





Migrating to IBM Information Server Version 8



Migrating to IBM Information Server Version 8

Note

Before using this information and the product that it supports, read the information in “Notices” on page 47.

Contents

Chapter 1. Planning your migration . . . 1

Installation layers.	1
Version Control support	2

Chapter 2. Migrating from earlier DataStage versions 3

DataStage release 7.0 to 7.5 installations	4
Basic method: Migrate using the installation program	5
Saving configuration files	5
Removing DataStage clients	6
Installing IBM Information Server 8.	6
Migrating DataStage projects	6
Restoring contents of configuration files	7
Complete method: migrate to a clean installation of DataStage	7
Exporting projects	8
Saving configuration files	9
Job dependencies	10
Removing previous versions of DataStage	11
Installing IBM Information Server 8	12
Importing DataStage projects	12
Restoring contents of configuration files	13
Restoring job dependencies	13
Recompiling jobs	14
Multiple server method (UNIX and Linux only)	14
Installing IBM Information Server in addition to existing DataStage servers	16
Installing WebSphere DataStage 8.0 clients in addition to existing DataStage clients	17
Starting and stopping multiple DataStage servers	18
Connecting to multiple DataStage servers	18
DataStage releases earlier than 7.0	22
Upgrade method: upgrade to DataStage 7.5.2 and then to IBM Information Server.	23
Saving configuration files.	23
Removing DataStage clients	24
Upgrading the server to release 7.5.2 (Windows)	24
Upgrading the server to Release 7.5.2 (UNIX or Linux)	25

Installing IBM Information Server 8	25
Migrating DataStage projects	26
Restoring contents of configuration files	26
Complete method: migrate to a clean installation of DataStage	26
Exporting projects	27
Saving configuration files.	28
Job dependencies	29
Removing previous versions of DataStage	30
Installing IBM Information Server 8	31
Importing DataStage projects	31
Restoring contents of configuration files	32
Restoring job dependencies	32
Recompiling jobs	33

Chapter 3. Migrating QualityStage jobs 35

Job migration in legacy operational mode	35
Job migration in expanded form	36
Rule set migration	36
Match specification migration	36
Installing the WebSphere QualityStage migration utility	36
Running the migration utility	37
Importing the migrated files into the Designer client	39
Preparing migrated jobs for operation	40
WebSphere QualityStage data quality stages	41
Upgrading the QualityStage plug-in DLLs	42

Accessing information about the product 45

Providing comments on the documentation.	45
--	----

Notices 47

Trademarks	49
----------------------	----

Index 51

Chapter 1. Planning your migration

Before you install IBM® Information Server 8, plan whether you need to upgrade components and migrate projects and files.

You might need to upgrade your software in the following circumstances:

- If you want to use an existing database to host the metadata repository for IBM Information server, ensure that you have the correct version, and upgrade as necessary. The versions that are required are as follows:
 - IBM DB2® 9.1. (DB2 is installed with IBM Information Server unless you specify that you want to use your own database.)
 - Oracle 10g
 - Microsoft® SQL Server 2005
- If you want to use an existing DB2 installation to host Federation Server, then you require DB2 9.1.
- If you want to use an existing installation of IBM WebSphere® Application Server, then you require release 6.0.2.11 or later. WebSphere Application Server release 6.0.2.11 is installed with IBM Information Server unless you specify that you want to use your own application server.

DB2 database systems

publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp

WebSphere products

publib.boulder.ibm.com/infocenter/wasinfo/v6r0/index.jsp

Related concepts

Chapter 2, “Migrating from earlier DataStage versions,” on page 3

You have a number of options for upgrading your DataStage™ installation and migrating the associated projects.

Chapter 3, “Migrating QualityStage jobs,” on page 35

If you have QualityStage, there are some migration tasks to complete when upgrading to IBM Information Server.

Installation layers

IBM Information Server comprises layers and product components.

IBM Information Server is installed in five layers: client, metadata repository, domain, engine, and documentation. Product components are installed in each layer, depending on the selections that you make when you run the installer program. The client layer is installed on a Windows® computer. The metadata repository, domain, engine, and documentation layers can be installed on a Windows computer, or a UNIX® or Linux® computer. You can install the metadata repository, domain, and engine layers on the same computer, or on separate computers.

If you install these layers on separate computers, then they must be installed in this order:

1. Metadata repository layer
2. Domain layer
3. Engine layer

Version Control support

The DataStage Version Control component is not supported by IBM Information Server 8.

If you previously used Version Control, you can migrate all of your projects by using the migration procedures, but versioning and promote capabilities are not available for these projects.

Chapter 2. Migrating from earlier DataStage versions

You have a number of options for upgrading your DataStage installation and migrating the associated projects.

The integration of the components that comprise IBM Information Server means that upgrading DataStage is different from previous DataStage upgrades. In particular, note the following points:

- WebSphere DataStage version 8 uses a metadata repository that is shared with other IBM Information Server components.
- An application server facilitates the communication of the different components within IBM Information Server.
- QualityStage is now available as an integrated part of the WebSphere DataStage design framework.
- The Version Control component that was previously available with DataStage is not supported by IBM Information Server 8.

The available options for upgrading DataStage depend on what version of DataStage is currently installed, and whether the DataStage server runs on a Windows computer or a UNIX or Linux computer.

DataStage 7.0 to 7.5.2 - Windows server

There are two available upgrade options:

- You can install IBM Information Server to overwrite the existing DataStage server. The installation program prepares your projects for migration and you migrate them after installation.
- You can save your DataStage design objects and associated data, then remove DataStage before you install IBM Information Server. You then restore your objects and data. You can also use this option to move DataStage to a different computer.

DataStage 7.0 to 7.5.2 - UNIX or Linux server

There are two available upgrade options and a co-existence option:

- You can install IBM Information Server to overwrite the existing DataStage server. The installation program prepares your projects for migration and you migrate them after installation.
- You can save your DataStage design objects and associated data, then remove DataStage before you install IBM Information Server. You then restore your objects and data. You can also use this option to move DataStage to a different computer.
- You can maintain your existing version of DataStage server, and install IBM Information Server in addition.

Earlier DataStage versions - Windows servers

If your version of DataStage is earlier than version 7.0, then you cannot upgrade directly to IBM Information Server. There are two available upgrade options:

- Upgrade your DataStage server to DataStage version 7.5.2 and then run the IBM Information Server installation program to upgrade to WebSphere DataStage.
- You can save your DataStage design objects and associated data, then remove DataStage before installing IBM Information Server. You then restore your objects and data. You can also use this option to move DataStage to a different computer.

Earlier DataStage versions - UNIX and Linux servers

If your version of DataStage is earlier than version 7.0, then you cannot upgrade directly to IBM Information server. There are two available upgrade options and a co-existence option:

- Upgrade your DataStage server to DataStage version 7.5.2 and then run the IBM Information Server installation program to upgrade to WebSphere DataStage.
- You can save your DataStage design objects and associated data, then remove your old installation of DataStage before installing IBM Information Server. You then restore your objects and data. You can also use this option to move DataStage to a different computer.
- You can maintain your existing version of DataStage server, and install IBM Information Server alongside it.

DataStage release 7.0 to 7.5 installations

You can upgrade an existing DataStage 7.0 to 7.5.2 installation to IBM Information Server 8.0.1.

You can upgrade DataStage and migrate projects by using one of these methods:

- To preserve the existing directory structure of your projects, use the basic method of migrating by using the installation program. You do not need to move any of the data files that are associated with your projects.
- To move your projects to the new IBM Information Server directory structure, use the complete method of migrating to a clean installation. With this method, you must move associated project dependencies yourself. You can also use this method to install IBM Information Server on a different computer and move your existing projects to that computer.
- If your DataStage server is on a UNIX or Linux computer, you can retain the version currently on that computer. Use the multiple server method.

Related concepts

“Basic method: Migrate using the installation program” on page 5

The basic method preserves the directory structure of your existing projects, and upgrades the projects to WebSphere DataStage 8.0.1.

“Complete method: migrate to a clean installation of DataStage” on page 7

The complete method moves all of your migrated projects to the WebSphere DataStage 8.0.1 directory structure.

“Multiple server method (UNIX and Linux only)” on page 14

If your DataStage server is on a UNIX or Linux computer, you can keep that instance and install IBM Information Server 8. You can also install multiple WebSphere DataStage 8 servers.

Basic method: Migrate using the installation program

The basic method preserves the directory structure of your existing projects, and upgrades the projects to WebSphere DataStage 8.0.1.

During the installation process, your projects are prepared for migration to the new metadata repository. You then use the WebSphere DataStage and QualityStage Administrator client to migrate each of your projects.

If you use this method, you do not need to move any of your project dependencies. The dependencies remain in the existing location during the migration.

To migrate by using the installation program:

1. Back up DataStage and save settings and configurations files.
2. Remove DataStage clients.
3. Install IBM Information Server, including WebSphere DataStage 8.0.1.
4. Migrate each of the existing projects by using the Administrator client.
5. Integrate the contents of your saved settings and configurations files.

Saving configuration files

Before you upgrade an existing DataStage server, you need to back up your configuration files.

After you install WebSphere DataStage 8.0.1, you can compare each of your saved configuration files with the equivalent new file. You can then integrate the two files. You must not replace the new files with the files that you saved from your existing installation, because you will overwrite new settings.

The files that you copy depend on whether your DataStage server is on a Windows computer, or a UNIX or Linux computer.

Related tasks

“Restoring contents of configuration files” on page 7

If you previously customized any of your DataStage configuration files, you might need to restore these customizations. You do not need to restore the dsparams file.

Saving configuration files (UNIX and Linux)

If you are migrating an existing DataStage server on a UNIX or Linux computer, you must save copies of certain configuration files.

The environment variable \$DSHOME identifies the engine directory of your DataStage installation. By default the engine directory is `/user_home/Ascential/DataStage/DSEngine`, where `user_home` is the home directory of the DataStage administrator who originally installed DataStage (by default, dsadm). Each project is in its own directory in the `/user_home/Ascential/DataStage/Projects` directory.

