

IBM InfoSphere DataStage and QualityStage
Version 11 Release 3

Administrator Client Guide



IBM InfoSphere DataStage and QualityStage
Version 11 Release 3

Administrator Client Guide



Note

Before using this information and the product that it supports, read the information in “Notices and trademarks” on page 37.

Contents

Chapter 1. What is the Administrator client? 1

- Who can use the Administrator client? 1
- What can you do from the Administrator client? 1

Chapter 2. Using the Administrator client 3

- Starting the Administrator client. 3
- Description of the Administrator client 3
- Setting the InfoSphere Information Server engine timeout 4
- Issuing InfoSphere Information Server engine commands 4

Chapter 3. Administering projects 7

- Adding projects 7
- Deleting projects 8
- Moving projects 8
 - Moving a project 8

Chapter 4. Setting project properties 11

- General page 11
 - Enabling job administration in the Director client 11
 - Enable runtime column propagation for parallel jobs 12
 - Enabling editing of internal reference information 12
 - Controlling the import of metadata from connectors. 13
 - Protecting a project. 13
 - Setting environment variables 14
 - Enabling the generation of operational metadata at the project level (parallel and server jobs) . . 15
 - Setting the workload management queue 16
- Permissions page 16
 - InfoSphere DataStage user roles 17
 - Assigning InfoSphere DataStage user roles . . . 17
 - Changing the operator's view of job log entries 18

- Enabling tracing on the InfoSphere Information Server engine. 18
- Specifying a scheduling user. 19
- Supplying mainframe information. 19
- Tunables page 21
 - Hashed file caching. 21
 - Row buffering 21
- Parallel page 22
- Sequence page 22
- Deploying on remote systems 23
- Logs page 23
 - Purging job log files 23

Chapter 5. Configuring projects for NLS 25

- Changing project maps 25
 - Server job project maps 25
 - Parallel job project maps 26
- Changing project locales 26
 - Server job locales 26
 - Parallel job locales 27
- Client and server maps 27

Appendix A. Product accessibility 29

Appendix B. Contacting IBM 31

Appendix C. Accessing the product documentation 33

Appendix D. Providing feedback on the product documentation 35

Notices and trademarks 37

Index 43

Chapter 1. What is the Administrator client?

You use the IBM® InfoSphere® DataStage® Administrator client to perform administrative tasks for individual projects located on a single InfoSphere DataStage

The Administrator client exists in parallel with the web client-based Suite Administrator. You can open the Suite Administrator from within the Administrator client by clicking on the Suite Admin hyperlink.

Who can use the Administrator client?

To use all the features of the Administrator client, you need to have been set up as a Administrator within the Suite Administrator.

(If you have been set up as an InfoSphere DataStage user you can open the Administrator client and view information and perform certain non-administrative functions.)

What can you do from the Administrator client?

The IBM InfoSphere DataStage Administrator client is used to perform several administrative tasks for InfoSphere DataStage projects.

The Administrator client allows you to perform the following tasks:

- Add, delete, and move InfoSphere DataStage projects (see "Project Administration")
- Set up user permissions for projects (see "Project Properties")
- Purge job log files (see "Project Properties")
- Set the timeout interval on the InfoSphere DataStage engine (see "Using the Administrator")
- Trace InfoSphere Information Server engine activity (see "Project Properties")
- Set job parameter defaults (see "Project Properties")
- Issue InfoSphere Information Server engine commands from the Administration client (see "Using the Administrator")
- Check license details (see "Using the Administrator")

All of these tasks relate to a particular InfoSphere Information Server engine installation.

Chapter 2. Using the Administrator client

These topics tell you how to start the Administrator, how to perform basic tasks and refers you to other topics for more detailed descriptions of advanced tasks.

Starting the Administrator client

You can start and configure the IBM InfoSphere DataStage Administrator client from the **Start** menu. To use all of the IBM InfoSphere DataStage Administrator client functions, you must first be specified as a DataStage Administrator.

Procedure

1. From the start menu, choose **IBM InfoSphere Information Server > IBM InfoSphere DataStage and QualityStage Administrator**.
2. Specify the host name for the services tier that the InfoSphere Information Server engine you want to administer belongs to.
3. Select the host name for the computer where the InfoSphere Information Server engine you want to administer resides.
4. Specify a **User name** and **Password** for connecting to the specified domain. To use all the Administrator client functions, you must already have been specified as a DataStage Administrator in the Administration tab of the Web console.
5. Click **OK**.
6. If you get a message that a security certificate from the server is not trusted, accept the certificate:
 - a. To view the security certificate, click **View Certificate**.
 - b. Click the **Certification Path** tab, and then select the root certificate.
 - c. Click the **General** tab.
 - d. Click **Install Certificate**, and then click **Next**.
 - e. Select **Place all certificates in the following store**.
 - f. Click **Browse**, and then select **Trusted Root Certification Authorities**.
 - g. Click **Next**, and then click **Finish** to import the certificate.

Description of the Administrator client

The Administrator client is used to set engine-wide properties and administer projects.

The Administrator client has three tabbed pages:

- Use the General page to set engine-wide properties. This page is enabled only when at least one project exists. The engine wide properties are:
 - Inactivity Timeout (see "Setting the InfoSphere Information Server Engine Timeout").
 - NLS client/server map (if NLS is enabled on this system) (see "Client/Server Maps").
- The Projects page lists all the InfoSphere DataStage projects currently known to the InfoSphere Information Server engine. Use this page to administer projects. From here you can:
 - Add and delete projects (see "Project Administration.")

- View and set project properties (see "Project Properties")
- Change the default character set maps and locales for projects (if NLS is enabled) (see "NLS Configuration.")
- Issue InfoSphere Information Server engine commands directly to the project (see "Issuing InfoSphere Information Server Engine Commands").

If no projects exist, only the **Add...** and **Command** buttons are enabled, and **Command** issues InfoSphere Information Server engine commands on the InfoSphere Information Server engine account.

Setting the InfoSphere Information Server engine timeout

You can change the default IBM InfoSphere Information Server engine timeout value.

About this task

By default, the connection between the InfoSphere DataStage client and InfoSphere Information Server engine on Windows times out after 86400 seconds (24 hours) inactivity. By default, the connection between the InfoSphere DataStage client and InfoSphere Information Server engine on UNIX does not time out.

Procedure

1. In the Administrator window, click the **General** tab to move the General page to the front.
2. In the **Inactivity Timeout** area, use the up and down buttons to change the timeout period, or enter the new timeout period in the **seconds** field.
3. To disable inactivity timeout, select the **Do not timeout** check box.
4. Click **Apply** to apply the new settings. The changes take effect immediately for any new client connections.

Results

If inactivity timeout is already disabled when the Administrator starts, the timeout reverts to the 86400-second default when you reenable it.

The timeout also affects the `dsjob -wait` command issued from the command line and the job control api function `DSWaitForJob`. If you are using these features you are advised to set a long timeout, or disable the timeout altogether. It also affects shared containers on the parallel canvas (which use `DSWaitForJob`).

