Informatica® Fast Clone
11.0

Installation and Upgrade Guide

Table of Contents

Preface

The *Informatica Fast Clone Installation Guide* describes how to install and uninstall Informatica Fast Clone.

This guide is written for the person who is responsible for installing the product. The guide assumes that you have knowledge of the data cloning environment, including operating systems, source and target databases, and system resources.

For instructions on upgrading from an earlier Fast Clone release, see the *Fast Clone Release Notes*.

Informatica Resources

Informatica Network

Informatica Network hosts Informatica Global Customer Support, the Informatica Knowledge Base, and other product resources. To access Informatica Network, visit https://network.informatica.com.

As a member, you can:

- Access all of your Informatica resources in one place.
- Search the Knowledge Base for product resources, including documentation, FAQs, and best practices.
- View product availability information.
- Review your support cases.
- Find your local Informatica User Group Network and collaborate with your peers.

Informatica Knowledge Base

Use the Informatica Knowledge Base to search Informatica Network for product resources such as documentation, how-to articles, best practices, and PAMs.

To access the Knowledge Base, visit https://kb.informatica.com. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Documentation

To get the latest documentation for your product, browse the Informatica Knowledge Base at https://kb.informatica.com/_layouts/ProductDocumentation/Page/ProductDocumentSearch.aspx.

If you have questions, comments, or ideas about this documentation, contact the Informatica Documentation team through email at infa_documentation@informatica.com.

Informatica Product Availability Matrixes

Product Availability Matrixes (PAMs) indicate the versions of operating systems, databases, and other types of data sources and targets that a product release supports. If you are an Informatica Network member, you can access PAMs at
https://network.informatica.com/community/informatica-network/product-availability-matrices.

Informatica Velocity

Informatica Velocity is a collection of tips and best practices developed by Informatica Professional Services. Developed from the real-world experience of hundreds of data management projects, Informatica Velocity represents the collective knowledge of our consultants who have worked with organizations from around the world to plan, develop, deploy, and maintain successful data management solutions.

If you are an Informatica Network member, you can access Informatica Velocity resources at
http://velocity.informatica.com.

If you have questions, comments, or ideas about Informatica Velocity, contact Informatica Professional Services at ips@informatica.com.

Informatica Marketplace

The Informatica Marketplace is a forum where you can find solutions that augment, extend, or enhance your Informatica implementations. By leveraging any of the hundreds of solutions from Informatica developers and partners, you can improve your productivity and speed up time to implementation on your projects. You can access Informatica Marketplace at https://marketplace.informatica.com.

Informatica Global Customer Support

You can contact a Global Support Center by telephone or through Online Support on Informatica Network.

To find your local Informatica Global Customer Support telephone number, visit the Informatica website at the following link:
http://www.informatica.com/us/services-and-training/support-services/global-support-centers.

If you are an Informatica Network member, you can use Online Support at http://network.informatica.com.

CHAPTER 1

Installing Fast Clone

This chapter includes the following topics:

Installation Overview

Fast Clone Installation Overview

Fast Clone rapidly unloads Oracle data and loads it to heterogenous databases and data warehouse appliances. You can use Fast Clone to back up a source or to materialize or synchronize a target. Typically, you install Fast Clone on the Oracle source system.

Informatica provides a platform-specific .zip file or .tar.gz file for installing Fast Clone on a supported Linux, UNIX, or Windows system. To install Fast Clone, unzip or untar the installation file. You do not need to run a separate installer to get the product directories, executables, and files.

Prior to installation, it is important that you work with your Oracle DBA, system administrator, and target database administrators to plan the best Fast Clone implementation for your environment.

Depending on the unload method and target type that you use, you might need to perform some additional tasks on the source or target system to set up user permissions or connectivity.

Fast Clone supports the following unload methods:

- Direct path unload. Reads data directly from Oracle data files, without using the Oracle Call Interface (OCI). This method is much faster than the conventional path unload method. However, you cannot use the direct path unload method to run unload jobs from a system that is remote from the Oracle source, perform SQL JOIN operations on two or more tables, or unload data from views or cluster tables.
- Conventional path unload. Uses the OCI to retrieve metadata and data from the Oracle source. This method is slower than the direct path unload method, but it does not have the limitations of the direct path unload method.

If you use the direct path unload method, you must install Fast Clone on the Oracle source system and grant some system and object permissions to the Fast Clone user.

Fast Clone Components

Fast Clone includes the following components:

Informatica Fast Clone Console

The Fast Clone Console is the graphical user interface from which you configure and manage data-cloning jobs. From the Fast Clone Console, you can generate a cloning configuration file and run data-cloning jobs on a local or remote system. The cloning configuration file name is unload.ini by default. The Fast Clone Console runs on Linux, UNIX, and Windows. You can run it on the Oracle source system or on a standalone system. To start the Fast Clone Console, run gui.cmd on Windows or gui.sh on Linux or UNIX.

Fast Clone Executable

The Fast Clone executable, FastReader.exe on Windows or FastReader on Linux or UNIX, runs data-cloning jobs. You can start the Fast Clone executable from the command line or from the Fast Clone Console. If you run the Fast Clone executable from the command line, you can manually enter parameters to control unload processing.

Fast Clone Server

The Fast Clone Server is an optional add-on component that you can purchase to enable network communication across systems in a distributed Fast Clone topology. These systems can include the Fast Clone Console and Fast Clone instances on source and target database systems.

The Fast Clone Server runs as a Windows service or as a daemon on Linux or UNIX.

Use the Fast Clone Server to initiate unloads of Oracle data and metadata on a remote system or to transmit the output files to the target system. For example, use the Fast Clone Server to retrieve Oracle data that was unloaded by scheduled Fast Clone unload operations on separate Oracle systems and then make that data available to another Fast Clone instance for loading to the target database.