To save your configuration:

1. Copy the \$DSHOME/dsenv file to a location outside of the DataStage file structure.
2. To keep a record of the existing configuration of your ODBC drivers, copy the following files to a location outside of the DataStage file structure:
 - \$DSHOME/.odbc.ini

- \$DSHOME/uvodbc.config
 - The uvodbc.config file located in each project directory
3. Copy the DSParams file from each project directory to a location outside of the DataStage file structure.

Saving configuration files (Windows)

If you are migrating an existing DataStage server on a Windows computer, you should save copies of certain configuration files.

Each project has its own directory located, by default, in the directory C:/Ascential/DataStage/Projects.

To preserve your configuration:

Copy the DSParams file from each project directory to a safe location.

Removing DataStage clients

You use Windows Add or Remove Programs to remove existing DataStage clients from a Windows computer.

Ensure that nobody is using DataStage and that no clients are connected to the DataStage server.

To remove the DataStage clients:

1. Open the Windows control panel.
2. In the control panel, select **Add or Remove Programs**.
3. In the list of installed programs, select **DataStage Clients**.
4. Click **Change/Remove**.

Installing IBM Information Server 8

When you have saved the required configuration files you can install IBM Information Server.

1. Plan your installation, and then run the installation program to install WebSphere DataStage server and clients, and any other IBM Information Server components that you purchased a license for.
2. Ensure that all of your existing projects are listed by the installer program when the program asks if you want to specify new projects. If the existing projects are not listed, cancel the installation.
3. If NLS is enabled on your previous DataStage installation, enable it when the installation program prompts you.
4. Configure the IBM Information Server and add users.

During the installation, the installation program upgrades your DataStage server and prepares each of your projects for migration.

Migrating DataStage projects

After you have installed IBM Information Server, you must migrate your projects before you can use them.

You require DataStage Administrator privileges to complete this task.

You migrate the projects from the WebSphere DataStage and QualityStage Administrator client.

To migrate the projects:

1. Open the Administrator client.
2. Open the Projects page.
3. In the list of project names, select the first project that you want to migrate. A **Migrate** button is visible when an eligible project is selected.
4. Click **Migrate**. The project is migrated.
5. Repeat steps 3 and 4 for each of the DataStage projects that you need to migrate.

After a project is migrated, it is available to the WebSphere DataStage and QualityStage Designer and Director clients.

Restoring contents of configuration files

If you previously customized any of your DataStage configuration files, you might need to restore these customizations. You do not need to restore the dsparams file.

To restore the files, do not copy your saved versions over the new ones. You must merge the contents of your saved files with the new files. You must only add or edit settings in the new files. Do not delete settings. The procedure described below can be used with any of the configuration files that you saved earlier.

1. Open the saved version of the configuration file in a text editor.
2. Open the new version of the same file in a text editor.
3. Compare the contents of the saved file with the new file.
4. Add any required entries to the new file.
5. Save the new file and close the text editor.

Related tasks

“Saving configuration files” on page 5

Before you upgrade an existing DataStage server, you need to back up your configuration files.

Complete method: migrate to a clean installation of DataStage

The complete method moves all of your migrated projects to the WebSphere DataStage 8.0.1 directory structure.

You can also use this method to migrate your DataStage installation to a different computer.

To migrate by using the complete method

1. Export your existing DataStage projects.
2. Save settings and configurations files.
3. Copy to the export directory any flat files, libraries, schema files, and hashed files that your project uses.
4. Remove your existing DataStage clients and server.
5. Install IBM Information Server, including WebSphere DataStage 8.0.1.
6. Import your projects.
7. Integrate the contents of your saved settings and configuration files.

8. Copy any project dependencies to the appropriate project directories.
9. Recompile the jobs in your projects.

If you have flat files, hashed files, or other dependencies that are located in your DataStage directory structure, or if you are migrating to a different computer, then you must move these files yourself. Such files are not saved as part of the export procedure.

Exporting projects

You export your DataStage projects before removing DataStage.

Ensure that nobody is using DataStage and that all clients are disconnected.

You can export projects using one of the following methods:

- Export from the DataStage Manager client.
- Export from the command line of a DataStage client computer.

Exporting projects from the DataStage Manager client

You can use the DataStage Manager client to export projects.

To export projects using the DataStage Manager client:

1. Open the DataStage Manager client, and attach to project that you want to export.
2. Ensure **View** → **Host View** is selected.
3. Select **Export** → **DataStage Components**.
4. In the Export window, specify details about the project that you want to export:
 - a. In the **Export to file** field, type or browse for the path name of the file in which to store the project. By default, export files have the suffix .dsx.
 - b. In the **Components** tab, select the **Whole Project** option.
5. Click **OK**. The project is exported to the file that you specified.
6. In the left pane of the DataStage Manager window, select the next project and repeat the export process from step 3, specifying a different name for the export file.

Exporting projects from the command line

You can export your projects using the command line on the client computer. You can create a script that runs these commands and exports all of your projects in a single operation.

To export a project from the command line:

1. Open a command line editor and go to the DataStage client directory (the default path is C:\Ascential\Program Files\Ascential\DataStage\version).
2. Enter the command:

```
dscmdexport /H=hostname /U=username /P=password project_name  
export_file_path [/V]
```

The arguments are as follows:

- *hostname* is the name of the DataStage server computer where the project is located.
- *username* is your user name on the server computer.
- *password* is the password for that user name.
- *project_name* is the name of the project that you are exporting.

- *export_file_path* is the path name of the destination file. By convention, export files have the suffix .dsx.
- /V is optional. Including /V turns the verbose option on so that you can follow the progress of the export procedure.

For example, the following command exports the project named monthlyaudit that is located on the DataStage server named R101 and writes the project to a file named monthlyaudit.dsx. The target file is located in a directory named migrated_projects on the client computer.

```
dscmdexport /H=R101 /U=BillG /P=paddock monthlyaudit
C:\migrated_projects\monthlyaudit.dsx
```

You can enter a command for each project that you want to export, or you can create a script that contains commands for all of the projects on your DataStage server computer.

Saving configuration files

Before you upgrade an existing DataStage server, you need to back up your configuration files.

After you install WebSphere DataStage 8.0.1, you can compare each of your saved configuration files with the equivalent new file. You can then integrate the two files. You must not replace the new files with the files that you saved from your existing installation, because you will overwrite new settings.

The files that you copy depend on whether your DataStage server is on a Windows computer, or a UNIX or Linux computer.

Related tasks

“Restoring contents of configuration files” on page 7

If you previously customized any of your DataStage configuration files, you might need to restore these customizations. You do not need to restore the dsparams file.

Saving configuration files (UNIX and Linux)

If you are migrating an existing DataStage server on a UNIX or Linux computer, you must save copies of certain configuration files.

The environment variable \$DSHOME identifies the engine directory of your DataStage installation. By default the engine directory is /user_home/Ascential/DataStage/DSEngine, where *user_home* is the home directory of the DataStage administrator who originally installed DataStage (by default, dsadm). Each project is in its own directory in the /user_home/Ascential/DataStage/Projects directory.

To save your configuration:

1. Copy the \$DSHOME/dsenv file to a location outside of the DataStage file structure.
2. To keep a record of the existing configuration of your ODBC drivers, copy the following files to a location outside of the DataStage file structure:
 - \$DSHOME/.odbc.ini
 - \$DSHOME/uvodbc.config
 - The uvodbc.config file located in each project directory
3. Copy the DSParams file from each project directory to a location outside of the DataStage file structure.

Saving configuration files (Windows)

If you are migrating an existing DataStage server on a Windows computer, you should save copies of certain configuration files.

Each project has its own directory located, by default, in the directory C:/Ascential/DataStage/Projects.

To preserve your configuration:

Copy the DSParams file from each project directory to a safe location.

Job dependencies

The jobs in your projects might depend on objects outside of DataStage. You might need to move these objects to a new location.

The DataStage export procedure exports DataStage objects, such as job designs, table definitions and custom routines. If the jobs in your projects have other dependencies, then you might need to move these dependencies. Dependencies include flat files, files containing schemas used by parallel jobs, and libraries that are called by parallel job custom routines. It does not include parallel job data sets or file sets. If you are using data sets or file sets to permanently store data, then you must design a job to write the data to a flat file or other data target.

You need to move dependencies in the following situations:

- If the files are in the DataStage server or client directory structures.
- If you are installing WebSphere DataStage on a different computer.

You can move flat files, libraries, and schema files by using normal operating system commands. You might need to modify job designs to identify the new locations of these files.

If there are server jobs in your projects that use hashed files, then you must use a dedicated procedure to move those files.

Related tasks

“Restoring job dependencies” on page 13

You must restore any job dependencies that you saved.

“Restoring hashed files” on page 13

You need to restore the hashed files that you saved.

Migrating hashed files

If your DataStage project contains server jobs that use hashed files, then you need to move the hashed files to a new location as part of the migration procedure.

Before moving the files, stop the DataStage server to ensure that dynamic hashed files are copied correctly. If you have previously added indices to your hashed files, then you should remove these indices.

To migrate hashed files:

1. Locate the hashed files in your directory structure:
 - Each static hashed file is represented by two operating system files. For example, a static hashed file named price_lookup is contained in the two files named price_lookup and d_price_lookup.

- Each dynamic hashed file is represented by a directory with the same name as the hashed file and a file named `d_hashed_file_name`. For example, a dynamic hashed file named `code_lookup` is represented by the directory named `code_lookup` and the file named `d_code_lookup`.
2. Copy the files or directories that represent each of your hashed files to a safe location outside of the DataStage directory structure.

You can restore these files after you install and configure WebSphere DataStage.

Related tasks

“Restoring hashed files” on page 13

You need to restore the hashed files that you saved.

Removing previous versions of DataStage

Removing DataStage completely removes your current version of DataStage from server and client computers.

Ensure that your DataStage server computer is backed up before you remove DataStage, and that no one is using DataStage.

There are separate procedures for removing DataStage clients and removing DataStage servers.