The timeout set in the Administrator client overrides the session timeout set in the IBM InfoSphere Information Server Web console. Connections between InfoSphere DataStage clients and the engine are not affected by session timeouts set in the Web console.

Issuing InfoSphere Information Server engine commands

You can issue an InfoSphere Information Server engine command from the Administration client:

About this task

The Administration client lets you issue InfoSphere Information Server engine commands directly from a selected project rather than having to use a Telnet session.

You cannot issue the following InfoSphere Information Server engine commands from the client: LOGOUT, LO, QUIT, Q, and OFF.

You should take care when issuing InfoSphere Information Server engine commands that manipulate the NLS settings of your system. Incorrect use of these commands could disrupt both InfoSphere DataStage and the rest of your system.

Procedure

1. Click the **Projects** tab in the DataStage Administration Window to move this page to the front.
2. Select the project.
3. Click **Command**. The Command Interface dialog box appears.
4. Either type the InfoSphere Information Server engine command you want to execute into the **Command** field, or double-click a command in the **Command History** list box. A command selected from the command history appears in the **Command** field, where you can edit it.
5. Click **Execute**. The command is added to the **Command History** list box and the Command Output window appears, displaying the results of the command. The command output pauses at the end of the first page. Click **Next** to scroll to the next page of output. If you want to switch paging off, clear the **Pause output at end of page** check box at the bottom of the Command Output dialog box.
6. If the command requires further input, the Command Output window displays a prompt. Enter the response in the field below the command output display, then click **Respond**.
7. When the command has run to completion, click **Close**. This window closes, and the Command Interface dialog box reappears.
8. If you want to save commands to the VOC file on the InfoSphere Information Server engine, select one or more commands from the **Command History** list box, then click **Save As...** . A single command is saved as a sentence and multiple commands as a paragraph. Enter the name of the sentence or paragraph in the Save As dialog box, then click **OK**.
You cannot save a sentence or paragraph that contains a call to itself.
9. Click **Close** to close the Command Interface dialog box.

Chapter 3. Administering projects

You can use the InfoSphere DataStage Administrator client to add, delete and move projects.

These topics describe how to do the following:

- Add new InfoSphere DataStage projects
- Delete projects
- Move projects

When you delete InfoSphere DataStage projects, always do so through the Administrator client to ensure that licensing information is updated correctly.

Adding projects

You can add further projects from the Administrator client as required.

About this task

When you add a new project, you can specify that the new project inherits users and their associated roles from an existing project.

You can also specify that a new project is protected. This is a special category of project and, normally, nothing can be added, deleted or changed in the project. Users can view objects in the project, and perform tasks that affect the way a job runs rather than the jobs design; specifically they can:

- Run jobs
- Set job properties
- Set job parameter default values

A newly created protected project is populated by importing developed jobs and components; only a user with the Production Manager role or Administrator role can perform the import, no other types of user can import into a special project. (To mark an existing project as protected, use the Project Properties dialog box - see "Protecting a Project").

Protected projects provide a way of maintaining the integrity of jobs intended to be run in a 'production' environment, that is, the finished jobs. Giving widespread access to the Production Manager role or Administrator role would obviate the benefits of using protected projects.

In addition to having the InfoSphere DataStage administrator role in order to add a project, you also require the correct operating system permissions on the computer on which the InfoSphere Information Server engine you are administering resides.

Procedure

1. Click the **Projects** tab in the Administrator window to move this page to the front.
2. Click the **Add...** button. The Add Project dialog box appears.

3. Enter the project name in the **Name** field. It is automatically appended to the default location for the new project.
You cannot add a project to a location that already exists.
4. To change the default location, enter a new location in the **Location on host** field, or use the **Browse...** button to select a new location.
5. If you want the project to be a protected one, select the **Create protected projected** check box.
6. If you want to copy the users who can access the new project, and their roles, from an existing project, select **Copy roles from existing project**, and select the project from which to copy from the list.
7. Click **OK**.

Deleting projects

You can use the InfoSphere DataStage Administrator client to delete projects.

Procedure

1. Ensure that no users are connected to the project you want to delete.
InfoSphere DataStage returns an error message if you try to delete a project while another user is connected to it.
2. Click the **Projects** tab on the Administrator window to move this page to the front.
3. Select the project you want to delete, then click **Delete**.
4. You are prompted to confirm the deletion. If you confirm, the project is deleted immediately.

Moving projects

To move an InfoSphere DataStage project, you must export it to a file and then import it to the new location.

You then need to delete the original project by using the Administrator client.

Before attempting to move a project in this way, ensure that no one is running any jobs in it. The easiest way to do this is to start the InfoSphere DataStage Director, attaching to the project you want to move, and check whether any jobs are currently running.

When you move projects, project level settings such as environment variable settings and project options are not included.

Moving a project

Move an InfoSphere DataStage project to a new location.

Procedure

1. Start the Designer client, attaching to the project that you want to move.
2. Click **Export > DataStage Components** from the main menu. The Repository Export dialog box appears.
3. Click the **Select all** hyperlink to select all the objects in the repository tree.
4. From the drop-down list, click **Export job designs with executables (where applicable)**.

5. Click the **Options** button to open the Export Options dialog box. Specify the following under the **Default > General** branch:
 - specify that source code is included with exported routines
 - specify that source code is included with job executables
 - specify that source content is included with data quality specifications
6. Select dsx as the **Type of export**.
7. Specify or select the file that you want to export to. If the file exists, you can click the **View** button to look at the file. The file is shown using either the default viewer for this file type specified in Windows, or a viewer you have specified in the Export Options dialog box.
8. Click **Export** to export the project to the specified file.
9. Create a project with the required name in the location where you want to move the existing project to (see "Adding Projects" for instructions).
10. Attach the Designer client to the new project you created.
11. Click **Import > DataStage Components...** . The DataStage Repository Import dialog box appears.
12. Enter the file name you used for the exported project.
13. Click the **Import all** option, then click **OK**. The project is imported.

Chapter 4. Setting project properties

You can view and change particular project properties.

Procedure

1. Click the **Projects** tab on the Administrator window to move the Projects page to the front.
2. Select the project.
3. Click the **Properties** button. The Project Properties window appears.

Results

You use the pages in the Project Properties window to do the following:

- **General.** Enable job administration from the Director client, enable runtime column propagation for parallel jobs, define a project-wide setting for auto-purge of the job log, set up environment variables, enable the automatic sharing of metadata when importing by using the connector import wizard, and generate operational metadata.
- **Permissions.** View and assign InfoSphere DataStage user roles to users and groups for a particular project. You assign suite user roles in the IBM InfoSphere Information Server Web console.
- **Tracing.** Enable or disable tracing on the InfoSphere Information Server engine.
- **Schedule.** Set up a user name and password to use for running scheduled InfoSphere DataStage jobs. The **Schedule** tab is enabled only if you have logged on to a Windows server.
- **Mainframe.** Set mainframe job properties and the default platform type. This page is enabled only if mainframe edition is installed.
- **Tunables.** Configure cache settings for Hashed File stages.
- **Parallel.** Set parallel job properties and defaults for date/time and number formats.
- **Sequence.** Set compilation defaults for job sequences.

