

# **Informix DataStage**

## **Administrator's Guide**

Version 3.6  
October 1999  
Part No. 000-6467

Published by Informix® Press

Informix Corporation  
4100 Bohannon Drive  
Menlo Park, CA 94025-1032

© 1999 Informix Corporation. All rights reserved. The following are trademarks of Informix Corporation or its affiliates, one or more of which may be registered in the United States or other jurisdictions:

Answers OnLine™; C-ISAM®; Client SDK™; DataBlade®; Data Director™; Decision Frontier™; Dynamic Scalable Architecture™; Dynamic Server™; Dynamic Server™, Developer Edition™; Dynamic Server™ with Advanced Decision Support Option™; Dynamic Server™ with Extended Parallel Option™; Dynamic Server™ with MetaCube®; Dynamic Server™ with Universal Data Option™; Dynamic Server™ with Web Integration Option™; Dynamic Server™, Workgroup Edition™; Dynamic Virtual Machine™; Enterprise Decision Server™; Formation™; Formation Architect™; Formation Flow Engine™; Frameworks for Business Intelligence™; Frameworks Technology™; Gold Mine Data Access®; i.Reach™; i.Sell™; Illustra®; Informix®; Informix® 4GL; Informix® COM Adapter™; Informix® Informed Decisions™; Informix® InquireSM; Informix® Internet Foundation.2000™; InformixLink®; Informix® Red Brick® Decision Server™; Informix Session Proxy™; Informix® Vista™; InfoShelf™; Interforum™; I-Spy™; Mediazation™; MetaCube®; NewEra™; Office Connect™; ON-Bar™; OnLine Dynamic Server™; OnLine/Secure Dynamic Server™; OpenCase®; Orca™; PaVER™; Red Brick® and Design; Red Brick® Data Mine™; Red Brick® Mine Builder™; Red Brick® Decisionscape™; Red Brick® Ready™; Red Brick Systems®; Regency Support®; Rely on Red BrickSM; RISQL®; Solution DesignSM; STARindex™; STARjoin™; SuperView®; TARGETindex™; TARGETjoin™; The Data Warehouse Company®; Universal Data Warehouse Blueprint™; Universal Database Components™; Universal Web Connect™; ViewPoint®; Visionary™; Web Integration Suite™. The Informix logo is registered with the United States Patent and Trademark Office. The DataBlade logo is registered with the United States Patent and Trademark Office.

#### GOVERNMENT LICENSE RIGHTS

Software and documentation acquired by or for the US Government are provided with rights as follows:

(1) if for civilian agency use, with rights as restricted by vendor's standard license, as prescribed in FAR 12.212;  
(2) if for Dept. of Defense use, with rights as restricted by vendor's standard license, unless superseded by a negotiated vendor license, as prescribed in DFARS 227.7202. Any whole or partial reproduction of software or documentation marked with this legend must reproduce this legend.

# Table of Contents

## **Preface**

|                                   |      |
|-----------------------------------|------|
| Organization of This Manual ..... | vii  |
| Documentation Conventions .....   | vii  |
| DataStage Documentation .....     | viii |

## **Chapter 1. Configuring DataStage**

|  |      |
|--|------|
| DataStage NLS .....                          | 1-1  |
| Who Can Administer DataStage? .....          | 1-1  |
| The DataStage Administrator .....            | 1-2  |
| Project Properties .....                     | 1-3  |
| Changing License Details .....               | 1-4  |
| DataStage User Administration .....          | 1-4  |
| DataStage User Categories .....              | 1-5  |
| Changing User Group Assignments .....        | 1-5  |
| Setting a User Name for Scheduled Jobs ..... | 1-6  |
| DataStage Project Administration .....       | 1-7  |
| Deleting Projects .....                      | 1-8  |
| Moving Projects .....                        | 1-8  |
| Adding Projects .....                        | 1-10 |
| Cleaning Up Project Files .....              | 1-10 |
| Job Log File Maintenance .....               | 1-11 |
| Setting Auto-Purge Defaults .....            | 1-11 |
| Purging Log Files Manually .....             | 1-12 |
| Setting Server Properties .....              | 1-13 |
| Setting the Server Timeout .....             | 1-13 |
| Enabling Tracing on the Server .....         | 1-14 |
| DataStage Job Administration .....           | 1-15 |

|   |      |
|---|------|
| Setting Job Parameter Defaults .....    | 1-16 |
| Changing Encrypted Parameters .....     | 1-18 |
| Changing a Default List of Values ..... | 1-19 |
| Customizing the Tools Menu .....        | 1-20 |
| Adding a Tool .....                     | 1-20 |
| Changing the Tools Menu Order .....     | 1-21 |
| Adding a Separator .....                | 1-21 |
| Issuing UniVerse Commands .....         | 1-22 |