Removing a DataStage server from a Windows computer

You use Windows Add or Remove Programs to remove an existing DataStage server from a Windows computer.

Ensure that nobody is using DataStage and that no clients are connected to the server.

To remove the DataStage server:

1. Open the control panel.
2. In the control panel, select **Add or Remove Programs**.
3. In the list of installed programs, select **DataStage server**.
4. Click **Change/Remove**.

Removing a DataStage server from a UNIX or Linux computer

You run an uninstallation program to remove an existing DataStage server from a UNIX or Linux computer.

To run the uninstallation program, you must mount the original installation CD, or copy the CD contents to your computer.

To remove the DataStage server:

1. Log in as the root user.
2. Change directory to the top level directory of the CD, or the directory to which you copied the CD contents.
3. Type the relevant uninstallation command:

Option	Description
Solaris	<code>uninstall.sh</code>
AIX®	<code>uninstall.sh</code>
HP-UX	<code>'UNINSTALL.SH;1'</code>

Option	Description
LINUX	'uninstall.sh'

The uninstallation program guides you through the procedure for removing the DataStage server.

Removing DataStage clients

You use Windows Add or Remove Programs to remove existing DataStage clients from a Windows computer.

Ensure that nobody is using DataStage and that no clients are connected to the DataStage server.

To remove the DataStage clients:

1. Open the Windows control panel.
2. In the control panel, select **Add or Remove Programs**.
3. In the list of installed programs, select **DataStage Clients**.
4. Click **Change/Remove**.

Installing IBM Information Server 8

When you have exported your projects and removed DataStage, you can install IBM Information Server 8.

Plan your installation, and then run the installation program to install WebSphere DataStage server and clients, and any other IBM Information Server components that you purchased a license for. When the installation is complete, configure IBM Information Server and add users.

Importing DataStage projects

After IBM Information Server is installed and configured, you can import the projects that you exported.

To import projects from earlier versions of DataStage:

1. Move the .dsx file or files that you created to the computer where the WebSphere DataStage clients are installed.
2. Open the Administrator client and create a project or projects to contain the objects from your exported projects:
 - a. On the Projects page, click **Add**.
 - b. In the Add Project window, type the name of the project that you want to create and specify a path name for it.
 - c. Click **OK** to create the project.
 - d. Repeat these steps for each project that you want to create. You might want to create a project for each of the projects that you exported and give it the same name as the original project.
3. Open the WebSphere DataStage and QualityStage Designer client and attach to the target project.
4. Select **Import** → **DataStage Components**.
5. Specify the name of the migration file that you want to import, and click **OK**.

Restoring contents of configuration files

If you previously customized any of your DataStage configuration files, you might need to restore these customizations. You do not need to restore the dsparams file.

To restore the files, do not copy your saved versions over the new ones. You must merge the contents of your saved files with the new files. You must only add or edit settings in the new files. Do not delete settings. The procedure described below can be used with any of the configuration files that you saved earlier.

1. Open the saved version of the configuration file in a text editor.
2. Open the new version of the same file in a text editor.
3. Compare the contents of the saved file with the new file.
4. Add any required entries to the new file.
5. Save the new file and close the text editor.

Related tasks

“Saving configuration files” on page 5

Before you upgrade an existing DataStage server, you need to back up your configuration files.

Restoring job dependencies

You must restore any job dependencies that you saved.

To restore job dependency files:

1. For each file, find the location for the file in the new WebSphere DataStage directory structure. For example, if the file was in the project directory in your previous DataStage installation, find the project directory to which you imported the project contents.
2. Use operating system commands to copy the saved file to the required location.
3. Ensure that the jobs that reference these files can locate them. If the file is referred to directly, you must edit the path name in the job design. If the file is referred to by a job parameter, you might need to edit the default value of the parameter to reference the new location for the file.

Related concepts

“Job dependencies” on page 10

The jobs in your projects might depend on objects outside of DataStage. You might need to move these objects to a new location.

Restoring hashed files

You need to restore the hashed files that you saved.

To restore hashed files, use WebSphere DataStage to create empty hashed files for each of the files that you want to restore. You then replace the empty files with the saved files.

To restore hashed files:

If you have a job in your migrated project that creates the hashed files, then follow this procedure:

1. Open the WebSphere DataStage and QualityStage Director client, and attach to the migrated project that uses the hashed files.
2. Run or validate the job that creates the hashed files.
3. Close all WebSphere DataStage clients and stop the DataStage services.

4. Locate the newly created files in your directory structure and copy the saved hashed files over the top of them.
5. Restart the DataStage services.

Your hashed files are now available for use.

If you need to design a job to create the new hashed files, then follow this procedure:

1. Open the Designer client, and attach to the migrated project that uses the hashed files.
2. Create a new server job.
3. For each hashed file that you need to restore, add a Sequential File stage linked to a Hashed File stage.
4. In the Sequential File stage, point to an empty text file and define one column in the Outputs page Columns tab.
5. In the Hashed File stage, specify the name of the hashed file that you are restoring, then select the **Create File** option and specify the file type of the hashed file that you are restoring.
6. When you have added stages for all the hashed files that you want to restore, compile and then validate or run the job to create the empty hashed files.
7. Close all WebSphere DataStage clients and stop the DataStage services.
8. Locate the newly-created files in your directory structure and copy your saved hash files over the top of them.
9. Restart the DataStage services.

Your hashed files are now available for use.

Recompiling jobs

When you move your projects, you must recompile the jobs and routines in these projects to create new executable jobs.

You use the Multiple Job Compile tool to recompile all of the jobs and routines in your projects.

To use the Multiple Job Compile tool:

1. Open the Designer client, and attach to a restored project.
2. Select **Tools** → **Multiple Job Compile**.
3. Follow the steps in the DataStage Compilation wizard to compile all of the jobs and routines in your project.

Multiple server method (UNIX and Linux only)

If your DataStage server is on a UNIX or Linux computer, you can keep that instance and install IBM Information Server 8. You can also install multiple WebSphere DataStage 8 servers.

If you used the multi-server facility to maintain different releases of DataStage server, or multiple instances of the same server, you can keep all of those instances. You can also upgrade each of these servers to WebSphere DataStage 8.0.1 and keep their existing identity information. You cannot install WebSphere DataStage 8.0 on

the same computer as WebSphere DataStage 8.0.1. If you are using WebSphere DataStage 8.0, you must upgrade to IBM Information server version 8.0.1 by using the complete method.


When you use the multi-server facility, each DataStage server is identified by a tag ID (an ITAG) and allocated a different port. You also specify a different home directory for each server. Each DataStage server has a different set of projects that you access through that server. You connect to different servers by specifying the port number in addition to the UNIX server ID.

The existing DataStage server instance retains the default port number, 31538, and uses the default ITAG, ADE. If you already have multiple server instances, then you already have servers using different ports and ITAGs.

Each version of DataStage requires clients of the same version. You must use the Multi-Client Manager program to switch between different client versions on your Windows client computers. You cannot have WebSphere DataStage 8.0 and 8.0.1 clients installed on the same computer.

You can use the multi-server facility to install multiple instances of the WebSphere DataStage 8.0.1 server. Each server requires a different instance of the WebSphere Application Server installed on a different computer, and each server uses a different metadata repository.

The following diagrams show three examples of multi-server installations, and illustrate the use of port numbers and ITAGs.



Scenario 2:
Two existing DataStage servers

DataStage 7.2 server

Itag ADE
Port 31538
/disk1/Ascential/DataStage

DataStage 7.5.1 server

Itag A23
Port 31546
/disk2/Ascential/DataStage

WebSphere DataStage 8.0.1 server

Itag 123
Port 31540
/opt/IBM/InformationServer/Server

Scenario 3:
Multiple WebSphere DataStage 8.0.1 servers

WebSphere DataStage 8.0.1 server

Itag ADE
Port 31538
/opt/IBM/InformationServer/Server

WebSphere DataStage 8.0.1 server

Itag BED
Port 31540
/opt2/IBM/InformationServer/Server

WebSphere DataStage 8.0.1 server

Itag 123
Port 31542
/opt3/IBM/InformationServer/Server

Installing IBM Information Server in addition to existing DataStage servers

You can maintain existing DataStage server instances while installing IBM Information Server.

You need the IBM Information Server DVD for UNIX and Linux systems and the *IBM Information Server Planning, Installation, and Configuration Guide*, which is in the

top-level directory of the DVD. Before you start the installation, plan where to install each of the IBM Information Server components.

The procedure uses the graphical version of the installation program on an XTERM console.

To install IBM Information Server in addition to existing DataStage servers

1. Log on to the system as the root user.
2. Turn off any firewall software that is installed on the system.
3. Mount the IBM Information Server DVD and copy the contents of the DVD to the computer.
4. In the directory where you copied the contents of the DVD, double-click **install** or run the following command from a shell prompt: `./install`. The installation program starts and guides you through the installation procedure.
5. When asked to specify an installation directory, select the **New installation** option and either use the default destination, or type a new one. The destination directory must be different from the installation directory that the existing DataStage server uses.
6. When asked to select IBM Information Server layers, select the Engine layer (select other layers as required by your installation plan).
7. When asked to select product modules and components, select WebSphere DataStage and QualityStage (select other modules as required by your installation plan).
8. When asked to specify a DataStage Server installation directory, select **New Installation**. The installation program detects an existing DataStage server installation and prompts you for an ITAG value and a port number to use for the new WebSphere DataStage server.
9. Enter a value for the ITAG. The ITAG value is a three-digit hexadecimal number. ACE, ADE, and DCD are reserved values. The ITAG value must be unique.
10. Enter a port number for the server. The port number 31538 and any other port numbers that existing servers use are reserved.
11. When the installation is finished, re-enable your firewall software.

Installing WebSphere DataStage 8.0 clients in addition to existing DataStage clients

When you maintain multiple versions of the DataStage server on UNIX or Linux computers, you must maintain corresponding clients on Windows computers.