General page

Use the General page to control a variety of features in the selected project.

Enabling job administration in the Director client

The administrator can enable or disable job administration features in the Director client using the Administrator client.

About this task

These features let InfoSphere DataStage operators release the resources of a job that has aborted or hung, and so return the job to a state in which it can be rerun when the cause of the problem has been fixed. The default is for these features to be disabled.

Procedure

1. Click the **Projects** tab in the Administrator window to move this page to the front.
2. Select the project.
3. Click **Properties**. The Project Properties window appears, with the General page displayed.
4. Select the **Enable job administration in Director** check box.
5. Click **OK**.

Results

This procedure enables two commands in the Director **Job** menu:

- **Cleanup Resources**
- **Clear Status File**

Cleanup Resources is also enabled in the Monitor window shortcut menu in the Director. **Cleanup Resources** lets the user:

- View and end job processes
- View and release the associated locks

Cleanup Resources is not compatible with tracing (see "Enabling Tracing on the InfoSphere DataStage Engine"). If you enable tracing and job administration, the Director displays an error message when **Cleanup Resources** is chosen.

Clear Status File removes the status records associated with all stages of the selected job. It should be used only if the user believes all the job processes have gone away and the job cannot be reset correctly.

These two commands give the user considerable control over job resources, and should be used with care.

Enable runtime column propagation for parallel jobs

Before you can enable runtime column propagation in a stage, runtime column propagation must be enabled for parallel jobs at the project level from the IBM InfoSphere DataStage and QualityStage Administrator. If you enable this feature, stages in parallel jobs can handle undefined columns that they encounter when the job is run, and propagate these columns through to the rest of the job.

About this task

This check box enables the feature, to actually use it you need to explicitly select the option on each stage. Enabling this feature makes the following sub-property available:

- **Enable Runtime Column Propagation for new links**. Select this to have runtime column propagation enabled by default when you add a new link in a InfoSphere DataStage job. If you do not select this, you will need to enable it for individual links in the stage editor when you are designing the job.

Enabling editing of internal reference information

Select **Enable edition of internal references in jobs** to enable the editing of the **Table definition reference** and **Column definition reference** fields in the column definitions of stage editors.

About this task

These settings affect how stage editors behave in the Designer client.

The two fields identify the table definition, and individual columns within that definition, that columns have been loaded from. These fields are enabled on the Stage Editor's **Columns** tab when you specify settings in the Grid Properties dialog box in the Designer client.

Controlling the import of metadata from connectors

Use the **Share metadata when importing from Connectors** option to specify what objects are created when you use the Designer client to import metadata via connectors.

About this task

This option is selected by default. This means that when you use the **Import > Table Definitions > Start Connector Import Wizard** command in the Designer client a table definition object is created in the project and a data collection is created in the dynamic repository. Other suite components, and other DataStage projects, can then access the data collection in the dynamic repository. This option is selected by default. If the option is not selected then only a table definition object in the project repository is created.

Protecting a project

If you are a Production Manager or Administrator user, you can convert the project to a protected project.

About this task

A protected project is a special category of project and, normally, nothing can be added, deleted, or changed in the project.

Note: Currently on UNIX systems only root or the administrative user can protect or unprotect a project.

Users can view objects in the project, and perform tasks that affect the way a job runs rather than the jobs design.

- Run jobs
- Set job properties
- Set job parameter default values

Users with Production Manager and Administrator status can import existing InfoSphere DataStage components into a protected project.

To convert the current project into a protected one, click the **Protect Project** button. A dialog box asks you to confirm the conversion. Click **OK** to proceed. The **Protect Project** button changes to an **Unprotect Project** button, so you can convert the project back to unprotected status if required.

After a project has been protected, only a Production Manager or Administrator user (who is root or the administrative user on UNIX) can unprotect it. Operators can add and amend environment variables in protected projects.

Protected projects are identified in the Administrator client by the string (Protected) after their name.

Setting environment variables

You can set project-wide defaults for general environment variables or ones specific to parallel jobs from this page.

Procedure

- You can also specify new variables. All of these are then available to be used in jobs. They are added to jobs in the same way as job parameters.
- To define environment variables, click the **Environment...** button. The Environment Variables dialog box appears.
 - Choose the type of environment variable for which you want to set a default from the tree in the left pane. A list of available variables appears in the right pane. You can select a new value for the variable in the **Value** column.
 - To define a new variable, choose **User Defined**. A dialog box appears asking you to specify the name of the variable and a default value. You can also set a type for a user defined variable. Choose between string (the default) and encrypted. If you choose encrypted, a further dialog box requests you enter the encrypted string value and then confirm it.
 - Click **Export** to export environment variables from a project to a file. See “Exporting environment variables.”
 - Click **Import** to import environment variables from a file. See “Importing environment variables” on page 15.
 - Click **Set to Default** to set the selected environment variable to its installed default value.
 - Click **All to Default** to set the environment variables currently visible to their installed default values.
 - Click **Variable Help** to get help information about the selected variable.

When you change the setting for an environment variable, all the jobs in the project are affected. If you want to change an environment variable for a particular job, leave the **Value** column empty, and specify the setting of the environment variable via a job parameter.

Exporting environment variables

You can export environment variables to a file (*.env) from the Projects page in the IBM InfoSphere DataStage and QualityStage® Administrator client. You can use the file to import environment variables from the file to other projects.

Procedure

1. Open the InfoSphere DataStage and QualityStage Administrator client.
2. On the Projects page, select the project that you want to export environment variables from, and click **Properties** to open the Project Properties window.
3. On the General page, click **Environment**.
4. In the Environment Variables dialog, click **Export to File**.
5. In the Environment Variable Export dialog, in the tree structure that shows all the environment variables in the project, select the check boxes for the environment variables that you want to export.
6. Click the **Export** button.

7. When prompted for the location to export the file, select a path on the client to export to and specify a file name. The default file name is **<project_name>.env**. Click **Save**.

Importing environment variables

You can import environment variables from a file into a project. The imported environment variables can then apply to all jobs within that project.

About this task

When you modify or import the value of an environment variable for a project, the new project setting can apply to all the jobs in the project that use the environment variable.

When you import environment variables, the following rules are applied to the project settings:

- For an environment variable with the same name and same value, no action is taken.
- For an environment variable with a different name from any existing environment variable, the import continues without prompting.
- If there is already an environment variable with the same name but a different value, you are prompted to overwrite the value.

Procedure

1. Open the InfoSphere DataStage and QualityStage Administrator client.
2. On the Projects page, select the project that you want to import environment variables to, and click **Properties** to open the Project Properties window.
3. On the General page, click **Environment**.
4. In the Environment Variables dialog, click **Import from File**.
5. When prompted, specify the path to the environment variables file (*.env) that you want to use.
6. In the Environment Variable Export dialog, in the tree structure that shows all the environment variables and their values in the file, select the check boxes for the environment variables that you want to import.
7. Click **Import** to import the selected environment variables.
8. In the report that displays when the import is complete, review the summary of environment variables that have been created, overwritten, or unaffected. Click **More** to see an expanded report of the results of the import. This expanded report also contains information for any lines in the file that could not be parsed because they do not match the expected format.