## **Chapter 2. Windows NT Server Administration**

|   |     |
|---|-----|
| Starting the Schedule Service .....           | 2-1 |
| Starting and Stopping the Server Engine ..... | 2-2 |
| Accessing ODBC Databases .....                | 2-2 |
| Accessing Remote UniVerse Systems .....       | 2-3 |
| Installing DataStage Packages .....           | 2-4 |
| Package Installer Command Line Options .....  | 2-4 |
| Installing MetaBrokers .....                  | 2-5 |

## **Chapter 3. UNIX Server Administration**

|  |     |
|--|-----|
| Installation Requirements .....                    | 3-1 |
| Tuning UNIX for DataStage .....                    | 3-2 |
| DataStage Engine Configuration .....               | 3-3 |
| System with UniVerse Installed .....               | 3-3 |
| System with UniVerse and GCI Installed .....       | 3-3 |
| The RPC Daemon ( <i>unirpcd</i> ) .....            | 3-4 |
| Environment Requirements for <i>unirpcd</i> .....  | 3-4 |
| Starting and Stopping the Server Engine .....      | 3-5 |
| Accessing Remote UniVerse Systems .....            | 3-5 |
| Accessing ODBC Databases .....                     | 3-6 |
| Installing Intersolv Drivers .....                 | 3-7 |
| Adding Data Sources to <i>uvodbc.config</i> .....  | 3-7 |
| Configuring the ODBC Driver Manager .....          | 3-7 |
| Using <i>relink.uvlibs</i> to Maintain Links ..... | 3-8 |
| Trial Version of ODBC Drivers .....                | 3-9 |

|                                      |      |
|--------------------------------------|------|
| Installing DataStage Packages .....  | 3-9  |
| BCPLoad Plug-In Restrictions .....   | 3-10 |
| Server Restrictions .....            | 3-10 |
| Connecting to the Server .....       | 3-10 |
| Scheduling Jobs .....                | 3-10 |
| Accessing Remote Files via NFS ..... | 3-11 |

## **Chapter 4. NLS Configuration**

|                                      |     |
|--------------------------------------|-----|
| Changing Project Maps .....          | 4-1 |
| Changing Project Locales .....       | 4-3 |
| Installing and Loading Locales ..... | 4-4 |
| Client/Server Maps .....             | 4-5 |

## **Chapter 5. Troubleshooting**

|   |     |
|---|-----|
| Error Messages .....                              | 5-1 |
| Scheduled Jobs .....                              | 5-1 |
| Windows NT Scheduling Problems .....              | 5-2 |
| UNIX Scheduling Problems .....                    | 5-3 |
| Job Termination Problems .....                    | 5-4 |
| ODBC Connection Problems on UNIX .....            | 5-4 |
| Check the Shared Library Environment .....        | 5-4 |
| Check Symbolic Links .....                        | 5-4 |
| DataStage Client to UNIX Server Connections ..... | 5-5 |
| Connecting to UniData Data Sources .....          | 5-5 |

## **Index**



# Preface

Informix DataStage is a tool set for designing, developing, and running jobs that populate tables in a data warehouse or data mart.

This manual describes DataStage administration, including:

- How to use the DataStage Administrator
- Essential system administration tasks that must be performed through the operating system
- Other housekeeping tasks, such as customizing menus and setting defaults

## Organization of This Manual

This manual contains the following:

Chapter 1 describes how to configure DataStage to suit the needs of your system.

Chapter 2 describes DataStage administration on a Windows NT server.

Chapter 3 describes DataStage administration on a UNIX server.

Chapter 4 describes how to use National Language Support (NLS) with DataStage.

Chapter 5 outlines troubleshooting procedures to solve possible problems.