DataStreamer

The DataStreamer component is an add-on component that you can purchase for Amazon Redshift, Greenplum, Netezza, Teradata, and Vertica targets. For Greenplum, Netezza, and Teradata targets, DataStreamer is an optional component to load data faster. For Amazon Redshift targets, DataStreamer is a required component that is used transparently and always enabled.

With DataStreamer, you must use the direct path unload method. Depending on the target type, DataStreamer streams the unloaded Oracle data to the target in one of the following ways:

- For Amazon Redshift targets, DataStreamer sends the unloaded data to the Amazon Simple Storage Service (Amazon S3). After the source data is in Amazon S3 storage, Fast Clone issues a copy command that copies the data to the Amazon Redshift target tables.
  If you plan to run unload jobs on a Windows system, you must install the PostgreSQL ODBC driver on the system. If you plan to run unload jobs on a Linux and UNIX system, use the DataDirect ODBC driver for PostgreSQL that Fast Clone provides.

- For Greenplum targets, DataStreamer sends the unloaded data directly to the Greenplum parallel file distribution server (gpfdist) for loading to the target.

- For Netezza targets, DataStreamer writes the unloaded data to the named pipes that represent the Netezza external tables. The Netezza ODBC driver reads the data from these pipes and loads the data to the Netezza target tables.
  To use the Netezza DataStreamer, you must install the Netezza ODBC driver on the system where you plan to run unload jobs.

- For Teradata targets, you must install the Teradata Parallel Transporter (TPT) libraries on the system where you plan to run unload jobs. DataStreamer sends unloaded source data directly to a TPT operator (Update Operator, Load Operator, or Stream Operator) to load data to the target.
- For Vertica targets, DataStreamer uses the COPY command on the server side or the LCOPY command on the client side to send the unloaded data directly to Vertica targets.

Fast Clone Targets

Fast Clone unloads bulk data from Oracle sources. Fast Clone loads data to many types of relational databases and data warehouse appliances.

Fast Clone supports the following target types:

- Amazon Redshift
- Cloudera
- Flat files
- Greenplum
- Hive
- Hortonworks
- Microsoft SQL Server
- Netezza
- Oracle
- Teradata
- Vertica

Target Database Load Utilities

Fast Clone generates scripts to load data to targets.

The following table describes the target database load utilities that Fast Clone uses to load data:

Target Type	Loader
Amazon Redshift	LibPQ utility
Greenplum	gpload or gpfdist utility
Microsoft SQL Server	Microsoft SQL Server osql utility
Netezza	nzload utility
Oracle	SQL*Loader (sqlldr) utility
Teradata Database	FastLoad (TPT Load Operator) or MultiLoad (TPT Update Operator)
Vertica	COPY command on the server side LCOPY command on the client side

Fast Clone Use with Informatica Data Replication

If you have the Informatica Data Replication product, you can use Fast Clone as a high-speed alternative to the Data Replication InitialSync component to load target tables the first time.

Use Fast Clone for initial synchronization of target tables, prior to starting the Data Replication load operations, or to resynchronize target tables. Fast Clone can read data directly from Oracle data files, which results in better performance.

Fast Clone and Data Replication are separate products, and they do not have any version compatibility issues.

Installation Considerations

Before you install Fast Clone, review the following considerations:

- You can install the Fast Clone Console on a system that is remote from the Oracle source system and Fast Clone executable that moves data. For example, you might want to run the Fast Clone Console separately to reduce overhead on the source system and to centralize Fast Clone configuration and management on a system without resource constraints.

- The Fast Clone instance on the Oracle source system requires the Oracle Call Interface oci.dll, which is available as part of the installed Oracle Client.

- From the Fast Clone Console, you can test connection settings before loading data. For all target types except flat files, the Fast Clone Console requires a JDBC driver to connect to the target. Fast Clone provides JDBC drivers for most targets in its lib subdirectory. However, for Netezza and Teradata targets, you must download the appropriate JDBC drivers.

- If you have Greenplum, Netezza, Teradata, or Vertica targets, you can optionally use the DataStreamer component. DataStreamer streams data to these targets to improve Fast Clone performance and reduce disk space use. DataStreamer also can perform some limited dynamic data transformation processing. To use DataStreamer, you must install Fast Clone on the Oracle source system. Also, verify that the following Greenplum or Teradata software for bulk loading is available on the targets:

  - For Teradata targets, DataStreamer requires the Teradata Parallel Transporter (TPT) libraries with load and update operators that are equivalent to FastLoad (TPT Load Operator) and MultiLoad (TPT Update Operator).

  - For Greenplum targets, DataStreamer requires the Greenplum parallel file distribution server (gpfdist). Also, verify that this server can connect to the FastClone port on the source to get data. Items such as firewalls can block this connection.

- If you want to use the optional Fast Clone Server, you can run it on a system that is remote from the Oracle source system. The server can communicate over the network with one or more Oracle source systems to get Fast Clone configuration information and initiate unload requests. A server also can receive output data files and control files from the Fast Clone instance on the source and send these files either to another Fast Clone Server or to the Fast Clone instance on the target. With a Fast Clone Server, bulk data movement latency might be prolonged because of increased network activity.

- To use the Fast Clone Server, the oci.dll must be available locally. To get it, you can download the 64-bit Oracle Instant Client from http://www.oracle.com/technetwork/database/features/instant-client/index.html. Then specify the oci.dll in the Path environment variable.

- You can use the TCP/IP protocol with the Secure Sockets Layer (TCPS) to establish secure network connections between Fast Clone and Oracle source and target systems. If the Fast Clone Console uses TCPS connections to remote sources and targets, you must install and configure the 64-bit Oracle Instant Client on the Fast Clone Console system. For more information, see the Oracle documentation.