If there are no recognizable environment variable definitions, the import returns a message that states that the file does not contain any recognizable environment variable definitions.

Enabling the generation of operational metadata at the project level (parallel and server jobs)

You can enable the generation of operational metadata at the project level in IBM InfoSphere DataStage and QualityStage server and parallel jobs.

About this task

When you run a job in the project, metadata is collected that describes the job run and the data warehouse resources that it affects. After the job runs, the OMDMonApp service imports this operational metadata into the metadata repository. You can analyze it in IBM InfoSphere Information Governance Catalog and report on the operational metadata in the Web console.

In the Director client, you can override the project-level setting for individual jobs.

Procedure

1. Open the InfoSphere DataStage and QualityStage Administrator client.
2. On the Projects page, select the project that you want to generate operational metadata for, and click **Properties** to open the Project Properties window.
3. Select **Generate operational metadata**.
4. Click **OK**.

What to do next

Whenever you run a job that is in the project, operational metadata is generated. For parallel and server jobs, you can override the default project-level setting before you run a job in InfoSphere DataStage and QualityStage Director.

Setting the workload management queue

You can set the default workload management queue for the project.

Procedure

1. Open the InfoSphere DataStage and QualityStage Administrator client.
2. On the Projects page, select the project that you want to export environment variables from, and click **Properties** to open the Project Properties window.
3. Select the default **Workload management Queue** for the project. If you do not specify a default queue for the project, the workload management default queue is used for jobs that do not specify a queue.
4. Click **OK**.

Permissions page

These topics describes InfoSphere DataStage user roles and how to change the assignment of these roles to users or groups.

Before any user can access InfoSphere DataStage they must be defined in the Suite Administrator tool as a DataStage Administrator or a DataStage User. As a DataStage administrator you can define whether a DataStage user can access a project, and if so, what category of access they have.

Using the Suite Administrator tool you can also add groups and assign users to groups. These groups are in turn allocated the role of DataStage Administrator or DataStage User. Any users belong to an administrator group will be able to administer InfoSphere DataStage. As a DataStage Administrator you can give a DataStage user group access to a project and assign a role to the group.

When setting up users and groups, these still have to have the correct permissions at the operating system level to access the folders in which the projects reside.

Users must also be administrative users on the Windows computer where InfoSphere DataStage clients are installed in order to be able to use those clients.

This section also explains how to change the default view of job log entries for those who have the DataStage Operator or DataStage Super Operator role.

The Permissions page contains the following controls:

- **Roles.** This window lists all the users and groups who currently have access to this project and lists their roles. Note that this window will always include users who have been defined as DataStage Administrators in the Suite Administrator tool, and you cannot remove such users from the list or alter their user role.
- **User Role.** This list contains the four categories of InfoSphere DataStage user you can assign. Choose one from the list to assign it to the user currently selected in the roles window.
- **Add User or Group.** Click this to open the Add Users/Groups dialog box in order to add a new user or group to the ones listed in the roles window.
- **Remove.** Click this to remove the selected user or group from those listed in the roles window.
- **DataStage Operator can view full log.** By default this check box is selected, letting an InfoSphere DataStage operator view both the error message and the data associated with an entry in a job log file. To hide the data part of the log file entry from operators, clear this check box. Access to the data is then restricted to users with a developer role or better.

After you have assigned users, groups, and roles for a project on the Permissions page, you can use this project as a template when adding new projects. New projects can inherit the users, groups, and roles from the project so that you do not have to repeat the process every time that you add a new project.

InfoSphere DataStage user roles

To prevent unauthorized access to InfoSphere DataStage projects, you must assign the InfoSphere DataStage users or groups on your system the appropriate InfoSphere DataStage user role.

For more information about user roles, see InfoSphere DataStage and QualityStage user roles.

Assigning InfoSphere DataStage user roles

You can view the InfoSphere DataStage users and groups in the Administrator client. You assign suite user roles in the IBM InfoSphere Information Server Web console.

Procedure

1. Open the Permissions page.
2. Click the **Projects** tab in the Administrator window to move this page to the front.
3. Select the project.
4. Click **Properties** to display the Project Properties window.
5. Click the **Permissions** tab to move this page to the front. By default the **Permissions** tab shows all the users or groups that have been defined as DataStage Administrators in the Administration tab of the Web console. Such

users automatically have access to the DataStage Administrator role for this project. You cannot remove that access right as the **User Role** drop-down list and **Remove** button are both disabled.

6. To add an InfoSphere DataStage user:
 - a. Open the Web console.
 - b. In the Navigation pane, click **Users and Groups**.
 - c. Click **Users**.
 - d. To assign a user, follow the instructions in the Web console help.
7. To remove an InfoSphere DataStage user or group you have previously added to this project:
 - a. From the Administration tab of the Web console, select the user in the Users table.
 - b. Click **Delete** in the task pane.

Changing the operator's view of job log entries

Entries in a job log file consist of an error message and data associated with the error by default. You can hide the associated data by changing the operator's view of job log entries.

About this task

By default, the IBM InfoSphere DataStage and QualityStage Director and Designer clients show both the message and the data when a DataStage operator or super operator views the details of an event. To override this setting so that only the error message is visible to an operator, clear the **DataStage Operator can view full log** check box on the Permissions page, then click **OK**. Access to the associated data is then restricted to users with Developer rights.

Enabling tracing on the InfoSphere Information Server engine

You can trace activity on the InfoSphere Information Server engine to help diagnose project problems.

About this task

The default is for tracing to be disabled. When you enable tracing, information about engine activity is recorded for any clients that attach to the project. This information is written to trace files, and users with in-depth knowledge of the system software can use it to help identify the cause of a client problem.

If tracing is enabled, users receive a warning message whenever they start an InfoSphere DataStage client.

Tracing is not compatible with the job administration command **Cleanup Resources**, which you also enable from the Administration client. See Enabling job administration in the Director client for more information.

Procedure

1. Click the **Projects** tab in the **Administrator** window to move this page to the front.
2. Select the project.
3. Click **Properties** to display the Project Properties window.

4. Click the **Tracing** tab to move this page to the front.
5. Click the **Enabled** option button in the **Server side tracing** area.
6. Click **OK** to activate tracing. Trace files are added to the **Trace files** list box on the Tracing page whenever a client attaches to the project.
7. To view a trace file, double-click the file name in the **Trace files** list box, or select the file name and then click the **View** button. The View trace file window displays the trace file. To copy trace information to the clipboard, select text in the View trace file window, then click **Copy**.
8. To delete a trace file, select one or more file names in the **Trace files** list box, and then click the **Delete** button.

Specifying a scheduling user

Where the engine is running on a Windows computer, InfoSphere DataStage uses the Windows Schedule service to schedule jobs.

About this task

This means that by default the job runs under the user name of the Schedule service, which defaults to NT system authority. You might find that the NT system authority does not have enough rights to run the job. To overcome this, you can define a user name to run scheduled jobs in a project, as described in this procedure.