You need the IBM Information Server DVD (Clients) and the *IBM Information Server Planning, Installation, and Configuration Guide*, which is in the top-level directory of the DVD).

You can install clients for WebSphere DataStage 8.0.1 on the same Windows computer as your existing DataStage clients, and use the Multi-client Manager to switch between the different versions. You can also install WebSphere DataStage 8.0.1 clients on a separate Windows computer. If you are installing on a computer with existing clients, the client installation program detects their presence and installs the Multi-Client Manager program.

To install WebSphere DataStage 8.0.1 clients:

1. Log on to the Windows computer as an administrator.

2. Turn off any firewall software that is installed on the computer.
3. Optional: Turn off your antivirus software.
4. Insert the IBM Information Server DVD or CD, or go to the root directory of the downloaded installation image.
5. In the root directory of the CD, DVD, or downloaded installation image, double-click **install.exe**. The installation program starts and guides you through the installation procedure.
6. When asked for an installation directory, select the **New installation** option and either use the default, or specify a new directory.
7. When asked to select the product modules and components, select WebSphere DataStage and QualityStage. Select other components as required by your installation plan.
8. When asked, select the **Typical** installation type.

Starting and stopping multiple DataStage servers

When multiple DataStage servers are on a single computer, you must identify the server that you want to start or stop.

You identify the DataStage server that you want to start or stop by setting the \$DSHOME environment variable to reference the DSEngine subdirectory of the server. For example, there are two servers on the same computer with the following DSEngine subdirectories:

- /disk1/Ascential/DataStage/DSEngine (existing 7.5.1 server)
- /opt/IBM/InformationServer/Server/DSEngine (WebSphere DataStage 8.0 server)

To stop and restart the WebSphere DataStage 8.0.1 server:

1. Set the \$DSHOME environment variable to point to the /opt/IBM/InformationServer/Server/DSEngine directory.
2. Stop the server by using the following command:
\$DSHOME/bin/uv -admin -stop
3. Wait thirty seconds to give the server time to stop.
4. Start the server by using the following command:
\$DSHOME/bin/uv -admin -start

Connecting to multiple DataStage servers

To connect to a DataStage server from a client when you have multiple servers, you must specify the server port number in addition to the name of the computer where the server is located.

You use a different procedure depending on the version of DataStage server and the client, and whether you are connecting from the command line. The possible scenarios are as follows:

- Connecting to DataStage version 7.5.2 or earlier from the Designer client, Manager client, or Director client.
- Connecting to DataStage version 7.5.2 or earlier from the Administrator client.
- Connecting to WebSphere DataStage version 8.0.1 from the Designer client or Director client.
- Connecting to WebSphere DataStage version 8.0.1 from the Administrator client.
- Connecting to any version of DataStage server from the command line.

Use the Multi-Client Manager to switch between different versions of the clients. The client version must match the server version.

Connecting to DataStage version 7.5.2 or earlier


Use this procedure to connect to a DataStage version 7.5.2 or earlier server from any client other than the Administrator client.

If multiple versions of the DataStage clients are on your computer, ensure that the same version as the server version is currently selected in the Multi-Client Manager.

To connect to DataStage version 7.5.2 or earlier:

1. Select the Manager, Designer, or Director client from the **Start** menu.
2. In the **Host System** field of the Attach to Project window, type the identity of the server computer in the form *hostname:portnumber* where *portnumber* is the port number that the server uses.
3. Type your user name and password.
4. Specify the name of the project that you want to attach to.

For example, here is the Attach to Project window used to attach to a DataStage 7.5.2 server that uses the default port:



The screenshot shows a Windows-style dialog box titled "Attach to Project". It has a blue title bar with a close button (X) in the top right corner. The dialog contains several input fields and buttons. On the right side, there are three buttons: "OK", "Cancel", and "Help". The "Host system:" label is above a dropdown menu that currently displays "R101:31538". Below this, there is a group box containing the "User name:" label above a text field with "BILLG", a "Password:" label above a text field with "xxxxxxx", and an "Omit" checkbox which is currently unchecked. At the bottom of the dialog, the "Project:" label is above a dropdown menu that displays "testproj".

Connecting to DataStage version 7.5.2 or earlier from the Administrator client

Use this procedure to connect to a DataStage version 7.5.2 or earlier server from the Administrator client.


If multiple versions of the DataStage clients are on your computer, ensure that the same version as the server version is currently selected in the Multi-Client Manager.

To connect to DataStage version 7.5.2 or earlier from the Administrator client:

1. Select the Administrator client from the **Start** menu.

2. In the **Host System** field of the Attach to DataStage window, type the identity of the server computer in the form *hostname:portnumber* where *portnumber* is the port number that the server uses.
3. Type your user name and password.

For example, here is the Attach to DataStage window used to attach to a DataStage 7.5.2 server that uses the default port:



Connecting to WebSphere DataStage version 8.0.1


Use this procedure to connect to a WebSphere DataStage 8.0.1 server from any client other than the Administrator client.

If multiple versions of the DataStage clients are on your computer, ensure that the same version as the server version is currently selected in the Multi-Client Manager.

To connect to WebSphere DataStage version 8.0.1:

1. Select the Designer client or Director client from the **Start** menu.
2. In the **Domain** field of the *Attach to Project* window, type the name of the domain to which your WebSphere DataStage server belongs in the form *Domainserver:9080*.
3. Type your user name and password.
4. In the **Project** field, specify the identity of the project that you want to attach to in the form *hostname:portnumber/project* where *portnumber* is the port number that the server uses.

For example, here is the Attach to Project window used to attach to a project called *datastage* on a WebSphere DataStage 8.0.1 server that uses the port 31548:



Connecting to WebSphere DataStage version 8.0.1 from an Administrator client


Use this procedure to connect to a WebSphere DataStage 8.0.1 server from the Administrator client.

If multiple versions of the DataStage clients are on your computer, ensure that the same version as the server version is currently selected in the Multi-Client Manager.

To connect to WebSphere DataStage version 8.0.1 from an Administrator client.

1. Select the Administrator client from the **Start** menu.
2. In the **Domain** field of the Attach to DataStage window, type the name of the domain to which your WebSphere DataStage server belongs in the form *Domainserver:9080*.
3. Type your user name and password.
4. In the **DataStage Server** field, specify the identity of the computer where the server that you want to attach to is located in the form *hostname:portnumber* where *portnumber* is the port number that the server uses.

For example, here is the Attach to DataStage window used to attach to a WebSphere DataStage 8.0.1 server on a computer called R101 that uses the port 31548:



Connecting to a DataStage server from the command line

You can use the command line interface to control DataStage jobs in different projects.

To connect to a project from the command line, you specify the server name and port number of the required instance using the `-server` argument in the form `-server server:portnumber` for local computers.

For remote computers that run WebSphere DataStage 8.0.1 server instances, you must specify the domain and the server name in the form `-domain domain:domain_portnumber -server server:portnumber`. The default domain port number is 9080.

For example, to run a job on the local computer on the server that uses port 31359, run this command:

```
dsjob -server r101:31359 -run myproj myjob
```

To run a job on the local computer on the default server, run this command:

```
dsjob -run myotherproj myotherjob
```

To run a job on a remote computer called R101 on the WebSphere DataStage 8.0.1 server that uses port 31360, you must also specify the host computer, the domain, and supply login information:

```
dsjob -domain mydomain:9080 -server r101:31360 -user billg -password  
paddock -run myproj myjob
```

DataStage releases earlier than 7.0

This topic describes the upgrade options for versions of DataStage earlier than Release 7.0.

If you have a version of DataStage that is earlier than Release 7.0, then you cannot directly migrate your projects to IBM Information Server 8.0.1.

To preserve the existing directory structure of your projects, use the upgrade method of upgrading to DataStage 7.5.2 and then upgrading to IBM Information Server. Your projects are migrated as part of the upgrade to IBM Information Server. You do not need to move any of the data files that are associated with your projects.

To move your projects to the new IBM Information Server directory structure, use the complete method of migrating to a clean installation. With this method, you must move associated project dependencies yourself. You can also use this method to install IBM Information Server on a different computer and move your existing projects to that computer.

Related concepts

“Complete method: migrate to a clean installation of DataStage” on page 7
The complete method moves all of your migrated projects to the WebSphere DataStage 8.0.1 directory structure.

Related tasks

“Upgrade method: upgrade to DataStage 7.5.2 and then to IBM Information Server”

This method preserves your existing DataStage project directory structure.

Upgrade method: upgrade to DataStage 7.5.2 and then to IBM Information Server

This method preserves your existing DataStage project directory structure.

To upgrade using the upgrade method:

1. Save settings and configuration files.
2. Uninstall your existing DataStage clients.
3. Upgrade your DataStage server to version 7.5.2.
4. Install IBM Information Server.
5. Migrate DataStage projects.
6. Restore configuration files.

Related concepts

“DataStage releases earlier than 7.0” on page 22

This topic describes the upgrade options for versions of DataStage earlier than Release 7.0.

Saving configuration files

Before you upgrade an existing DataStage server, you need to back up your configuration files.

After you install WebSphere DataStage 8.0.1, you can compare each of your saved configuration files with the equivalent new file. You can then integrate the two files. You must not replace the new files with the files that you saved from your existing installation, because you will overwrite new settings.

The files that you copy depend on whether your DataStage server is on a Windows computer, or a UNIX or Linux computer.

Related tasks

“Restoring contents of configuration files” on page 7

If you previously customized any of your DataStage configuration files, you might need to restore these customizations. You do not need to restore the dsparams file.

Saving configuration files (UNIX and Linux)

If you are migrating an existing DataStage server on a UNIX or Linux computer, you must save copies of certain configuration files.

The environment variable \$DSHOME identifies the engine directory of your DataStage installation. By default the engine directory is `/user_home/Ascential/DataStage/DSEngine`, where `user_home` is the home directory of the DataStage administrator who originally installed DataStage (by default, dsadm). Each project is in its own directory in the `/user_home/Ascential/DataStage/Projects` directory.