## Documentation Conventions

This manual uses the following conventions:

| <b>Convention</b> | <b>Usage</b>   |
|-------------------|--|
| <b>Bold</b>       | In syntax, bold indicates commands, function names, keywords, and options that must be input exactly as shown. In text, bold indicates keys to press, function names, and menu selections. |

| Convention          | Usage   |
|---------------------|---|
| <i>Italic</i>       | In syntax, italic indicates information that you supply. In text, italic also indicates operating system commands and options, filenames, and pathnames.  |
| Courier             | Courier indicates examples of source code and system output.  |
| <b>Courier Bold</b> | In examples, courier bold indicates characters that the user types or keys the user presses (for example, <Return>).  |
| ➤                   | A right arrow between menu options indicates you should choose each option in sequence. For example, “Choose <b>File</b> ➤ <b>Exit</b> ” means you should choose <b>File</b> from the menu bar, then choose <b>Exit</b> from the File pull-down menu. |

## DataStage Documentation

DataStage documentation includes the following:

***DataStage Developer’s Guide:*** This guide describes the DataStage Manager and Designer, and how to create, design, and develop a DataStage application.

***DataStage Operator’s Guide:*** This guide describes the DataStage Director and how to validate, schedule, run, and monitor DataStage applications.

***DataStage Administrator’s Guide:*** This guide describes DataStage setup, routine housekeeping, and administration.

These guides are also available online in PDF format. You can read them with the Adobe Acrobat Reader supplied with DataStage. See *DataStage Installation Instructions* in the DataStage CD jewel case for details on installing the manuals and the Adobe Acrobat Reader.

# 1

## Configuring DataStage

This chapter describes how to set up DataStage when it is first installed, and some routine administration tasks, including:

- Changing license details
- Setting up DataStage users
- Deleting, moving, and adding DataStage projects
- Cleaning up project files
- Purging job log files
- Setting the timeout interval on the server computer
- Tracing server activity
- Adding entries to the Tools menu
- Setting job parameter defaults
- Issuing UniVerse commands from the Administration client

### DataStage NLS

If you installed DataStage NLS (National Language Support), DataStage will support the language you specified during the install without any further configuration. But if your requirements change, you can reconfigure NLS to support different languages. For more information, see Chapter 4, “NLS Configuration.”

### Who Can Administer DataStage?

You can do most of the administration tasks described in this chapter if you have been defined as a DataStage Developer. You do not need to have specific administration rights. However, to change license details or to set user group assignments you do need to have logged on to DataStage using a user name that gives you administrator status.

*For Windows NT servers:* You must be a member of the Windows NT Administrators group.

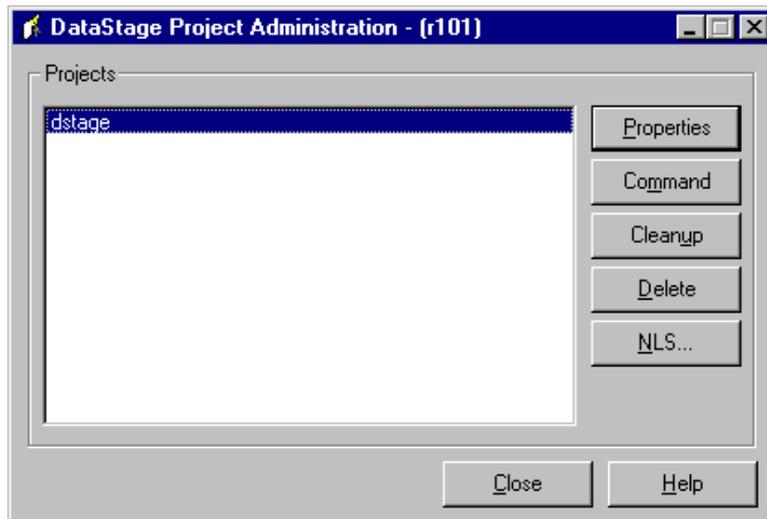
*For UNIX servers:* You must be *root* or superuser. You must also be *root* or superuser to carry out UNIX server administration tasks described in Chapter 3, “UNIX Server Administration.”

## The DataStage Administrator

Most DataStage configuration tasks are carried out using the DataStage Administrator, a client program provided with DataStage. To access the DataStage Administrator:

1. From the Ardent DataStage program folder, choose **DataStage Administrator**.
2. Log on to the server. If you do so as an Administrator (for Windows NT servers) or as *root* (for UNIX servers) you have unlimited administrative rights; otherwise your rights are restricted as described in the previous section.