Procedure

1. Click the **Projects** tab in the Administrator window to move this page to the front.
2. Select the project.
3. Click **Properties** to display the Project Properties window.
4. Click the **Schedule** tab to move this page to the front.
5. Enter the user name and password you want to use to run the scheduled jobs.
6. Click **Test** to test that the user name and password can be used successfully. This involves scheduling and running a command on the InfoSphere Information Server engine, so the test might take some time to complete.
7. Click **OK** to save the user name and password.

Supplying mainframe information

When mainframe jobs are uploaded from the Designer to a mainframe computer, a JCL script is also uploaded. You must define the mainframe job properties and specify the default platform type for inclusion in the script

About this task

The script provides information required when the mainframe job is compiled. If you have a project that supports mainframe jobs, you need to define the mainframe job properties for inclusion in the script. You also need to set the default platform type. This platform type is the default when new objects, such as a mainframe routine, are created in the Designer. The page also allows you to specify details about how nulls are defined in flat files.

Procedure

1. Click the **Projects** tab in the Administrator window to move this page to the front.
2. Select a project that supports mainframe jobs.
3. Click **Properties** to display the Project Properties window.
4. Click the **Mainframe** tab to move this page to the front.
The **Mainframe** tab is enabled only if the selected project supports mainframe jobs.
5. Select the default date format for the project from the Date Format drop-down list. (The project default can be overridden at job level in the **Job properties** dialog box in the Designer client.)
6. Select the type of platform from the **Platform Type** list.
7. Enter the database system name, your user name, and password.
8. To change the default maximum block size and maximum blocking factor for the selected platform, enter new values in the **Max. Blocking Factor** and **Max. Block Size** fields.
(On an OS/390[®] platform, the maximum blocking factor can be a value from 1 to 32767. The maximum block size can be zero or a value from 4096 to 32767. If you specify zero, the operating system will determine the optimal block size for you.)
9. To make the selected platform the default, click the **Set as Default** button. "Default" appears in parentheses after the platform type.
10. Select **Perform expression semantic checking** to turn semantic checking on for the expression editor by default in all mainframe jobs in the project. You can override this setting at job level if required.
11. Select **Support extended decimal** to enable the use of extended decimal types in the jobs in the project. The **Maximum decimal size field** is then enabled.
12. Select **Generate operational metadata** to have mainframe jobs in the project generate operational metadata by default. You can override this setting at job level if required.
13. If you have selected **Support extended decimal**, in the **Maximum decimal size** field enter the maximum value that can be specified in the **Length** field of metadata used in mainframe jobs in the project. This can be 18 (the default) or 31.
14. From the **Null indicator location** drop-down list, select **Before column** or **After column** to specify the position of NULL indicators in mainframe column definitions.
15. In the **Null indicator value** field, specify the character used to indicate nullability of mainframe column definitions. NULL indicators must be single-byte, printable characters. Specify one of the following:
 - A single character value (1 is the default)
 - An ASCII code in the form of a three-digit decimal number from 000 to 255
 - An ASCII code in hexadecimal form of %Hnn or %hnn where 'nn' is a hexadecimal digit (0-9, a-f, A-F)
16. In the **Non-Null Indicator Value** field, specify the character used to indicate non-NULL column definitions in mainframe flat files. NULL indicators must be single-byte, printable characters. Specify one of the following:
 - A single character value (0 is the default)
 - An ASCII code in the form of a three-digit decimal number from 000 to 255

- An ASCII code in hexadecimal form of %Hnn or %hnn where 'nn' is a hexadecimal digit (0-9, a-f, A-F)
17. Click **OK** to save your changes.

Tunables page

The Tunables page allows you to set up caching details for hashed file stages and to set up row buffering to improve performance of server jobs.

Hashed file caching

When a Hashed File stage writes records to a hashed file, there is an option for the write to be cached rather than written to the hashed file immediately.

About this task

Similarly, when a Hashed File stage is reading a hashed file there is an option to pre-load the file to memory, which makes subsequent access very much faster and is used when the file is providing a reference link to a Transformer stage. (UniData Stages also offer the option of preloading files to memory, in which case the same cache size is used.)

The **Hashed file stage** area of the Tunables page lets you configure the size of the read and write caches.

Procedure

1. To specify the size of the read cache, enter a value between 0 and 999 in the Read cache size (MB) field. The value, which is in megabytes, defaults to 128.
2. To specify the size of the write cache, enter a value between 0 and 999 in the Write cache size (MB) field. The value, which is in megabytes, defaults to 128.
3. Click **OK** to save your changes.

Row buffering

The use of row buffering can greatly enhance performance in server jobs.

Select the **Enable row buffer** check box to enable this feature for the whole project.

There are two types of mutually exclusive row buffering:

- **In process.** You can improve the performance of most jobs by turning in-process row buffering on and recompiling the job. This allows connected active stages to pass data via buffers rather than row by row.
- **Inter process.** Use this if you are running server jobs on an SMP parallel system. This enables the job to run using a separate process for each active stage, which will run simultaneously on a separate processor.

You cannot use row-buffering of either sort if your job uses COMMON blocks in transform functions to pass data between stages. This is not recommended practice, and it is advisable to redesign your job to use row buffering rather than COMMON blocks.

When you have enabled row buffering, you can specify the following:

- **Buffer size.** Specifies the size of the buffer used by in-process or inter-process row buffering. Defaults to 128 Kb.

- **Timeout.** Only applies when inter-process row buffering is used. Specifies the time one process will wait to communicate with another via the buffer before timing out. Defaults to 10 seconds.

Parallel page

The parallel page allows you to specify certain defaults for parallel jobs in the project.

If you select the **Generated OSH visible for Parallel jobs in ALL projects** option, you will be able to view the code that is generated by parallel jobs at various points in the Designer and Director:

- In the Job Properties dialog box for parallel jobs.
- In the job run log message.
- When you use the View Data facility in the Designer.
- In the Table Definition dialog box.

Note that selecting this option enables this feature for all projects, not just the one currently selected.

The **Advanced runtime options for Parallel Jobs** field allows experienced Orchestra[®] users to enter parameters that are added to the OSH command line. Under normal circumstances this should be left blank. You can use this field to specify the `-nosortinsertion` or `-nopartinsertion` options. These prevent the automatic insertion of sort or partition operations where InfoSphere DataStage considers they are required. This applies to all jobs in the project.

Message Handler for Parallel Jobs allows you to specify a message handler for all the parallel jobs in this project. You define message handlers in the Director. They allow you to specify how certain warning or information messages generated by parallel jobs are handled. Choose one of the pre-defined handlers from the drop-down list.

The **Format defaults** area allows you to override the system default formats for dates, times, timestamps, and decimal separators. To change a default, clear the corresponding **System default** check box, then either select a new format from the drop down list or type in a new format.

Sequence page

Use the Sequence page to set compilation defaults for job sequences. You can optionally have InfoSphere DataStage add checkpoints to a job sequence so that, if part of the sequence fails, you do not necessarily have to start again from the beginning.