To save your configuration:

1. Copy the \$DSHOME/dsenv file to a location outside of the DataStage file structure.
2. To keep a record of the existing configuration of your ODBC drivers, copy the following files to a location outside of the DataStage file structure:
 - \$DSHOME/.odbc.ini
 - \$DSHOME/uvodbc.config
 - The uvodbc.config file located in each project directory
3. Copy the DSParams file from each project directory to a location outside of the DataStage file structure.

Saving configuration files (Windows)

If you are migrating an existing DataStage server on a Windows computer, you should save copies of certain configuration files.

Each project has its own directory located, by default, in the directory `C:/Ascential/DataStage/Projects`.

To preserve your configuration:

Copy the DSParams file from each project directory to a safe location.

Removing DataStage clients

You use Windows Add or Remove Programs to remove existing DataStage clients from a Windows computer.

Ensure that nobody is using DataStage and that no clients are connected to the DataStage server.

To remove the DataStage clients:

1. Open the Windows control panel.
2. In the control panel, select **Add or Remove Programs**.
3. In the list of installed programs, select **DataStage Clients**.
4. Click **Change/Remove**.

Upgrading the server to release 7.5.2 (Windows)

Use this procedure if your server is on a Windows computer.

Ensure that nobody is using DataStage and that no clients are connected. You must back up your system before upgrading.

Do not add any new packages or plug-ins during the upgrade procedure, you can make any required additions during the IBM Information Server installation procedure.

To upgrade your DataStage server to DataStage 7.5.2:

1. Insert the DataStage 7.5.2 CD into the CD drive.
2. If the installation program window does not appear automatically, double-click **install.exe** in the top-level directory of the CD.
3. In the installation program window, select **Install Server Software**.

Upgrading the server to Release 7.5.2 (UNIX or Linux)

Use this procedure if your DataStage server is on a UNIX or Linux computer.

Ensure that no one is using DataStage and that no clients are connected. You must use the sh or ksh shell to upgrade. Back up your system before you upgrade.

Do not add any new packages or plug-ins during the upgrade procedure, you can make any required additions during the IBM Information Server installation procedure.

To upgrade your DataStage server to release 7.5.2:

1. Log in as the root user.
2. Mount the DataStage 7.5.2 CD.
3. Go to the directory where the CD is mounted and run the install.sh program in that directory. Type the relevant installation command:

Option	Description
Solaris	install.sh
AIX	install.sh
HP-UX	'INSTALL.SH;1'
LINUX	'install.sh'

4. Supply details to the installation program as requested. Do not add any new packages or plug-ins during the upgrade procedure, you can make any required additions during the IBM Information Server installation procedure.

Installing IBM Information Server 8

When you have saved the required configuration files you can install IBM Information Server.

1. Plan your installation, and then run the installation program to install WebSphere DataStage server and clients, and any other IBM Information Server components that you purchased a license for.
2. Ensure that all of your existing projects are listed by the installer program when the program asks if you want to specify new projects. If the existing projects are not listed, cancel the installation.
3. If NLS is enabled on your previous DataStage installation, enable it when the installation program prompts you.
4. Configure the IBM Information Server and add users.

During the installation, the installation program upgrades your DataStage server and prepares each of your projects for migration.

Migrating DataStage projects

After you have installed IBM Information Server, you must migrate your projects before you can use them.

You require DataStage Administrator privileges to complete this task.

You migrate the projects from the WebSphere DataStage and QualityStage Administrator client.

To migrate the projects:

1. Open the Administrator client.
2. Open the Projects page.
3. In the list of project names, select the first project that you want to migrate. A **Migrate** button is visible when an eligible project is selected.
4. Click **Migrate**. The project is migrated.
5. Repeat steps 3 and 4 for each of the DataStage projects that you need to migrate.

After a project is migrated, it is available to the WebSphere DataStage and QualityStage Designer and Director clients.

Restoring contents of configuration files

If you previously customized any of your DataStage configuration files, you might need to restore these customizations. You do not need to restore the dsparams file.

To restore the files, do not copy your saved versions over the new ones. You must merge the contents of your saved files with the new files. You must only add or edit settings in the new files. Do not delete settings. The procedure described below can be used with any of the configuration files that you saved earlier.

1. Open the saved version of the configuration file in a text editor.
2. Open the new version of the same file in a text editor.
3. Compare the contents of the saved file with the new file.
4. Add any required entries to the new file.
5. Save the new file and close the text editor.

Related tasks

“Saving configuration files” on page 5

Before you upgrade an existing DataStage server, you need to back up your configuration files.

Complete method: migrate to a clean installation of DataStage

The complete method moves all of your migrated projects to the WebSphere DataStage 8.0.1 directory structure.

You can also use this method to migrate your DataStage installation to a different computer.

To migrate by using the complete method

1. Export your existing DataStage projects.

2. Save settings and configurations files.
3. Copy to the export directory any flat files, libraries, schema files, and hashed files that your project uses.
4. Remove your existing DataStage clients and server.
5. Install IBM Information Server, including WebSphere DataStage 8.0.1.
6. Import your projects.
7. Integrate the contents of your saved settings and configuration files.
8. Copy any project dependencies to the appropriate project directories.
9. Recompile the jobs in your projects.

If you have flat files, hashed files, or other dependencies that are located in your DataStage directory structure, or if you are migrating to a different computer, then you must move these files yourself. Such files are not saved as part of the export procedure.

Exporting projects

You export your DataStage projects before removing DataStage.

Ensure that nobody is using DataStage and that all clients are disconnected.

You can export projects using one of the following methods:

- Export from the DataStage Manager client.
- Export from the command line of a DataStage client computer.

Exporting projects from the DataStage Manager client

You can use the DataStage Manager client to export projects.

To export projects using the DataStage Manager client:

1. Open the DataStage Manager client, and attach to project that you want to export.
2. Ensure **View** → **Host View** is selected.
3. Select **Export** → **DataStage Components**.
4. In the Export window, specify details about the project that you want to export:
 - a. In the **Export to file** field, type or browse for the path name of the file in which to store the project. By default, export files have the suffix **.dsx**.
 - b. In the **Components** tab, select the **Whole Project** option.
5. Click **OK**. The project is exported to the file that you specified.
6. In the left pane of the DataStage Manager window, select the next project and repeat the export process from step 3, specifying a different name for the export file.

Exporting projects from the command line

You can export your projects using the command line on the client computer. You can create a script that runs these commands and exports all of your projects in a single operation.

To export a project from the command line:

1. Open a command line editor and go to the DataStage client directory (the default path is C:\Ascential\Program Files\Ascential\DataStage`version`).
2. Enter the command:

```
dscmdexport /H=hostname /U=username /P=password project_name  
export_file_path [/V]
```

The arguments are as follows:

- *hostname* is the name of the DataStage server computer where the project is located.
- *username* is your user name on the server computer.
- *password* is the password for that user name.
- *project_name* is the name of the project that you are exporting.
- *export_file_path* is the path name of the destination file. By convention, export files have the suffix .dsx.
- */V* is optional. Including */V* turns the verbose option on so that you can follow the progress of the export procedure.

For example, the following command exports the project named *monthlyaudit* that is located on the DataStage server named *R101* and writes the project to a file named *monthlyaudit.dsx*. The target file is located in a directory named *migrated_projects* on the client computer.

```
dscmdexport /H=R101 /U=BillG /P=paddock monthlyaudit  
C:\migrated_projects\monthlyaudit.dsx
```

You can enter a command for each project that you want to export, or you can create a script that contains commands for all of the projects on your DataStage server computer.

Saving configuration files

Before you upgrade an existing DataStage server, you need to back up your configuration files.

After you install WebSphere DataStage 8.0.1, you can compare each of your saved configuration files with the equivalent new file. You can then integrate the two files. You must not replace the new files with the files that you saved from your existing installation, because you will overwrite new settings.

The files that you copy depend on whether your DataStage server is on a Windows computer, or a UNIX or Linux computer.

Related tasks

“Restoring contents of configuration files” on page 7

If you previously customized any of your DataStage configuration files, you might need to restore these customizations. You do not need to restore the *dsparams* file.

Saving configuration files (UNIX and Linux)

If you are migrating an existing DataStage server on a UNIX or Linux computer, you must save copies of certain configuration files.

The environment variable *\$DSHOME* identifies the engine directory of your DataStage installation. By default the engine directory is */user_home/Ascential/DataStage/DSEngine*, where *user_home* is the home directory of the DataStage administrator who originally installed DataStage (by default, *dsadm*). Each project is in its own directory in the */user_home/Ascential/DataStage/Projects* directory.

To save your configuration:

1. Copy the \$DSHOME/dsenv file to a location outside of the DataStage file structure.
2. To keep a record of the existing configuration of your ODBC drivers, copy the following files to a location outside of the DataStage file structure:
 - \$DSHOME/.odbc.ini
 - \$DSHOME/uvodbc.config
 - The uvodbc.config file located in each project directory
3. Copy the DSParams file from each project directory to a location outside of the DataStage file structure.

Saving configuration files (Windows)

If you are migrating an existing DataStage server on a Windows computer, you should save copies of certain configuration files.

Each project has its own directory located, by default, in the directory C:/Ascential/DataStage/Projects.

To preserve your configuration:

Copy the DSParams file from each project directory to a safe location.

Job dependencies

The jobs in your projects might depend on objects outside of DataStage. You might need to move these objects to a new location.

The DataStage export procedure exports DataStage objects, such as job designs, table definitions and custom routines. If the jobs in your projects have other dependencies, then you might need to move these dependencies. Dependencies include flat files, files containing schemas used by parallel jobs, and libraries that are called by parallel job custom routines. It does not include parallel job data sets or file sets. If you are using data sets or file sets to permanently store data, then you must design a job to write the data to a flat file or other data target.

You need to move dependencies in the following situations:

- If the files are in the DataStage server or client directory structures.
- If you are installing WebSphere DataStage on a different computer.

You can move flat files, libraries, and schema files by using normal operating system commands. You might need to modify job designs to identify the new locations of these files.

