
password=
controlDbDriver=com.ibm.db2.jcc.DB2Driver
controlDbUrl=jdbc:db2:127.0.0.1:446/RDBNAME:driverType=4;deferPrepares=f
alse;
useAuthentication=true
TABLEINFO=false
#
#
# OPERATIONAL NOTIFICATIONS PARAMETERS
#
notificationLevel=OFF
smtpServer=
to=admin@mydomain.com
from=sqdrplus@mydomain.com
starquestNotification=SEVERE
#
#
# CHANGE DATA LOGGING PARAMETERS
#
managedLogs=false
deleteLogsWithoutSave=true
archivedLogPath=
# maxLogFileSize=512000
# maxNumLogFiles=5
#
#
# CHANGE DATA NOTIFICATION PARAMETERS
#
# broadcastAddress=192.1.1.255
# udpPort=2728
# minimumNotificationDelay=5000
```

Appendix B: Format of the SQ_STATISTICS control table

The SQDR Plus software maintains a table called SQ_STATISTICS which stores statistical data related to the operation of the software. Understanding how this data is collected and used can be helpful in evaluating the performance of the software on the host database.

The statistics in the SQ_STATISTICS table are periodic snapshots of in-memory statistics maintained by the SQDR Plus Capture Agent. These in-memory statistics are written/flushed to the SQ_STATISTICS table every two minutes.

The data contained in this table can be easily viewed in CAMAINT by choosing **Option 4 Display current statistics** or **Option 5 Display all statistics**. The statistics reflect the activity for the last 7 days, in hourly increments. Displaying the current statistics shows all of the statistics in the current hour. For example, if you check the statistics at 10:18 AM, the statistics shown are those that were captured in the last 18 minutes of activity (between 10:00AM and 10:18AM). If you would like to view statistics that were recorded before the current hour, use the **Display all statistics** option.

Active replications:

The TOTAL "Active replications" value is the number of source tables managed by a particular journal (and log reader). The journal name is indicated in parentheses. AVG and MAX values are also displayed.

Lag time:

The "Lag time" TOTAL represents in seconds the difference between the time when the source table received updates and the time when the data is moved to the staging table. If the lag time is close to zero, no significant backups have occurred during the hour. A large value indicates that the SQDR Plus log reader is lagging behind in processing changed data and may indicate that the source table received a very large number of changes in a short period of time and may be better suited for snapshot replication.

If "Lag time" MIN value is 2147483647and MAX value is –2147483648, then no statistics have been collected for the time period. If MIN/MAX have other values, those are the true MIN/MAX values collected for the time period covered by the statistic "bucket". An AVG value is also displayed.

Records processed:

The "Records processed" value is the number of records processed for the particular journal reader. This count includes records that were both committed and roll-backed and does not necessarily indicate the number of new updates that have occurred on the source table(s). If this value is lower than expected (e.g., the source tables for this journal received 100 new records in the last hourly unit of time, but the records processed value is only 10), it is reasonable to expect that this would be associated with a non-zero lag time.

Staged Rows:

The "Staged Rows" value is the aggregate number of rows left in the staging tables at the end of a pruning operation. The default behavior is for the Capture Agent to prune the staging tables every 30 minutes, so this value is typically updated twice every hour.

Prune Cycles:

The number of times pruning has run in the current hour. For default pruning settings (every 30 minutes), this value will be 0, 1, or 2.

Prune duration:

The shortest, longest and most recent pruning duration times (MIN, MAX and TOTAL values, in seconds) during the hour bucket.

Rows pruned:

The number of rows pruned in the current hour.

In addition to displaying statistics with CAMAINT, you can use an ODBC-enabled data analysis application and an ODBC driver such as StarSQL (or a JDBC application and a JDBC driver such as StarSQL for Java) to analyze the SQ_STATISTICS table.

The SQ_STATISTICS table contains the following fields:

Table 7:

COMPONENT	The qualified journal name or 'Pruner	VARCHAR
STATNAME	The type of statistic (Lag time, Records processed, Active replications, Staged Rows, Prune Cycles, Prune duration, Rows pruned). Note that these names are case-sensitive	VARCHAR
STATDATE	The date of the statistic entry	DATE
STATHOUR	The hour of the statistic entry (0-23)	INTEGER
STATTYPE	The type of the statistic (e.g., MIN, MAX, AVG, TOTAL)	VARCHAR
INTVALUE	The value of the statistic (except for AVG values)	INTEGER
DOUBLEVALUE	The value of the average (AVG) statistic.	DOUBLE

Index