You can fix the problem and rerun the sequence from the point at which it failed. You can also specify that InfoSphere DataStage automatically handle failing jobs within a sequence (this means that you do not have to have a specific trigger for job failure).

The remaining options on the Sequence page allow you to specify that job sequences, by default, log a message in the sequence log if they run a job that finishes with warnings or fatal errors, or a command or routine that finishes with an error status. You can also have the log record a status report for a job immediately the job run finishes.

Deploying on remote systems

You can specify remote deployment of parallel jobs.

About this task

Where you are deploying parallel jobs on other, deployment-only systems, this page lets you:

- Provide a location for the remote deployment.
- Specify names for deployment directories.
- Specify further actions to be carried out at the end of a deployment compile.

Procedure

1. In the **Base directory name** field, provide a home directory location for deployment; in this directory there will be one directory for each job. This location has to be accessible from the InfoSphere Information Server engine machine, but does not have to be a disk local to that machine. Providing a location here enables the job deployment features.
2. In the **Deployed job directory template** field, optionally specify an alternative name for the deployment directory associated with a particular job. This field is used in conjunction with **Base directory name**. By default, if nothing is specified, the name corresponds to the internal script directory used on the InfoSphere Information Server engine project directory, *RT_SCjobnum*, where *jobnum* is the internal job number allocated to the job. Substitution strings provided are:

- %j - jobname
- %d - internal number

The simplest case is just "%j" - use the jobname. A prefix can be used, that is, "job_%j". The default corresponds to RT_SC%d.

3. In the **Custom deployment commands** field, optionally specify further actions to be carried out at the end of a deployment compile. You can specify UNIX programs and /or calls to user shell scripts as required.

This field uses the same substitution strings as the directory template. For example:

```
tar -cvf ../%j.tar * ; compress ../%j.tar
```

creates a compressed *.tar archive of the deployed job, named after the job.

Logs page

Use the logs page to control how the jobs in your project log information when they run.

Purging job log files

Every InfoSphere DataStage job has a log file, and every time you run a job, new entries are added to the log file. To prevent the files from becoming too large, they must be purged from time to time.

About this task

To prevent the files from becoming too large, they must be purged from time to time. You can set project-wide defaults for automatically purging job logs, or purge them manually. If you change the defaults in the Administrator client, the new

settings are reflected in the jobs in the project, unless a job has overridden the default settings (which can be done from the Director client), in which case it keeps the override values.

To set automatic purging for a project:

Procedure

1. Click the **Projects** tab in the Administrator window to move this page to the front.
2. Select the project.
3. Click **Properties**. The Project Properties window appears, with the General page displayed.
4. Click the **Logs** tab.
5. In the Logs page, select the **Auto-purge of job log** check box.
6. Select the Auto-purge action. You can purge jobs over the specified number of days old, or specify the number of jobs you want to retain in the log. For example, if you specify 10 job runs, entries for the last 10 job runs are kept.
7. Click **OK** to set the auto-purge policy. Auto-purging is applied to all new jobs created in the project. You can set auto-purging for existing jobs from the Clear Log window. Select **Job > Clear Log...** from the Director window to access this window.

Results

You can override automatic job log purging for an individual job by choosing **Job > Clear Log...** from the Director window.

You can delete job log files and job instances from the Job Status view in the Director client.

Chapter 5. Configuring projects for NLS

IBM InfoSphere DataStage has built-in National Language Support (NLS).

This means InfoSphere DataStage can:

- Process data in a wide range of languages
- Use local formats for dates, times, and money
- Sort data according to local rules

Using NLS, InfoSphere DataStage holds data in Unicode format. This is an international standard character set that contains many of the characters used in languages around the world. InfoSphere DataStage maps data to or from Unicode format as required.

Each InfoSphere DataStage project has a map and a locale assigned to it during installation. The map defines the character set that the project can use. The locale defines the local formats for dates, times, sorting order, and so on (sorting order only for parallel jobs), that the project should use. The InfoSphere DataStage client and engine components also have maps assigned to them during installation to ensure that data is transferred in the correct format.

InfoSphere DataStage has different mechanisms for implementing NLS for server and parallel jobs, and so you set map and locale details separately for the two types of job. Under normal circumstances, the two settings will match.

From the Administrator window, you can check which maps and locales were assigned during installation and change them as required.

Changing project maps

You can view or change a project map.

Procedure

1. Click the **Projects** tab in the Administrator window to move this page to the front.
2. Select the project.
3. Click **NLS...**. The Project NLS Settings window appears.
If the **NLS...** button is not active, you do not have NLS installed.
4. Choose whether you want to set the project map for server jobs or parallel jobs and choose the **Server Maps** or **Parallel Maps** tab accordingly.

Server job project maps

The **Default map name** field shows the current map that is used for server jobs in the project.

About this task

By default, the list shows only the maps that are loaded and ready to use in InfoSphere DataStage. You can examine the complete list of maps that are supplied with InfoSphere DataStage by clicking **Show all maps**.

To change the default map name for the project, click the map name you want to use, then click **OK**.

To install a map into InfoSphere DataStage, click **Install** to see additional options on the Maps page.

The **Available** list shows all the character set maps that are supplied with InfoSphere DataStage. The **Installed/loaded** list shows the maps that are currently installed. To install a map, select it from the **Available** list and click **Add**. The map is loaded into InfoSphere DataStage ready for use the next time the server is restarted. If you want to use the map immediately, you must restart the server engine.

To remove an installed map, select it from the **Installed/loaded** list and click **Remove**. The map is unloaded the next time the server is rebooted or the server engine is restarted.

Parallel job project maps

The **Default map name** field shows the current map that is used for parallel jobs in the project.

About this task

The list shows only the maps that are loaded and ready to use in InfoSphere DataStage. Double-click the map you want to make the default map.

Changing project locales

To view or change default project locales, having opened the Project NLS Settings Window, click the **Server Locales** tab or **Parallel Locales** tab as appropriate.

Server job locales

You can view or change default locales for all server jobs in a project by using the IBM InfoSphere DataStage Administrator client. You choose a locale from the drop down list of installed locales on the **Server Locales** tab in the **Project NLS Settings** Window.

About this task

This page shows fields for the default project locales in five categories:

- **Time/Date** - The format for dates and times, for example, 31 Dec 1999 or 12/31/99 are two ways of expressing the same date that might be used in different locales.
- **Numeric** - The format used for numbers, including the thousands separator and radix (decimal) delimiter.
- **Currency** - The format for monetary strings, including the type and position of the currency sign (\$, £, €, ¥, and so on).
- **CType** - The format for character types. This includes defining which characters can be uppercase or lowercase characters in a language.
- **Collate** - The sort order used for a language.

By default, each field has a drop-down list of the locales that are loaded and ready to use. To change a locale in any category, select the locale you want from the drop-down list. Click **OK** when you have completed your changes. You can

examine the complete list of locales that are supplied with InfoSphere DataStage by clicking **Show all locales**, then clicking a category drop-down list. These locales must be installed and loaded into InfoSphere DataStage before you can use them.

Installing and loading locales

To install a locale, click **Install** to display further options on the Locales page.