If there are server jobs in your projects that use hashed files, then you must use a dedicated procedure to move those files.

Related tasks

“Restoring job dependencies” on page 13

You must restore any job dependencies that you saved.

“Restoring hashed files” on page 13

You need to restore the hashed files that you saved.

Migrating hashed files

If your DataStage project contains server jobs that use hashed files, then you need to move the hashed files to a new location as part of the migration procedure.

Before moving the files, stop the DataStage server to ensure that dynamic hashed files are copied correctly. If you have previously added indices to your hashed files, then you should remove these indices.

To migrate hashed files:

1. Locate the hashed files in your directory structure:
 - Each static hashed file is represented by two operating system files. For example, a static hashed file named `price_lookup` is contained in the two files named `price_lookup` and `d_price_lookup`.
 - Each dynamic hashed file is represented by a directory with the same name as the hashed file and a file named `d_hashed_file_name`. For example, a dynamic hashed file named `code_lookup` is represented by the directory named `code_lookup` and the file named `d_code_lookup`.
2. Copy the files or directories that represent each of your hashed files to a safe location outside of the DataStage directory structure.

You can restore these files after you install and configure WebSphere DataStage.

Related tasks

“Restoring hashed files” on page 13

You need to restore the hashed files that you saved.

Removing previous versions of DataStage

Removing DataStage completely removes your current version of DataStage from server and client computers.

Ensure that your DataStage server computer is backed up before you remove DataStage, and that no one is using DataStage.

There are separate procedures for removing DataStage clients and removing DataStage servers.

Removing a DataStage server from a Windows computer

You use Windows Add or Remove Programs to remove an existing DataStage server from a Windows computer.

Ensure that nobody is using DataStage and that no clients are connected to the server.

To remove the DataStage server:

1. Open the control panel.
2. In the control panel, select **Add or Remove Programs**.
3. In the list of installed programs, select **DataStage server**.
4. Click **Change/Remove**.

Removing a DataStage server from a UNIX or Linux computer

You run an uninstallation program to remove an existing DataStage server from a UNIX or Linux computer.

To run the uninstallation program, you must mount the original installation CD, or copy the CD contents to your computer.

To remove the DataStage server:

1. Log in as the root user.

2. Change directory to the top level directory of the CD, or the directory to which you copied the CD contents.
3. Type the relevant uninstallation command:

Option	Description
Solaris	uninstall.sh
AIX	uninstall.sh
HP-UX	'UNINSTLL.SH;1'
LINUX	'uninstall.sh'

The uninstallation program guides you through the procedure for removing the DataStage server.

Removing DataStage clients

You use Windows Add or Remove Programs to remove existing DataStage clients from a Windows computer.

Ensure that nobody is using DataStage and that no clients are connected to the DataStage server.

To remove the DataStage clients:

1. Open the Windows control panel.
2. In the control panel, select **Add or Remove Programs**.
3. In the list of installed programs, select **DataStage Clients**.
4. Click **Change/Remove**.

Installing IBM Information Server 8

When you have exported your projects and removed DataStage, you can install IBM Information Server 8.

Plan your installation, and then run the installation program to install WebSphere DataStage server and clients, and any other IBM Information Server components that you purchased a license for. When the installation is complete, configure IBM Information Server and add users.

Importing DataStage projects

After IBM Information Server is installed and configured, you can import the projects that you exported.

To import projects from earlier versions of DataStage:

1. Move the .dsx file or files that you created to the computer where the WebSphere DataStage clients are installed.
2. Open the Administrator client and create a project or projects to contain the objects from your exported projects:
 - a. On the Projects page, click **Add**.
 - b. In the Add Project window, type the name of the project that you want to create and specify a path name for it.
 - c. Click **OK** to create the project.
 - d. Repeat these steps for each project that you want to create. You might want to create a project for each of the projects that you exported and give it the same name as the original project.

3. Open the WebSphere DataStage and QualityStage Designer client and attach to the target project.
4. Select **Import** → **DataStage Components**.
5. Specify the name of the migration file that you want to import, and click **OK**.

Restoring contents of configuration files

If you previously customized any of your DataStage configuration files, you might need to restore these customizations. You do not need to restore the dsparams file.

To restore the files, do not copy your saved versions over the new ones. You must merge the contents of your saved files with the new files. You must only add or edit settings in the new files. Do not delete settings. The procedure described below can be used with any of the configuration files that you saved earlier.

1. Open the saved version of the configuration file in a text editor.
2. Open the new version of the same file in a text editor.
3. Compare the contents of the saved file with the new file.
4. Add any required entries to the new file.
5. Save the new file and close the text editor.

Related tasks

“Saving configuration files” on page 5

Before you upgrade an existing DataStage server, you need to back up your configuration files.

Restoring job dependencies

You must restore any job dependencies that you saved.

To restore job dependency files:

1. For each file, find the location for the file in the new WebSphere DataStage directory structure. For example, if the file was in the project directory in your previous DataStage installation, find the project directory to which you imported the project contents.
2. Use operating system commands to copy the saved file to the required location.
3. Ensure that the jobs that reference these files can locate them. If the file is referred to directly, you must edit the path name in the job design. If the file is referred to by a job parameter, you might need to edit the default value of the parameter to reference the new location for the file.

Related concepts

“Job dependencies” on page 10

The jobs in your projects might depend on objects outside of DataStage. You might need to move these objects to a new location.

Restoring hashed files

You need to restore the hashed files that you saved.

To restore hashed files, use WebSphere DataStage to create empty hashed files for each of the files that you want to restore. You then replace the empty files with the saved files.

To restore hashed files:

If you have a job in your migrated project that creates the hashed files, then follow this procedure:

1. Open the WebSphere DataStage and QualityStage Director client, and attach to the migrated project that uses the hashed files.
2. Run or validate the job that creates the hashed files.
3. Close all WebSphere DataStage clients and stop the DataStage services.
4. Locate the newly created files in your directory structure and copy the saved hashed files over the top of them.
5. Restart the DataStage services.

Your hashed files are now available for use.

If you need to design a job to create the new hashed files, then follow this procedure:

1. Open the Designer client, and attach to the migrated project that uses the hashed files.
2. Create a new server job.
3. For each hashed file that you need to restore, add a Sequential File stage linked to a Hashed File stage.
4. In the Sequential File stage, point to an empty text file and define one column in the Outputs page Columns tab.
5. In the Hashed File stage, specify the name of the hashed file that you are restoring, then select the **Create File** option and specify the file type of the hashed file that you are restoring.
6. When you have added stages for all the hashed files that you want to restore, compile and then validate or run the job to create the empty hashed files.
7. Close all WebSphere DataStage clients and stop the DataStage services.
8. Locate the newly-created files in your directory structure and copy your saved hash files over the top of them.
9. Restart the DataStage services.

Your hashed files are now available for use.

Recompiling jobs

When you move your projects, you must recompile the jobs and routines in these projects to create new executable jobs.

You use the Multiple Job Compile tool to recompile all of the jobs and routines in your projects.

To use the Multiple Job Compile tool:

1. Open the Designer client, and attach to a restored project.
2. Select **Tools** → **Multiple Job Compile**.
3. Follow the steps in the DataStage Compilation wizard to compile all of the jobs and routines in your project.

Chapter 3. Migrating QualityStage jobs

If you have QualityStage, there are some migration tasks to complete when upgrading to IBM Information Server.

If QualityStage 7.0 to 7.5 is installed, you need to migrate your jobs to IBM Information Server.

If you were using the QualityStage plug-in stage in DataStage jobs, the jobs are migrated by the installation process. To run these jobs after migration, however, you must configure the QualityStage plug-in to work with WebSphere DataStage 8.0.1.

IBM WebSphere QualityStage provides a migration utility to migrate QualityStage 7.0 to 7.5 jobs, match specifications, and standardization rule sets to the WebSphere DataStage and QualityStage Designer client environment.

The migration utility uses information in the QualityStage 7.0 to 7.5 server project directory to build the .dsx file format that WebSphere DataStage requires to import jobs.

The migration utility migrates the following QualityStage 7.0 to 7.5 objects:

- Standardization rule set
- Jobs in full legacy operational mode
- Jobs in expanded form, in which some legacy operations are replaced by QualityStage 8.0 stages
- Match specification

The migration utility creates a .dsx file that contains migrated jobs, rule sets, and match specifications in the Temp directory in the QualityStage 7.0 to 7.5 project directory.

Job migration in legacy operational mode

When you migrate QualityStage 7.0 to 7.5 jobs in full legacy operational mode, the original job is replaced with a single instance of a QualityStage Legacy stage and Sequential File stages linked as source and target stages.

If you migrate a QualityStage 7.0 - 7.5 job in legacy operational mode, you can make only minor changes to the resulting QualityStage Legacy stage. Use this option only for stable jobs that were never modified or jobs that are due to be replaced.

Do not use this option if you are migrating a job that contains the following QualityStage 7.0 to 7.5 stages because their functionality is not supported by the QualityStage Legacy stage:

- Postal stages such as CASS and SERP
- Program stage
- Multinational Standardize stage
- WAVES stage
- Format Convert stage

Job migration in expanded form

When you migrate a QualityStage 7.0 to 7.5 job in expanded form, the utility replaces the original job with one or more parallel job stages for each stage in the 7.0 to 7.5 job, and with Sequential File stages that are linked together to represent the 7.0 to 7.5 job flow.

A job migrated in expanded form might have a mixture of data quality, processing, and Legacy stages when you open the job in the Designer client. For complex jobs, you can move the stages around on the canvas to make the job more intelligible. You can also replace a QualityStage Legacy stage with a data quality stage that has equivalent functions.

Rule set migration

The migration of a QualityStage 7.0 to 7.5 standardization rule set option migrates the rule sets explicitly by name.

When a QualityStage 7.0 to 7.5 job is migrated, the migration utility detects the dependent rule sets and displays them. You can then choose to include the rule sets in the .dsx file with the job.