Installation Task Flow

To create a Fast Clone installation, you must complete some tasks before and after installing the software.

Use the following high-level task flow:

1. Perform the pre-installation tasks:

    - Get a license key.

    - To use the Fast Clone Console, install the Java Runtime Environment (JRE).

    - To run the Fast Clone Console on Linux or UNIX, install the X Window Server.

    - Configure the ORACLE_HOME and ORACLE_SID environment variables, if not already defined, on the Oracle source system.

2. Install Fast Clone on the Oracle source system.

3. Set up the system user under which Fast Clone runs.

4. Optional. If you do not want to run the Fast Clone Console on the Oracle source system, install it on another system from which you want to configure and manage bulk data movement jobs. Complete the normal Fast Clone installation procedure.

5. Optional. To use the Fast Clone Server, install it on a Linux, UNIX, or Windows system in your environment. Complete the normal Fast Clone installation. On Windows, also run install_server.cmd.

6. Perform post-installation configuration tasks:

    - Verify that the Fast Clone directories, executables, and files are present. See Appendix A, "Installed Fast Clone Directories and Files" on page 26.

    - Verify that you can start the Fast Clone Console.

    - Enter the license key.

    - Verify the Fast Clone version.

    - If you use the direct path unload method, set up Fast Clone user permissions in Oracle that allow read access to the physical data files.

    - Optional. If you want to test connection settings for targets from the Fast Clone Console, get the JDBC drivers that the Fast Clone Console uses to connect to the targets. Fast Clone provides drivers for many targets in its lib subdirectory. For other targets, you must download the drivers.

    - If you use DataStreamer with Teradata or Greenplum targets, download the Teradata TPT libraries to the system where DataStreamer runs or verify that you have the Greenplum parallel file server (gpfdist) installed.

After you complete the installation tasks, you can use the Fast Clone Console or command line interface to configure bulk data movement jobs. For more information, see the *Informatica Fast Clone User Guide*.

Before You Install

Before You Install Tasks

Before you install Fast Clone, you need to complete some tasks in your environment.

Perform the following tasks:

1. Get a license key and an installation file for your platform.

2. To run the Fast Clone Console, install the Java Runtime Environment (JRE) on the system where the Fast Clone Console will run.

    Note: You can use the command line interface instead.

3. To run the Fast Clone Console on a Linux or UNIX system, set up an X Window environment.

Getting a License Key

Get a valid license key for Fast Clone.

If you download the installation files from the Informatica Electronic Software Download (ESD) site, the license key is delivered in an email message from Informatica.

If you did not get a license key, contact Informatica Global Customer Support.

Restrictions of the Trial License

In addition to the limited trial period, the trial license limits the number of records that the Fast Clone executable can unload from a source database.

The following table describes the restrictions on the number of rows that the trial version of Fast Clone can unload per job:

Restriction	Maximum Result Set
Rows per table that uses WHERE clauses	1000
Rows per table that does not use WHERE clauses	10000
Total rows per unload job	100000

Important: To calculate the total number of unloaded rows, Fast Clone always uses the maximum number of rows that the trial license permits to unload from a given table. For example, the trial version of Fast Clone can unload a maximum of 10 tables that do not use WHERE clauses per unload job. For each of these tables, Fast Clone unloads a maximum of 10000 rows.

Installing the Java Runtime Environment

The Informatica Fast Clone Console is a Java application that requires the Java Runtime Environment (JRE) 1.8. For Hadoop and Hive Targets, the Fast Clone executable requires 64-bit JRE.

If the JRE is not installed, you can download it for free from http://www.oracle.com/technetwork/java/javase/downloads/index.html. Then install it on the system where the Fast Clone Console will run.

Also, set the JRE_HOME environment variable to the JRE base directory. On Linux or UNIX, include JRE_HOME/bin in the PATH environment variable. On Windows, include %JRE_HOME%\bin in the PATH environment variable.

To verify that the correct JRE version is installed, issue the following command from the command line:

```
java -version
```

Setting Up an X Window Environment

To run on Linux or UNIX, Informatica Fast Clone Console requires an X Window environment. If an X Window environment does not exist, you must set it up.

X Window uses a client-server model, where an X server communicates with client applications such as the Fast Clone Console to provide graphical user interface display and I/O services.

Set up an X server on the workstation where you want to run the Fast Clone Console. To determine which X server your organization uses, contact your system administrator.

Also, configure the DISPLAY environment variable. This variable specifies the IP address of the system with the X server plus a display number. The default display number is 0. Use :0.0 if the X server runs on the local system. After the X server is running, it opens a TCP/IP listener port, usually port 6000.

If the Fast Clone Console client application is remote from the workstation where the X server runs, configure a DISPLAY environment variable on the client system that points to the X server on your workstation, where you want to display the Fast Clone Console.

The following table shows the commands to set the DISPLAY environment variable:

Shell	Command
C	setenv DISPLAY *X_server_ip_address*:0
Bash/Korn	export DISPLAY="*X_server_ip_address*"
Bourne	DISPLAY="*X_server_ip_address*:0" export display

If you run Fast Clone gui.sh and get the following error message, the X Window environment is not properly configured:

```
Exception in thread "main" java.lang.InternalError: Can't connect to X11 window server
using 'x.y.z.w' as the value of the DISPLAY variable.
```

Installing Fast Clone

Installation Overview

Install a Fast Clone instance on the Oracle source system from which data will be unloaded.

You can run the Fast Clone Console on another Linux, UNIX, or Windows system if you have a busy Oracle system and want to avoid increasing its overhead.