About this task

The **Available** list shows all the locales that are supplied with InfoSphere DataStage. The **Installed/loaded** list shows the locales that are currently installed. To install a locale, select it from the **Available** list and click **Add**. The locale is loaded into InfoSphere DataStage ready for use the next time the server engine is restarted. If you want to use the locale immediately, you can restart the server engine.

To remove an installed locale, select it from the **Installed/loaded** list and click **Remove**. The locale is unloaded the next time the server engine is restarted.

Parallel job locales

Only the collate category is used for parallel jobs. Choose a locale from the drop down list of installed locales.

About this task

The **Browse** button allows you to browse for text files that define other collation sequences.

Client and server maps

When you installed the InfoSphere Information Server engine, you specified the language that you wanted InfoSphere DataStage to support. InfoSphere DataStage automatically sets the language supported on the InfoSphere DataStage clients to match what you specified for the server.

About this task

If you access the InfoSphere Information Server engine from a different client, data might not be mapped correctly between the client and the engine. To ensure that the data is mapped correctly, you might need to change the client maps. You can view the current mapping by doing the following steps.

Procedure

1. Click the **General** tab on the Administrator window to move this page to the front.
2. Click **NLS...** The General NLS Settings window opens.

Results

The **Current ANSI code page** field is informational only, and contains the current Microsoft code page of the client. The code page is independent of the current project or engine. The **Client/Server map in use** field shows the name of the map being used on the engine computer for the current client session. The list shows all loaded maps.

If you select a map and click **Apply**, InfoSphere DataStage attempts to set this map for all clients connecting to the current server that use the code page shown. The mapping is tested, and might be rejected if it is not appropriate.

To install further maps into InfoSphere DataStage, click **Install** to show further options on the Client page.

InfoSphere DataStage uses special maps for client/server communication, with names ending in "-CS" (for Client Server). Always choose one of these maps for this purpose.

The **Available** list shows all the character set maps that are supplied with InfoSphere DataStage. The **Installed/loaded** list shows the maps that are currently installed. To install a map, select it from the **Available** list and click **Add**. The map is loaded into InfoSphere DataStage ready for use at the next time the server is restarted. If you want to use the map immediately, you can restart the server engine.

To remove an installed map, select it from the **Installed/loaded** list and click **Remove**. The map is unloaded the next time the server is rebooted or the server engine is restarted.

Appendix A. Product accessibility

You can get information about the accessibility status of IBM products.

The IBM InfoSphere Information Server product modules and user interfaces are not fully accessible.

For information about the accessibility status of IBM products, see the IBM product accessibility information at http://www.ibm.com/able/product_accessibility/index.html.

Accessible documentation

Accessible documentation for InfoSphere Information Server products is provided in an information center. The information center presents the documentation in XHTML 1.0 format, which is viewable in most web browsers. Because the information center uses XHTML, you can set display preferences in your browser. This also allows you to use screen readers and other assistive technologies to access the documentation.

The documentation that is in the information center is also provided in PDF files, which are not fully accessible.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

Appendix B. Contacting IBM

You can contact IBM for customer support, software services, product information, and general information. You also can provide feedback to IBM about products and documentation.

The following table lists resources for customer support, software services, training, and product and solutions information.

Table 1. IBM resources

Resource	Description and location
IBM Support Portal	You can customize support information by choosing the products and the topics that interest you at www.ibm.com/support/entry/portal/Software/Information_Management/InfoSphere_Information_Server
Software services	You can find information about software, IT, and business consulting services, on the solutions site at www.ibm.com/businesssolutions/
My IBM	You can manage links to IBM Web sites and information that meet your specific technical support needs by creating an account on the My IBM site at www.ibm.com/account/
Training and certification	You can learn about technical training and education services designed for individuals, companies, and public organizations to acquire, maintain, and optimize their IT skills at http://www.ibm.com/training
IBM representatives	You can contact an IBM representative to learn about solutions at www.ibm.com/connect/ibm/us/en/

Appendix C. Accessing the product documentation

Documentation is provided in a variety of formats: in the online IBM Knowledge Center, in an optional locally installed information center, and as PDF books. You can access the online or locally installed help directly from the product client interfaces.

IBM Knowledge Center is the best place to find the most up-to-date information for InfoSphere Information Server. IBM Knowledge Center contains help for most of the product interfaces, as well as complete documentation for all the product modules in the suite. You can open IBM Knowledge Center from the installed product or from a web browser.

Accessing IBM Knowledge Center

There are various ways to access the online documentation:

- Click the **Help** link in the upper right of the client interface.
- Press the F1 key. The F1 key typically opens the topic that describes the current context of the client interface.

Note: The F1 key does not work in web clients.

- Type the address in a web browser, for example, when you are not logged in to the product.

Enter the following address to access all versions of InfoSphere Information Server documentation:

```
http://www.ibm.com/support/knowledgecenter/SSZJPZ/
```

If you want to access a particular topic, specify the version number with the product identifier, the documentation plug-in name, and the topic path in the URL. For example, the URL for the 11.3 version of this topic is as follows. (The ⇒ symbol indicates a line continuation):

```
http://www.ibm.com/support/knowledgecenter/SSZJPZ_11.3.0/⇒  
com.ibm.swg.im.iis.common.doc/common/accessingiidoc.html
```

Tip:

The knowledge center has a short URL as well:

```
http://ibm.biz/knowctr
```

To specify a short URL to a specific product page, version, or topic, use a hash character (#) between the short URL and the product identifier. For example, the short URL to all the InfoSphere Information Server documentation is the following URL:

```
http://ibm.biz/knowctr#SSZJPZ/
```

And, the short URL to the topic above to create a slightly shorter URL is the following URL (The ⇒ symbol indicates a line continuation):

```
http://ibm.biz/knowctr#SSZJPZ_11.3.0/com.ibm.swg.im.iis.common.doc/⇒  
common/accessingiidoc.html
```

Changing help links to refer to locally installed documentation

IBM Knowledge Center contains the most up-to-date version of the documentation. However, you can install a local version of the documentation as an information center and configure your help links to point to it. A local information center is useful if your enterprise does not provide access to the internet.

Use the installation instructions that come with the information center installation package to install it on the computer of your choice. After you install and start the information center, you can use the **iisAdmin** command on the services tier computer to change the documentation location that the product F1 and help links refer to. (The `⇒` symbol indicates a line continuation):

Windows

```
IS_install_path\ASBServer\bin\iisAdmin.bat -set -key ⇒  
com.ibm.iis.infocenter.url -value http://<host>:<port>/help/topic/
```

AIX® Linux

```
IS_install_path/ASBServer/bin/iisAdmin.sh -set -key ⇒  
com.ibm.iis.infocenter.url -value http://<host>:<port>/help/topic/
```

Where `<host>` is the name of the computer where the information center is installed and `<port>` is the port number for the information center. The default port number is 8888. For example, on a computer named `server1.example.com` that uses the default port, the URL value would be `http://server1.example.com:8888/help/topic/`.