The migration utility renames the rule sets within the .dsx file to prevent a naming duplication with a built-in WebSphere QualityStage 8.0 rule set. The utility uses the following naming convention:

QS-7.0to7.5-Ruleset-Name_QS-7.0to7.5-Project-Name

Match specification migration

When you migrate a job containing match and unduplicate stages, the Match specifications are migrated with the job.

The migration utility includes the match processing information in the .dsx file with the job, if you selected the option "plus dependencies" to migrate the job. Once the job is imported, you can locate the match specification in the DataStage **Repository** → **Match Specifications** folder.

As with rule sets, match specifications are renamed when the information is imported. The match specification name has the following form:

QS-7.0 to 7.5-Match-or-Undup-Stage-Name_QS7.0 to 7.5-Project-Name

Installing the WebSphere QualityStage migration utility

Use the migration utility to migrate QualityStage Version 7.0 or later to WebSphere QualityStage Version 8.0.1. The migration utility is installed when you install IBM Information Server, or you can download it and install it separately.

To install the WebSphere QualityStage migration utility:

1. Choose one of the following options, depending on Operating System type and whether you want to install the WebSphere DataStage and QualityStage server on the same computer as the QualityStage 7.0 to 7.5 server.

Option	Description
Installing the WebSphere DataStage and QualityStage server on the same computer as the QualityStage 7.0 or 7.5 server	Install IBM Information Server, installing the engine layer on the computer where the QualityStage 7.0 to 7.5 server is installed. The migration utility is installed with the engine layer.
Installing the WebSphere DataStage and QualityStage server on a different computer from the QualityStage 7.0 or 7.5 server (UNIX or Linux only)	Download and unzip the standalone migration utility on the computer where the QualityStage Version 7.0 to 7.5 server is installed. You can download the migration utility from http://www.alphaworks.ibm.com/ .
Installing the WebSphere DataStage and QualityStage server on a different computer from the QualityStage 7.0 or 7.5 server (Windows or UNIX or Linux)	Install IBM Information Server on the target computer. The migration utility is installed with the engine layer. Before you run the migration utility on the target computer, make the QualityStage 7.0 to 7.5 project metadata available to the target computer. Make the metadata available by copying the entire QualityStage 7.0 to 7.5 project directory and all its contents to the target computer, or by making the directory available via the network (for example, by mapping a drive on Windows).

- If you are migrating from a Unicode-enabled version of QualityStage Version 7.0 to 7.5, make the following changes to the qsmig.env file (located in the directory where the migration tool is installed).
 - Change the line, FLDEXTPR=0 to FLDEXTPR=1
 - For the Japanese language, add the line, QSLCHARSET="CP932"
 - For the Chinese language, add the line QSLCHARSET="CP936"
- On Windows, if the QualityStage Version 7.0 to 7.5 project directories are on a different drive from the migration utility, add the following line to the qsmig.env file:
 - APPCLIB=\${PROJ}/Controls

After you install the migration utility, you can migrate QualityStage Version 7.0 to 7.5 jobs to WebSphere QualityStage Version 8.0.1.

Related tasks

"Importing the migrated files into the Designer client" on page 39
 After you complete the file migration, you import the files into the WebSphere DataStage repository.

Running the migration utility

Run the migration utility to create a file that contains the project metadata stored on a QualityStage 7.0 to 7.5 server.

The migration utility runs on UNIX and Linux. For Windows, the script requires the MKS (Mortice Kerns Systems) Toolkit.

The migration utility is automatically installed when you install the WebSphere DataStage and QualityStage component of the IBM Information Server suite. By default, the utility is in the /opt/IBM/InformationServer/Server/PXEngine/bin

directory on a UNIX or Linux computer or C:\IBM\InformationServer\Server\PXEngine\bin on a Windows computer. The utility is also available as a standalone program on supported platforms. You specify the location of the standalone program when you install it.

To run the migration utility:

1. Ensure that the QualityStage 7.0 to 7.5 project metadata is accessible from the computer on which the migration utility is installed.
2. From the migration utility directory, run the script to start the utility.
 - On UNIX or Linux computers, enter `./qsmigrate.sh`.
 - On Windows, double-click the **qsmigrate.bat** file.
3. When prompted, enter the full path name of the QualityStage 7.0 to 7.5 project directory that contains the jobs that you want to migrate. The utility returns a list of the jobs and rule sets in the project.
4. Enter the number of an option from the list of options.
5. Select one of the following procedures, depending on the migration option that you selected:

Option	Description
1. Migrate multiple QualityStage 7.x jobs plus dependencies 2. Migrate multiple QualityStage 7.x jobs in expanded form plus dependencies	1. Enter a name for the output file that the migration utility produces. All the objects that you migrate are written to this file. The utility reports the path name of this file each time it informs you of the success or failure of an object migration. 2. When prompted with a job name, enter Y to migrate that job or N to skip that job. Repeat this step until you have migrated or skipped all the jobs. 3. When prompted with the name of a rule set or a match specification, enter Y to migrate that object or N to skip that object. Repeat this step until you have migrated or skipped all the objects.
3. Migrate individual QualityStage 7.x job 4. Migrate individual QualityStage 7.x job in expanded form	When prompted, enter the name of the job that you want to migrate. The output file name is derived from the job name.
5. Migrate individual QualityStage 7.x job plus dependencies 6. Migrate individual QualityStage 7.x job in expanded form plus dependencies	1. When prompted, enter the name of the job that you want to migrate. The output file name is derived from the job name. 2. When prompted with the name of a rule set or a match specification, enter Y to migrate that object or N to skip that object. Repeat this step until you have migrated or skipped all the objects
7. Migrate individual QualityStage 7.x standardization rule set	When prompted, enter the name of the rule set or match specification that you want to migrate. The output file name is derived from this name.

After you complete the migration of all of your jobs and objects, move the .dsx files that were created by the utility to the computer where the WebSphere DataStage client is installed.

Importing the migrated files into the Designer client

After you complete the file migration, you import the files into the WebSphere DataStage repository.

To import the contents of the migrated files into the Designer client:

1. Move the .dsx file or files that you created when you ran the migration script to the computer where the WebSphere DataStage clients are installed.
2. Open the Designer client and attach to the project where you want to save the QualityStage jobs.
3. Select **Import** → **DataStage Components**.
4. Specify the name of the migration file and click **OK**. The migrated jobs, rules sets, and match specifications are saved in the following folders:
 - *Project_name* → **Jobs** folder.
 - *Project_name* → **Standardization Rules** → **Imported Rules** → **Rule Sets** folder.
 - *Project_name* → **Match Specifications** folder.

Related tasks

“Installing the WebSphere QualityStage migration utility” on page 36
Use the migration utility to migrate QualityStage Version 7.0 or later to WebSphere QualityStage Version 8.0.1. The migration utility is installed when you install IBM Information Server, or you can download it and install it separately.

Provisioning imported rule sets

You need to provision imported rule sets to the Designer client before a job that uses them can be compiled.

To provision imported rule sets:

1. In the Designer client, find the rule set within the repository tree *Project* → **Standardization Rules** → **Imported Rules** → **Rule Sets** folder.
2. Select the rule set.
3. Right-click and select **Provision All** from the menu.

You can compile and run any job that uses the rule set.

Preparing imported match specifications for use

You need to save the imported match specification in the WebSphere DataStage Designer client and then provision the specification before you can use it in a job.

To prepare imported match specifications for use:

1. In the Designer client, find the match specification within the repository tree *Project* → **Match Specifications** folder.
2. Select the match specification and double-click to open the Match Designer.
3. Select **Save** → **All Passes**.
4. Select **Save** → **Specification**.
5. Click **OK** to close the Match Designer.
6. From the repository, select the match specification.
7. Right-click the match specification and select **Provision All** from the menu.

You can now use the match specification in a Match job.

Preparing migrated jobs for operation

You must prepare migrated jobs before they can be run. The steps can vary depending on the migration option that you selected.

For jobs that you migrated in Legacy operational mode (options 1, 3, or 5), compile the job to prepare it.

To compile a job migrated in the Legacy operational mode:

1. Double-click the job in the Designer client repository tree to open it on the Designer canvas. When you run the migrated job, the results will vary depending on how you ran the job in previous versions of QualityStage. If you used any mode other than the parallel extender mode, the results might be significantly different from previous runs.
2. If you did not previously run the job in parallel extender mode, insert a sort operation into the job design:
 - a. Double-click the target Sequential File stage to open the **Input → Partitioning** page.
 - b. Select **Sort Merge** from the **Collector type** list.
 - c. Under the Sorting section, click **Perform sort**.
3. Click **OK** to close the window.
4. Select **File → Compile** to compile the job.

Preparing migrated jobs in the expanded format

You must prepare jobs migrated in expanded format before they can be run.

To prepare jobs migrated in expanded format:

1. Double-click the job in the Designer client repository tree view to open it. The job contains both QualityStage Legacy stages and Data Quality stages.
2. If you have Standardize or Survive stages, double-click each stage to open it and then click **OK**.
3. Review any migration warnings that are displayed at the bottom of the job and resolve these issues.
4. Save the job.
5. Select **File → Compile**. When you run the migrated job, the results will vary depending on how you ran the job in previous versions of QualityStage. If you used any mode other than the parallel extender mode, the results might be significantly different from previous runs.
6. If you did not previously run the job in parallel extender mode, insert a sort operation into the job design:
 - a. Double-click the target Sequential File stage to open the **Input → Partitioning** page.
 - b. Select **Sort Merge** from the **Collector type** list.
 - c. Under the Sorting section, click **Perform sort**.
7. Optional: Replace QualityStage plug-in stages with the equivalent data quality or processing stage.
 - a. Double-click the QualityStage plug-in stage to open the Properties window.
 - b. Find the stage that offers functionality that is equivalent to the plug-in stage functionality from the Data Quality section of the palette.

- c. Substitute the QualityStage plug-in stage with the equivalent data quality stage or stages. To optimize your job, it is more efficient to replace the QualityStage plug-in stages.
- d. Configure the new stage or stages.
- e. Compile the job.