Also, you can install the optional Fast Clone Server on a remote system in the network. The Fast Clone Server can receive unload.ini configuration files from one or more sources and initiate unload processing. The Fast Clone Server also can receive output data and control files and forward the files to another Fast Clone Server or to the Fast Clone instance on the target.

For each location where you install a Fast Clone component, complete the normal Fast Clone installation. You unzip or untar the platform-specific installation file. You do not need to run a separate installer program.

Installing Fast Clone on Linux or UNIX

To install Fast Clone on Linux or UNIX, use the platform-specific installation .tar.gz file that Informatica provides for your operating system.

You can install and run Fast Clone under the Linux or UNIX user account "Oracle." This Oracle Linux or UNIX account allows read access to Oracle physical data files, as required for bulk data movement. If you cannot use this account, you must configure an administrative user.

1. Transfer the platform-specific installation tar.gz file that you received from Informatica to a directory on the Linux or UNIX system with the Oracle source instance.

   Informatica recommends that you retain this file in case you need to reinstall the product.

2. Create a directory for the Fast Clone version that you are installing, for example:

   ```
   /Informatica/FastClone10
   ```

3. Copy the installation tar.gz file to the new directory.

4. Decompress the tar.gz file into the new directory. Use a program such as gunzip.

   Depending on which decompression program you use, decompression might result in a .tar file or directly extract the Fast Clone files.

5. If decompression resulted in a .tar file, untar the file to extract the Fast Clone files. Use the standard untar command syntax for your platform.

6. Define the FAST_READER_HOME environment variable to identify the Fast Clone installation root directory.

7. If you are installing Fast Clone on an AIX operating system, set the following environment variables:

   ```
   AIXTHREAD_SCOPE=S
   AIXTHREAD_STK=8000000
   ```

   Note: If you are using the Informatica Fast Clone Console start script (unload.sh), this step is optional. If you are not using this script, you must set these environment variables. For more information, see the *Informatica Fast Clone User Guide*.

Installing Fast Clone on Windows

To install Fast Clone on Windows, use the platform-specific installation .zip file that Informatica provides for your operating system.

1. Transfer the installation .zip file to a directory on the Windows system with the Oracle source instance.

   Informatica recommends that you retain this file in case you need to reinstall the product.

2. Create a directory for the Fast Clone version that you are installing, for example:

   ```
   C:\Informatica\FastClone10
   ```

3. Extract the contents of the installation .zip file into the new directory.

   Select the **Use folder names** check box to maintain the directory structure.

4. Optional. Create a shortcut on the desktop for starting the Fast Clone Console.

   Right-click the gui.cmd file in the top-level installation directory and select **Create shortcut**. Then drag the shortcut to the desktop.

5. Define the FAST_READER_HOME environment variable to identify the Fast Clone installation root directory.

Setting Up the System User Under Which Fast Clone Runs

You must run Fast Clone programs under a system user account that has the required authority level.

Use one of the following types of user accounts:

- Use the Linux or UNIX user account named "Oracle," also referred to as "Oracle UNIX account" in this documentation.
- Create an administrative user in the DBA group and make this user the owner of the Fast Clone files. You must also configure some permissions on the Fast Clone executable files.

Configuring Fast Clone to Run Under an Administrative User on Solaris

If you cannot install and run Fast Clone under the Oracle UNIX account, you can create an administrative user in the DBA group. Then copy the Fast Clone files under the ownership of this user.

Note: For other supported Linux and UNIX systems, the steps vary slightly.

1. On the Solaris system, log in as the root user.
2. Create an administrative user in the DBA group. For example:

   ```
   as root useradd -g dba -s /usr/bin/tcsh -m -d /home/fastclone fastcln
   ```
3. Copy the Fast Clone files under ownership of the administrative "fastcln" user.
4. Add the $ORACLE_HOME/lib directory and the Fast Clone directories to the /var/ld/ld.config configuration file for the dynamic linker. For example, use the following command:

   ```
   crle -s /usr/lib:${ORACLE_HOME}/lib:'pwd'/support -l/usr/lib:${ORACLE_HOME}/
   lib:'pwd'/support
   ```
   Note: Alternatively, you can use the LD_LIBRARY_PATH environment variable for invoking the Fast Clone Console or Fast Clone Server.
5. Change the effective user to the administrative user:

   ```
   su -fastcln
   ```
6. Set full read and execution permissions on Fast Clone program files and shared object (.so) files in the Fast Clone installation path. Use the chmod command:

   ```
   cd Fast_Clone_install_directory
   chmod 755 FastReader
   chmod 755 QZip
   chmod 755 support/*so*
   ```
7. Set the group sticky bit on individual Fast Clone executables to enable them to run under the "dba" group instead of the group of the user that executes the program. Use the chmod command with g+s. For example:

   ```
   chmod g+s FastReader
   ```

Configuring Fast Clone to Run Under an Administrative User on Linux and UNIX Systems Other Than Solaris

If you cannot install and run Fast Clone under the Oracle UNIX account, you can create an administrative user in the DBA group. Then copy the Fast Clone files under the ownership of that user.

Note: For Solaris systems, a few more steps are required.

1. Create an administrative user in the DBA group. For example:

   ```
   as root useradd -g dba -s /usr/bin/tcsh -m -d /home/fastclone fastcln
   ```
2. Copy the Fast Clone files under ownership of the administrative "fastcln" user. You can use the chown command.

Installing Fast Clone Server

Fast Clone Server is an optional network server that you can install to communicate with remote Fast Clone instances to perform unloads and store files centrally.

On Linux and UNIX, Fast Clone Server is a daemon that runs as a background process. To install it, just untar the platform-specific Fast Clone installation file. The server.sh file for starting Fast Clone Server resides in the top-level installation directory. You do not need to install the server separately.