The DataStage Project Administration dialog box appears:

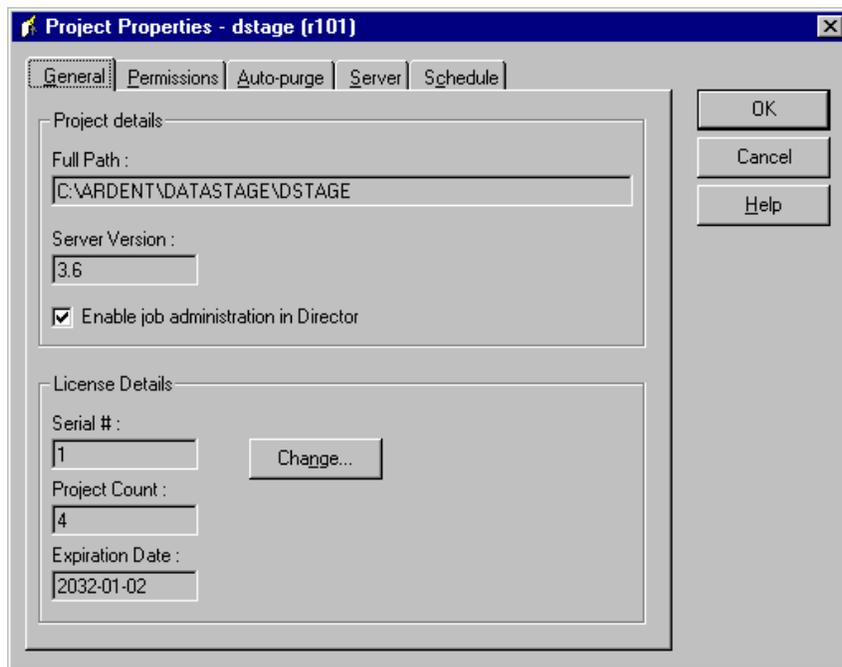


This is the starting point for many procedures described in this manual.

## Project Properties

You will often use the Project Properties window to administer DataStage projects. The descriptions of the individual procedures in the manual explain the relevant parts of the Project Properties window, but you will find it helpful to be aware of what you are going to see.

To display the Project Properties window, click the **Properties** button in the Project Administration dialog box.



You may not see all the tabs across the top of the window that are shown here.

The **Permissions** tab is visible only if you have logged on to DataStage using a name that gives you administrator status, and the **Schedule** tab is visible only if you have logged on to a Windows NT server.

The **Change...** button on the **General** page is enabled only if you have administrator status.

You use each page in the Project Properties window to do the following:

- **General.** View project and licence details and relicense the DataStage server.
- **Permissions.** Assign user categories to operating system user groups.
- **Auto-purge.** Define a project-wide setting for auto-purge of the job log.
- **Server.** Set the timeout on the connection between the server and client, and enable or disable tracing on the server.
- **Schedule.** Set up a user name and password to use for running scheduled DataStage jobs.

## Changing License Details

You can update your DataStage license details from the DataStage Administrator provided that your current license has not expired. Once the license expires, you can only relicense from the installation program.

To relicense DataStage:

1. From the DataStage Project Administration dialog box, click **Properties**. The **General** page on the Project Properties window appears.  
The **Change...** button is enabled only if you have administrator status (see “Who Can Administer DataStage?” on page 1-1).
2. Click **Change...** . The Change License Details dialog box appears.
3. Enter the new details exactly as shown on the Ardent license authorization.
4. Click **OK** to relicense the DataStage server.

## DataStage User Administration

This section describes:

- DataStage user categories and how to change the assignment of these categories to operating system user groups
- How to set up a user name and password to use for running scheduled DataStage jobs

## DataStage User Categories

To prevent unauthorized access to DataStage projects, you must assign the users on your system to the appropriate DataStage user category. To do this, you must have administrator status (see “Who Can Administer DataStage?” on page 1-1).

There are three categories of DataStage user:

- DataStage Developer, who has full access to all areas of a DataStage project
- DataStage Operator, who has permission to run and manage released DataStage jobs
- <None>, who does not have permission to log on to DataStage

You cannot assign individual users to these categories. You have to assign the operating system user group to which the user belongs. For example, a user with the user ID *peter* belongs to a user group called *clerks*. To give DataStage Operator status to user *peter*, you must assign the *clerks* user group to the DataStage Operator category.