WebSphere QualityStage data quality stages

You can use the table of replacement stages to select the WebSphere DataStage stage to substitute for the QualityStage 7.0 to 7.5 stage.

The following table lists replacement stages for previous versions of QualityStage stages.

Table 1. Replacement WebSphere DataStage and QualityStage stages for migrated QualityStage stages.

QualityStage 7.0 to 7.5 stage	Purpose	Replacement stage
Abbreviate	Creates match keys from company names.	No direct replacement. Use the Standardize stage to reformat company names and pair with an appropriate match.
Build	Rebuilds a single record from multiple records that are created with a Parse stage.	No direct replacement. Build was often used with Parse to analyze multi-domain data fields. Use Standardize to accomplish the same function in one step.
Collapse	Generates a list of each unique value in single-domain data fields.	Sort stage
Collapse	Generates frequency counts of data values in a field or a group of fields.	Aggregate stage
Format Convert	Reformats files from delimited to fixed-length and vice versa.	Sequential File stage
Format Convert	Provides I/O to an ODBC database.	ODBC stage or database specific stage
Investigate	Analysis of data quality.	Investigate stage and the Reporting tab for the WebConsole for IBM Information Server.
Match	Identifying data duplicates in a single file using fuzzy match logic.	Unduplicate Match stage in conjunction with the Match Frequency stage.
Match	Pairing records from one file with those in another using fuzzy match logic.	Reference Match stage in conjunction with the Match Frequency stage.
Multinational Standardize	Standardize multinational address data.	MNS stage
Parse	Tokenizes a text field by resolving free-form text fields into fixed-format records that contain individual data elements.	No direct replacement. Parse was often used with Build to analyze multi-domain data fields. Use the Standardize stage to accomplish the same function in one step.

Table 1. Replacement WebSphere DataStage and QualityStage stages for migrated QualityStage stages. (continued)

QualityStage 7.0 to 7.5 stage	Purpose	Replacement stage
Program	Invokes a customer-written program.	Depends on the functionality of the customer-written program. Possibilities include adding a Parallel Build, Custom, or Wrapped stage type.
Select	Conditionally routes records that are based on values in selected fields.	Switch and Filter stages
Sort	Sorts a list.	Sort stage
Standardize	Breaks down multi-domain data columns into a set of standardized single-domain columns.	Standardize stage
Survive	Produces the best results record from a group of related records.	Survive stage
Transfer	Rearranges and reformats columns in a record.	No separate stage is required to do this in QualityStage 8.0.
Transfer	Acts as a gatekeeper for files in non-standard formats (variable length records, non-standard code page, binary or packed data).	Sequential File or Complex Flat File stage
Transfer	Produces multiple output records from a single input record.	Splitting records can be achieved by Copy stage followed by Funnel stage
Transfer	Adds record keys that consists of sequence number plus an optional fixed "file identifier."	Surrogate Key Generator stage
Unijoin	Join records from two files based on a key.	Join stage or Lookup stage
Unijoin	Pairing records from one file with those in another using fuzzy match logic.	Reference Match stage in conjunction with Match Frequency stage.
Unijoin	Merges data from multiple records into one.	Join stage and Merge stage
Unijoin	Manipulate and transform data record.	Transformer stage
WAVES	Standardize and validate multinational address data.	WAVES stage

Upgrading the QualityStage plug-in DLLs

If you previously ran DataStage jobs that included the QualityStage plug-in stage, then you need to configure the QualityStage plug-in that provides this stage after you upgrade to WebSphere DataStage and QualityStage 8.0.

If you want to run existing WebSphere DataStage jobs that use the QualityStage plug-in stage, you must use the 7.5.1 version of the QualityStage plug-in. To use the 7.5.1 version of the plug-in, you must rename the DLLs of the plug-in that are in the IBM/InformationServer/Server/DSComponents/bin directory. The default version of the plug-in in WebSphere DataStage is 7.5.0. Follow these instructions to disable the 7.5.0 dlls and enable the 7.5.1 dlls.

1. Rename the irtc.dll file to irtc.dll.750.
2. Rename the irtc.dll.751 file to irtc.dll.
3. Rename the irtc_u.dll file to irtc_u.dll.750.
4. Rename the irtc_u.dll.751 file to irtc_u.dll.
5. Rename the qsdds.dll file to qsds.dll.750.
6. Rename the qsds.dll.751 file to qsds.dll.
7. Rename the following files.

If your installation is	Do this
English	1. Rename the qsdsenu.dll file to qsdsenu.dll.750. 2. Rename the qsdsenu.dll.751 file to qsdsenu.dll.
Japanese	1. Rename the qsdsjpn.dll file to qsdsjpn.dll.750. 2. Rename the qsdsjpn.dll.751 file to qsdsjpn.dll.

Accessing information about the product

IBM has several methods for you to learn about products and services.

You can find the latest information on the Web:
www.ibm.com/software/data/integration/info_server/

To access product documentation, go to publib.boulder.ibm.com/infocenter/iisinfsv/v8r0/index.jsp.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at www.ibm.com/shop/publications/order.
- To order publications by telephone in the United States, call 1-800-879-2755.

To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at www.ibm.com/planetwide.

Providing comments on the documentation

Please send any comments that you have about this information or other documentation.

Your feedback helps IBM to provide quality information. You can use any of the following methods to provide comments:

- Send your comments using the online readers' comment form at www.ibm.com/software/awdtools/rcf/.
- Send your comments by e-mail to comments@us.ibm.com. Include the name of the product, the version number of the product, and the name and part number of the information (if applicable). If you are commenting on specific text, please include the location of the text (for example, a title, a table number, or a page number).

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing 2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

(C) (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. (C) Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM trademarks and certain non-IBM trademarks are marked at their first occurrence in this document.

See www.ibm.com/legal/copytrade.shtml for information about IBM trademarks.

The following terms are trademarks or registered trademarks of other companies:

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT®, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel®, Intel Inside® (logos), MMX and Pentium® are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product or service names might be trademarks or service marks of others.

Index

A

Abbreviate stage, replacing 41
accessibility 45
Administrator client 5, 7, 12, 26, 31
attaching to DataStage servers 18

B

Build stage, replacing with new 41

C

client 17
 Administrator 5, 12, 18, 31
 Designer 12, 13, 14, 18, 31, 32, 33, 35, 39
 Director 13, 18, 32
 Manager 8, 18, 27
 Multi-client Manager 17
command line interface
 connecting to a project from 22
 exporting from 8, 27
compiling
 jobs 14, 33
 routines 14, 33
components 1
configuration files 5, 7, 9, 13, 23, 26, 28, 32
connecting to DataStage servers 18
connecting to multiple DataStage servers
 Version 7.5.1 and earlier
 Administrator client 19
 Version 7.5.2 and earlier
 Designer, Manager, and Director clients 19
 Version 8.0.1
 Administrator client 21
 Designer and Director clients 20
conversion script 40

D

DataStage 1
DataStage release 7.0 to release 7.5 4
DataStage releases earlier than 7.0 23
DB2 1
Designer client 35
 compiling jobs and routines 14, 33
 importing migrated files (QualityStage) 39
 importing projects (DataStage) 12, 31
 restoring hashed files 13, 32
Director client 13, 32
documentation
 accessible 45
dscmdexport 8, 27
dsenv 5, 9, 24, 28
DSHOME 5, 9, 24, 28
dsjob 22

DSPParams 5, 6, 9, 10, 24, 28, 29

E

exporting
 job dependencies 10, 29
 projects (DataStage) 8, 27

F

flat files 10, 13, 29, 32
Format Convert stage, replacing 41

H

hashed files 10, 13, 29, 30, 32

I

importing
 projects (DataStage) 12, 31
installation program 6, 12, 16, 25, 31
instance tag 14, 16
integrating existing set-ups 5, 9, 23, 28
ITAG 14, 16

J

job migration (QualityStage)
 expanded 36
 legacy 35
 match specification 36
 rule sets 36

L

layers
 installing 1
legacy migration (QualityStage)
 Legacy stage 35
 QualityStage jobs 35
 replacing legacy operators 41
 replacing Legacy stages 40
legal notices 47
Linux server 11, 30

M

Manager client 8, 27
match specification
 migration 36
metadata repository 1
Microsoft SQL Server 1
migrated files (QualityStage)
 importing 39
 provisioning 39
migration options 3

migration utility
 running 37
moving
 hashed files 10, 30
Multi-Client Manager 17
Multiple Job Compile tool 14, 33
multiple servers 14

N

NLS 6, 25

O

odbc.ini 5, 9, 24, 28
Oracle 1

P

Parse stage, replacing 41
Program stage, replacing 41
projects (DataStage)
 exporting 8, 27
 importing 12, 31
 migrating 7, 26

Q

QualityStage 1, 35
 migrating jobs 36
 migration utility
 installing 36
QualityStage 7.0 to 7.5 job migration
 expanded form 36
 legacy form 35

R

replacing legacy operators (QualityStage) 41
restoring
 configuration files 7, 13, 26, 32
 integrating existing set-ups 7, 13, 26, 32
roadmap
 migrating to a clean installation 7, 26
 migrating using the installation program 5
rule sets 36
provisioning 39

S

schema files 10, 13, 29, 32
screen readers 45
Select stage, replacing 41
servers
 multiple versions 14

- Sort stage, replacing 41
- SQL Server 1
- starting
 - DataStage server 18
- stopping
 - DataStage server 18

T

- tiers
 - installing 1
- trademarks 49
- Transfer stage, replacing 41

U

- Unijoin stage, replacing 41
- uninstalling
 - clients 6, 12, 24, 31
 - server 11, 30
- UNIX and Linux
 - installing migration utility 36
 - running migration utility 37
- UNIX server 11, 30
- upgrade methods
 - using installation program 5
- upgrade options 3
- uvodbc.config 5, 9, 24, 28

V

- Version Control 2, 3

W

- WebSphere Application Server 1
- Windows
 - installing migration utility 36
 - running migration utility 37
- Windows server 11, 30



Printed in USA

SC18-9924-01