On Windows, Fast Clone Server runs as a Windows service. To install it, unzip the platform-specific installation file and then run the install_server.cmd file that is located in the top-level installation directory.

Tip: For the install_server.cmd executable to run, the Oracle Client oci.dll must be available. Make sure that you include oci.dll in the PATH environment variable.

After You Install

Post-Installation Tasks

After installation, verify the installation and configure the environment for the new installation.

Complete the following tasks:

- Verify that the expected Fast Clone directories and files were installed.
- Verify that you can start the Fast Clone Console.
- Enter the license key.
- Verify the Fast Clone version.
- Create a database user ID to use for connecting to the Oracle source database.
- Configure Fast Clone user permissions in Oracle that allow Fast Clone to read physical data files.
- If you want to test target connection settings from the Fast Clone Console, get the JDBC drivers that the Fast Clone Console uses to connect to the data targets.
- If you use DataStreamer with Teradata or Greenplum targets, download the Teradata TPT libraries to the system where DataStreamer runs or verify that you have the required Greenplum parallel file server (gpfdist) installed.
- If you use TDE-encrypted columns and tablespaces in the Oracle source database, get the Oracle .jar files for Oracle Transparent Data Encryption.
- If you want to use the TCPS protocol to provide security for network connections between the Fast Clone Console and Oracle sources or targets, get the Oracle .jar files for Secure Sockets Layer (SSL).
- If Fast Clone is installed on a remote computer from the Oracle source database, or if Fast Clone is installed on the same computer as a 32-bit Oracle source database, verify that the 64-bit Oracle Client is installed.

Verifying Installed Directories and Files

After you unzip or untar the installation file, verify that the expected Fast Clone directories, executables, and files appear on your system.

For more information, see .

Testing the Fast Clone Console

Verify that you can start the Fast Clone Console.

On Windows, perform one of the following actions:

- Double-click the Fast Clone Console shortcut on the desktop, if you created one.
- Double-click gui.cmd in the top-level installation directory.
- Enter the full path to gui.cmd in the Run window.

On Linux or UNIX, cd to the Fast Clone installation directory and run gui.sh.

Entering the License Key

To use Fast Clone, you must enter the Informatica license key value that you received in the FastClone.key file.

Note: You do not have to register the license key with Informatica in a separate process.

Either enter the license key from the Informatica Fast Clone Console or manually copy the key value into the FastClone.key file, replacing the entire contents.

To enter the license key value in the Fast Clone Console, perform the following steps:

1. Start the Informatica Fast Clone Console.
2. On the menu bar, click **Help** > **Enter License Key**.
3. In the **License Key Entry** dialog box, enter the license key value that you received for Fast Clone.
4. Click **OK**.

   The license key value is written to the FastClone.key file.

Verifying the Fast Clone Version

From the Informatica Fast Clone Console, you can verify the Fast Clone version that is installed.

▶ On the menu bar, click **Help** > **About**.

   The **About** dialog box displays the Fast Clone version followed by the licensee name and Fast Clone copyright.

Creating a Database User for Connecting to the Oracle Source from Fast Clone

You must define a database user ID that has the authority to connect to the Oracle source.

On Linux or UNIX, you can use the "oracle" user. Alternatively, you can create a database user ID, for example:

```
CREATE USER FASTCLONE_USER PROFILE "DEFAULT"
IDENTIFIED BY FASTCLONE_USER_PASSWORD DEFAULT TABLESPACE "USERS"
TEMPORARY TABLESPACE "TEMP"
QUOTA UNLIMITED on "USERS" ACCOUNT UNLOCK;
```

Granting Privileges to the Oracle Database User

Grant privileges to the Oracle database user that you defined. Grant the user privileges that Fast Clone requires to connect to the Oracle database, retrieve source metadata, and unload data.

1. To perform a checkpoint before Fast Clone unloads the source data, grant the ALTER SYSTEM privilege or create a stored procedure.

    - To grant the ALTER SYSTEM privilege, issue the following statement:

        ```
        GRANT ALTER SYSTEM TO FASTCLONE_USER;
        ```

    - To configure Fast Clone to use a stored procedure for checkpointing, perform the following steps:

        - Create the stored procedure. Use the following statements:

            ```
            CREATE Procedure FastCloneCheckPoint
            IS
            BEGIN
            execute immediate 'ALTER SYSTEM CHECKPOINT GLOBAL';
            END;
            ```

        - Grant the EXECUTE privilege to this stored procedure. Use the following statement:

            ```
            GRANT EXECUTE ON FastCloneCheckPoint TO FASTCLONE_USER;
            ```

        - Edit the unload.ini file to use the stored procedure. You must add the non_standard_ckpt_block parameter. Use the following statement:

            ```
            non_standard_ckpt_block="begin procedure_owner.FastCloneCheckPoint(); end;"
            ```

        Note: The user who owns and creates the stored procedure must have the ALTER SYSTEM privilege or DBA role.

2. To connect to the Oracle source database, grant the CREATE SESSION or CONNECT privilege. Use one of the following statements:

    ```
    GRANT CREATE SESSION TO FASTCLONE_USER;
    GRANT CONNECT TO FASTCLONE_USER;
    ```

3. To unload source data with the conventional path unload method, grant the SELECT ANY TABLE privilege. Use the following statement:

    ```
    GRANT SELECT ANY TABLE TO FASTCLONE_USER;
    ```

    Tip: You can grant the SELECT privilege on each source table that you want to unload instead of granting the SELECT ANY TABLE privilege.