Obtaining PDF and hardcopy documentation

- The PDF file books are available online and can be accessed from this support document: <https://www.ibm.com/support/docview.wss?uid=swg27008803&wv=1>.
- You can also order IBM publications in hardcopy format online or through your local IBM representative. To order publications online, go to the IBM Publications Center at <http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss>.

Appendix D. Providing feedback on the product documentation

You can provide helpful feedback regarding IBM documentation.

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

- To provide a comment about a topic in IBM Knowledge Center that is hosted on the IBM website, sign in and add a comment by clicking **Add Comment** button at the bottom of the topic. Comments submitted this way are viewable by the public.
- To send a comment about the topic in IBM Knowledge Center to IBM that is not viewable by anyone else, sign in and click the **Feedback** link at the bottom of IBM Knowledge Center.
- Send your comments by 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, include the location of the text (for example, a title, a table number, or a page number).

Notices and trademarks

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

Notices

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 character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © 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.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session or persistent cookies. If a product or component is not listed, that product or component does not use cookies.

Table 2. Use of cookies by InfoSphere Information Server products and components

Product module	Component or feature	Type of cookie that is used	Collect this data	Purpose of data	Disabling the cookies
Any (part of InfoSphere Information Server installation)	InfoSphere Information Server web console	<ul style="list-style-type: none"> • Session • Persistent 	User name	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
Any (part of InfoSphere Information Server installation)	InfoSphere Metadata Asset Manager	<ul style="list-style-type: none"> • Session • Persistent 	No personally identifiable information	<ul style="list-style-type: none"> • Session management • Authentication • Enhanced user usability • Single sign-on configuration 	Cannot be disabled

Table 2. Use of cookies by InfoSphere Information Server products and components (continued)

Product module	Component or feature	Type of cookie that is used	Collect this data	Purpose of data	Disabling the cookies
InfoSphere DataStage	Big Data File stage	<ul style="list-style-type: none"> • Session • Persistent 	<ul style="list-style-type: none"> • User name • Digital signature • Session ID 	<ul style="list-style-type: none"> • Session management • Authentication • Single sign-on configuration 	Cannot be disabled
InfoSphere DataStage	XML stage	Session	Internal identifiers	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
InfoSphere DataStage	IBM InfoSphere DataStage and QualityStage Operations Console	Session	No personally identifiable information	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
InfoSphere Data Click	InfoSphere Information Server web console	<ul style="list-style-type: none"> • Session • Persistent 	User name	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
InfoSphere Data Quality Console		Session	No personally identifiable information	<ul style="list-style-type: none"> • Session management • Authentication • Single sign-on configuration 	Cannot be disabled
InfoSphere QualityStage Standardization Rules Designer	InfoSphere Information Server web console	<ul style="list-style-type: none"> • Session • Persistent 	User name	<ul style="list-style-type: none"> • Session management • Authentication 	Cannot be disabled
InfoSphere Information Governance Catalog		<ul style="list-style-type: none"> • Session • Persistent 	<ul style="list-style-type: none"> • User name • Internal identifiers • State of the tree 	<ul style="list-style-type: none"> • Session management • Authentication • Single sign-on configuration 	Cannot be disabled
InfoSphere Information Analyzer	Data Rules stage in the InfoSphere DataStage and QualityStage Designer client	Session	Session ID	Session management	Cannot be disabled

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

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

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Microsoft, Windows and Windows NT are trademarks of Microsoft 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.

Java[™] and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The United States Postal Service owns the following trademarks: CASS, CASS Certified, DPV, LACS^{Link}, ZIP, ZIP + 4, ZIP Code, Post Office, Postal Service, USPS and United States Postal Service. IBM Corporation is a non-exclusive DPV and LACS^{Link} licensee of the United States Postal Service.

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

Index

A

- Add Project dialog box 7
- adding
 - projects 7
- Administrator 3
- Administrator client 1
- Administrator window 23
- auto-purge defaults 24

C

- changing
 - locales 26
 - maps 25
- character types 26
- Cleanup Resources option 12
- Clear Status File option 12
- client and server maps 27
- code page 27
- Command Interface dialog box 5
- Command Output window 5
- commands
 - InfoSphere Information Server engine 5
- currency signs 26
- customer support
 - contacting 31

D

- DataStage Administration window 25
 - General page 3
- DataStage Director
 - enabling job administration in 11
- DataStage Repository Import dialog box 9
- DataStage users, setting up 17
- dates 26
- DatsStage engine tracing 18
- deleting projects 8
- deployment systems 23
- dialog boxes
 - Add Project 7
 - Command Interface 5
 - DataStage Repository Import 9

E

- enabling
 - job administration in Director 11
 - tracing on the server 18
- enabling tracing on the DataStage engine 18
- environment variables
 - exporting to a file 14
 - importing from a file 15

F

- files
 - cleaning up 11
 - log 23, 24
 - trace 18

G

- General NLS Settings window 27

I

- InfoSphere Information Server engine
 - issuing commands 5
- InfoSphere® DataStage® 23
- installing
 - locales 27
 - maps 28
- issuing InfoSphere Information Server engine commands 5

J

- job administration 11
- job properties, mainframe 19
- jobs
 - cleaning up files 11
 - clearing status file 12
 - log files 23
 - scheduling 19

L

- legal notices 37
- locales 26
 - changing 26
 - installing 27
- log files
 - automatic purging 23
 - manual purging 24
 - setting up operator's view 18
- Logs page 23

M

- mainframe job properties, setting 19
- maps 25
 - changing 25
 - client and server 27
 - installing 28
- message handlers 22
- moving projects 8

N

- National Language Support, see NLS 25
- NLS (National Language Support)
 - client and server maps 27

NLS (National Language Support)
(continued)

- locales 26
- maps 25
- overview 25
- NLS option
 - project settings 25
 - server settings 27
- NT system authority 19
- numbers 26

O

- operational metadata 16

P

- Permissions page 17
- product accessibility
 - accessibility 29
- product documentation
 - accessing 33
- Project NLS Settings window 25
- Project Properties window 11
- project properties, editing 11
- projects 23
 - adding 7
 - deleting 8
 - moving 8
- properties
 - mainframe 19
 - project 11
- purging, automatic 24

R

- radix 26
- remote deployment 23

S

- Schedule page of Project Properties window 19
- scheduling jobs 19
- server timeout, setting 4
- server tracing 18
- setting
 - auto-purge defaults 24
 - mainframe job properties 19
 - server timeout 4
 - user name for scheduled jobs 19
- setting up DataStage users 16
- software services
 - contacting 31
- sort order 26
- Suite Administrator 1
- support
 - customer 31

T

- testing user name for scheduled jobs 19
- thousands separator 26
- times 26
- tracing server activity 18
- trademarks
 - list of 37

U

- Unicode 25
- user roles 17

V

- View trace file window 19
- viewing
 - locales 26
 - maps 25

W

- Web console
 - Administration tab 3
- windows
 - Command Output 5
 - DataStage Administration 3, 25
 - General NLS Settings 27
 - Project NLS Settings 25
 - Project Properties 11
 - View trace file 19
- Windows Schedule service 19
- workload management queue
 - setting 16



Printed in USA

SC19-4254-00