4. To retrieve metadata about the structure of the source tables, grant the SELECT ANY DICTIONARY privilege. For Oracle 12*c* sources, also grant the SELECT privilege for the SYS.USER$ table. If the Oracle 12*c* sources use Oracle Transparent Data Encryption (TDE), grant the SELECT privilege on the SYS.ENC$ table too.

    a. Grant the SELECT ANY DICTIONARY privilege. Use the following statement:

        ```
        GRANT SELECT ANY DICTIONARY TO FASTCLONE_USER;
        ```

    b. For Oracle 12*c* sources, grant the SELECT privilege for the SYS.USER$ table. Use the following statement:

        ```
        GRANT SELECT ON SYS.USER$ TO FASTCLONE_USER;
        ```

    c. For Oracle 12*c* sources that use Oracle TDE on columns and tablespaces, grant the SELECT privilege for the SYS.ENC$ table. Use the following statement:

        ```
        GRANT SELECT ON SYS.ENC$ TO FASTCLONE_USER;
        ```

5. If you configure Fast Clone to lock the tables that are unloaded with the direct path unload method, grant the LOCK ANY TABLE privilege or grant the SELECT privilege on each table that you want to unload. Use one of the following statements:

    ```
    GRANT LOCK ANY TABLE TO FASTCLONE_USER;
    GRANT SELECT ON TABLE TABLE_NAME TO FASTCLONE_USER;
    ```

Note: The SELECT ANY TABLE privilege is not sufficient to lock the source tables when unloading data.

6. If you also use Informatica Data Replication, grant the EXECUTE ON DBMS_FLASHBACK privilege. Use the following statement:

```
GRANT EXECUTE ON DBMS_FLASHBACK TO FASTCLONE_USER;
```

Fast Clone requires this privilege to get the current SCN of the source database.

Installing JDBC Drivers for the Fast Clone Console Connectivity to Targets

If you want to test connection settings for data targets from the Fast Clone Console before loading data, verify that you have the JDBC drivers that the Fast Clone Console requires to connect to the targets. Fast Clone provides JDBC drivers for many targets in its lib subdirectory. For other targets, you must download the appropriate JDBC drivers.

Note: This task is optional and for test purposes only.

The Fast Clone lib subdirectory includes JDBC drivers for the following targets:

- Microsoft SQL Server
- Oracle
- Greenplum and Vertica targets

For the following targets, you must download the JDBC drivers from the appropriate locations to the Fast Clone lib subdirectory:

- For Netezza targets, download and install the Netezza JDBC driver.
- For Teradata targets, download and install the Teradata JDBC driver. Ensure that you get the tdgssconfig.jar and terajdbc4.jar files.

Note: The Fast Clone Console does not need a JDBC driver to connect to flat files.

Downloading Teradata Client Libraries

If you have a Teradata target, you must download the Teradata Tools and Utilities (TTU) package from the Teradata web site and configure ODBC connectivity for the Teradata target on the system where you run data unload jobs.

1. Record the following information about the target system:
   - Operating system and CPU type
   - Teradata database version
2. Download a compatible version of the Teradata Tools and Utilities (TTU) from the Teradata web site.
3. Install the following Teradata client-side libraries:
   - Teradata BTEQ Utility
   - Teradata Call-Level Interface (CLIv2)
   - Teradata Data Connector API
   - Teradata FastLoad utility
   - Teradata GSS client package
   - Teradata MultiLoad utility

- Teradata ODBC driver

- Teradata Parallel Transporter (TPT) Base

- TPT API

- TPT Infrastructure

- TPT Load Operator

- TPT Stream Operator

- TPT Update Operator

Important: If you do *not* plan to use DataStreamer, install all of the following tools and utilities: FastLoad utility, MultiLoad utility, and TPT libraries. If you plan to always use DataStreamer, you can install the TPT libraries only.

4. On Linux and UNIX, configure the DataDirect ODBC driver manager that Fast Clone provides in the *FactClone_installation*/dd subdirectory.

RELATED TOPICS:

- "Configuring the DataDirect ODBC Driver Manager" on page 20

Configuring Connectivity for Netezza DataStreamer

If you use Fast Clone DataStreamer to stream data to a Netezza target, you must configure ODBC connectivity for the Netezza target on the system where you run data unload jobs.

- On Windows, download and install the 64-bit ODBC driver for Netezza targets.

- On Linux and UNIX, complete the following tasks:

  - Download and install the 64-bit DataDirect ODBC driver for Netezza targets. Put the driver files in the *FastClone_installation*/dd/lib subdirectory.

  - Configure the DataDirect ODBC driver manager that Fast Clone provides in the *FactClone_installation*/dd subdirectory.

RELATED TOPICS:

- "Configuring the DataDirect ODBC Driver Manager" on page 20

Configuring the DataDirect ODBC Driver Manager

On Linux and UNIX, you must configure the DataDirect ODBC driver manager that DataStreamer requires to connect to Netezza and Teradata targets.

The DataDirect ODBC driver manager provides an ODBC API that Fast Clone uses to communicate with the ODBC drivers.

1. Verify the driver settings for Netezza and Teradata targets in the *FastClone_installation*/dd/odbcinst.ini file, including the path to the DataDirect ODBC drivers.

2. Set the DD_INSTALLDIR environment variable to point to the *FastClone_installation*/dd directory.

3. Add DataDirect driver manager libraries to the system path.

   Fast Clone provides the DataDirect driver manager libraries in the *FastClone_installation*/dd/lib directory. Use a system path environment variable to point to the directory that contains the DataDirect driver libraries.

The following table describes system path environment variables for the supported Linux and UNIX operating systems:

Operating System	Environment Variable
AIX	LIBPATH
HP-UX	SHLIB_PATH
Linux	LD_LIBRARY_PATH
Solaris	Use the LD_LIBRARY_PATH_64 environment variable if it is already defined in the system. Otherwise, use the LD_LIBRARY_PATH environment variable.

4. Set the ODBCINST environment variable to point to the *FastClone_installation*/dd/odbcinst.ini file. For example:

```
export ODBCINST=$DD_INSTALLDIR/odbcinst.ini
```

Getting the Oracle JAR Files Required for Transparent Data Encryption

If you use Oracle Transparent Data Encryption (TDE) to encrypt columns and tablespaces in Oracle sources, you must copy some Oracle .jar files to the *FastClone_installation*\lib subdirectory.

The Fast Clone Console uses these .jar files to open the Oracle wallet and save the wallet keys in the cloning configuration file. Fast Clone uses the keys to decrypt the source data during unload processing.

1. Copy the following libraries from %ORACLE_HOME%/jlib subdirectory to the *FastClone_installation*/lib subdirectory and replace the existing .jar files with those that have the same names:

    - oraclepki.jar
    - osdt_cert.jar
    - osdt_core.jar

2. Restart the Fast Clone Console.

Getting the Oracle JAR Files Required for Secure Sockets Layer (SSL) Authentication and Encryption

To use the TCP/IP protocol with the SSL to connect to Oracle sources and targets from the Fast Clone Console, you must copy some Oracle .jar files to the *FastClone_installation*\lib subdirectory on the Fast Clone Console system.

1. Copy one of the following .jar files from the %ORACLE_HOME%/jdbc/lib subdirectory to the *FastClone_installation*/lib subdirectory, depending on your Oracle version:

    - ojdbc5.jar for Oracle 11*g*
    - ojdbc6.jar for Oracle 11*g*
    - ojdbc7.jar for Oracle 12*c*

    Note: For Oracle 11*g*, copy either ojdbc5.jar or ojdbc6.jar. You do not need both files.

2. Restart the Fast Clone Console.

Note: The Fast Clone Console requires the 64-bit Oracle Instant Client to be installed and configured to use the TCPS protocol to connect to remote Oracle sources and targets. For more information, see the Oracle documentation.

Preparing Amazon Redshift Targets

To load source data to the Amazon Redshift target, you must complete a few tasks to prepare the Amazon Simple Storage Service (Amazon S3).

1. Create an Amazon S3 account.
2. Ensure that you have valid Amazon Web Services (AWS) credentials to access the Amazon S3 API.
3. Create an Amazon S3 bucket for the temporary files that will contain the source data to be loaded to the Amazon Redshift target.

Preparing Fast Clone for Hadoop and Hive Targets

To load source data to a Hadoop distribution, you must complete several tasks to prepare the Fast Clone system. Fast Clone can load data to Cloudera CDH, Hive, and Hortonworks targets.

1. Install the 64-bit Java Runtime Environment (JRE) 1.8.
2. Define the JAVA_HOME environment variable to point to the root Java installation directory.
3. Add a Java library to the system path.
   - On Windows, add the directory that contains the jvm.dll library to the PATH environment variable. For example, use the following command:

     ```
     PATH=%PATH%;%JAVA_HOME%\jre\bin\server
     ```
   - On Linux and UNIX, add the directory that contains the libjvm.so library to the library path environment variable for your operating system. The library path environment variables are:

     - LD_LIBRARY_PATH for Linux and Solaris systems

     - LIBPATH for AIX systems

     - SHLIB_PATH for HP-UX systems

     For example, use the following command:

     ```
     LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$JAVA_HOME/jre/lib/amd64/server
     ```
4. Unpack the .jar file package to the *FastClone_installation* directory.

   If you do not want to copy the .jar files to the *FastClone_installation* directory, define the IFC_HDFS_LIBS environment variable. In this variable, list the paths to the .jar files, including the file names. These paths can be either full paths or relative to the *FastClone_installation* directory.
   - On Windows, use the semicolon as a separator in the list of paths to the .jar files. For example:

     ```
     SET IFC_HDFS_LIBS=lib/hadoop/hadoop-core-0.20.2-cdh3u4.jar;lib/hadoop/
     guava-11.0.2.jar
     ```
   - On Linux and UNIX, use a colon as the separator in the list of paths to the .jar files. For example:

     ```
     IFC_HDFS_LIBS=lib/hadoop/hadoop-core-0.20.2-cdh3u4.jar:lib/hadoop/guava-11.0.2.jar
     ```

Uninstalling Fast Clone

About Uninstalling

If you no longer plan to use Fast Clone or a specific Fast Clone version, you can uninstall it.

First verify that no bulk data movement jobs are running or are scheduled to run with the old Fast Clone version.

Then uninstall the Fast Clone instance on the Oracle source system. If you run the Fast Clone Console and Fast Clone Server on separate systems, also uninstall those components.

Uninstalling Fast Clone

To totally uninstall Fast Clone on the Oracle source system or another system, manually delete the top-level Fast Clone installation directory.

Note: On a Windows system, you cannot remove the Fast Clone program from the Control Panel.

If you created a Fast Clone user role in Oracle with the permissions that are required for the direct path unload method, you can also remove that user.

Uninstalling the Fast Clone Server

If you installed the optional Fast Clone Server and no longer want to use it, run uninstall_server.cmd to remove it as a Windows service.

This file is in the top-level Fast Clone installation directory, along with the cmd files for installing, starting, and stopping the server.

If you are uninstalling Fast Clone totally, also delete the top-level installation directory.

CHAPTER 2

Upgrading Fast Clone

This chapter includes the following topics:

Overview

To upgrade Fast Clone to the current release, Informatica recommends that you extract or decompress the Fast Clone installation file for the current release into the existing Fast Clone installation directory.

If you run Fast Clone on multiple systems, upgrade Fast Clone on all of the systems. Do not run the current release on one system and run an earlier Fast Clone release on another system in the same topology.

Important: If you configured Hadoop or Hive targets in a previous version of Fast Clone, you must create new configurations for these targets after upgrading Fast Clone to the current release.

Upgrading Fast Clone

Upgrade Fast Clone on each system where you run a Fast Clone component such as the Fast Clone Console, Fast Clone executable, and Fast Clone server. Use this procedure for an existing installation.

1. Transfer the platform-specific Fast Clone installation file for the current release that you received from Informatica to a directory other than the existing Fast Clone installation directory on your system.

    Tip: Retain the installation file to upgrade other Fast Clone instances or redo an upgrade.

2. Wait for all of the Fast Clone unload jobs that are running to complete.

3. Stop all Fast Clone Server instances that are running under the earlier Informatica Fast Clone version.

    - On Windows, run the stop_server.cmd script, or run the following commands from the command line:

      ```
      cd Fast_Clone_installation_directory
      FastReader.exe RUN_AS_SERVICE=true SERVER_CONFIG=server.ini -k
      ```

    - On Linux or UNIX, run the following command from the command line:

      ```
      kill -9 fast_clone_server_process_id
      ```

4. Back up the existing Fast Clone installation, including the cloning configuration files.

5. Extract or decompress the installation file into the existing Fast Clone installation directory:

- On Linux, decompress the FastClone_*v_r_m*.linux.x86_64.tar.gz file.

- On Windows, extract the FastClone_*v_r_m*.win.64t.zip file.

Confirm that the files for the previous Fast Clone version were overwritten.

6. Copy the new license key file, FastClone.key, to the Fast Clone installation directory.

7. Verify that the FAST_READER_HOME environment variable still points to the Fast Clone installation.

8. Verify that the required version of Java Runtime Environment (JRE) is installed. To determine the installed version, issue the following command from the command line:

```
java -version
```

9. If you use the 32-bit Oracle Client in either of the following situations, you must install the 64-bit Oracle Client:

- Fast Clone and the 32-bit Oracle Client are installed on a computer that is remote from the Oracle source database.

- Fast Clone is installed on the same computer as the Oracle source database, and the Oracle version is 32-bit.

10. If you are installing Fast Clone on an AIX operating system, set the following environment variables:

```
AIXTHREAD_SCOPE=S
AIXTHREAD_STK=8000000
```

Note: If you are using the Informatica Fast Clone Console start script (unload.sh), this step is optional. If you are not using this script, you must set these environment variables. For more information, see the *Informatica Fast Clone User Guide*.

11. Restart all Fast Clone Server instances.

- To start the Fast Clone Server as a service on Windows, run the start_server.cmd script or issue the following commands from the command line:

```
cd Fast_Clone_installation_directory
FastReader.exe RUN_AS_SERVICE=true SERVER_CONFIG=server.ini -s
```

- To start the Fast Clone Server as a daemon on Linux or UNIX, run the server.sh script.

Installed Fast Clone Directories and Files

This appendix includes the following topic:

Fast Clone Directories and Files

After you complete installation, verify that the Fast Clone root directory and subdirectories contain the expected files.

Subdirectories

The following subdirectories appear under the Fast Clone root directory that you created at installation:

ddl

Contains some data type mappings for creating DDL statements. Use these files only at the direction of Informatica Global Customer Support.

dd

On Linux and UNIX, contains the DataDirect ODBC driver libraries.

doc

Contains Fast Clone documentation, including the release notes.

lib

Contains Java archive (.jar) files for the Fast Clone Console, including the JDBC drivers that the Fast Clone Console uses to optionally connect to target databases.

logs

Contains logs of Fast Clone Console error and diagnostic information.

support

Contains runtime libraries that Fast Clone uses.

uiconf

Contains global defaults for configuration files, and data type mapping and conversion rules.

Executables and Script Files

The following files are in the Fast Clone root directory. The .sh files occur on Linux or UNIX, and the .cmd and .exe files occur on Windows.

cmd_unload_example.sh or cmd_unload_example.cmd

An example of how to run Fast Clone from the command line. The command syntax includes Fast Clone configuration parameters.

Fast Clone or FastReader.exe

The Fast Clone engine for bulk data movement that is started from the unload.sh or unload.cmd file. You can also run it from the command line interface.

FastReaderTD.exe

The Fast Clone engine with the optional DataStreamer component that streams bulk data to a Teradata target.

gui.sh or gui.cmd

Runs the Fast Clone Console.

install_server.cmd

Installs the optional Fast Clone Server component as a Windows service.

Qzip or Qzip.exe

Runs the Fast Clone decompression utility. Use this utility on unloaded data that was compressed with a high-speed compression routine of type QZIP during unload processing.

server.sh or start_server.cmd

Starts the Fast Clone Server component.

On Linux or UNIX, server.sh starts the Fast Clone Server as a background process. On Windows, start_server.cmd starts the Fast Clone Server as a Windows service.

stop_server.cmd

Stops the Fast Clone Server component on Windows.

Note: On Linux or UNIX, you can use the `kill -9` *pid* command to stop the Fast Clone Server process.

unload.sh or unload.cmd

Sets the Fast Clone environment to the current directory. Then starts the Fast Clone executable, FastReader or FastReader.exe, based on the settings in the unload.ini configuration file to perform unload processing.

uninstall_server.cmd

Uninstalls the Fast Clone Server on Windows.

Other Files

FastClone.key

Contains the Fast Clone license key that you enter manually or from the Fast Clone Console.

server.ini

Contains configuration parameters for the optional Fast Clone Server component.

unload.ini

Contains configuration parameters for Fast Clone. This file is not present at installation completion. It appears after you configure a bulk data movement job in the Fast Clone Console.

INDEX

V

version, product
    viewing version from the Fast Clone Console <u>17</u>

X

X Window environment
    setting up for Fast Clone Console on Linux or UNIX <u>13</u>