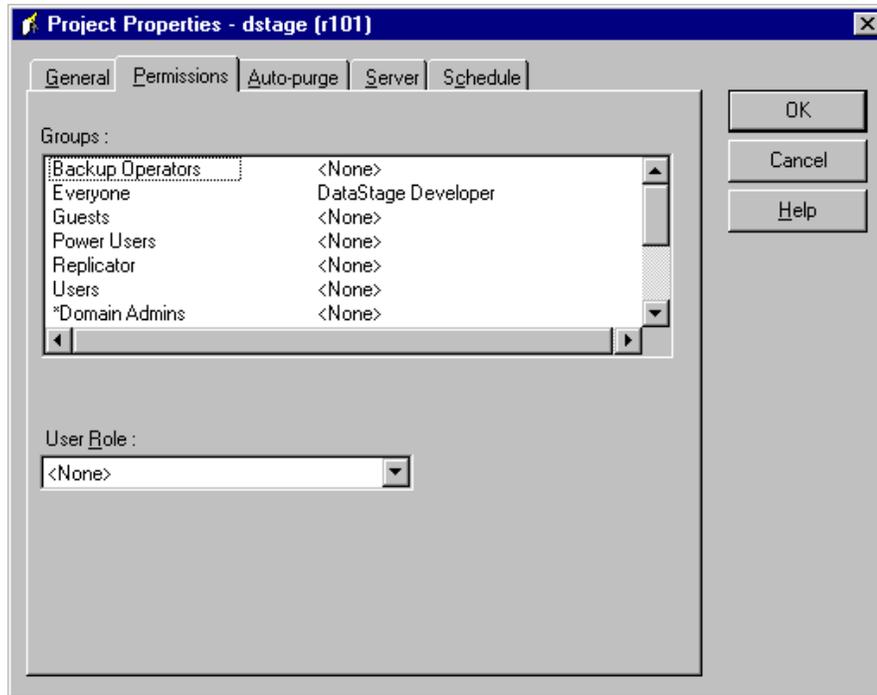
**Note:** When you first install DataStage, all users are assigned the category DataStage Developer by default. This means that every user has full access to DataStage. You must change the user group assignments to ensure they meet your needs.

You can prevent members of a group from logging on to DataStage by assigning the group to the <None> category.

## Changing User Group Assignments

To change the current user group assignments:

1. From the DataStage Project Administration dialog box (see page 1-2), click **Properties**. The Project Properties window appears. The **Permissions** tab is enabled only if you have administrator status.
2. Click the **Permissions** tab to move the **Permissions** page to the front as shown in the following screen:



The first column in the **Groups** list box shows all the operating system user groups on the server. On UNIX servers, only user groups listed in `/etc/groups` are shown in the list. Domain groups from the primary domain server are preceded by an asterisk (\*).

The second column shows the DataStage user category assigned to the operating system user group.

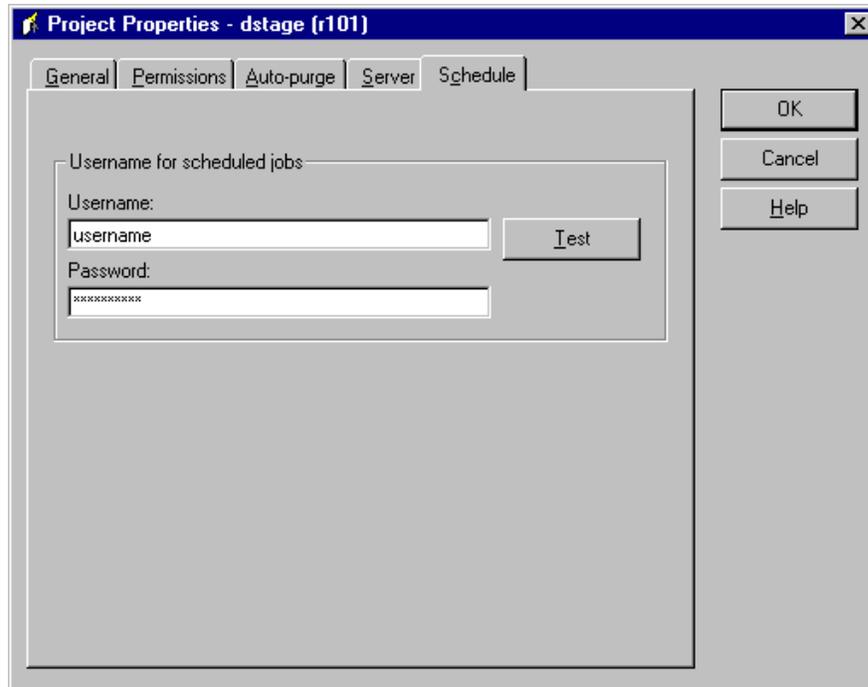
3. To change the DataStage user category for a user group, select the user group in the list box, then choose the category required from the drop-down list.
4. Repeat with other groups as required, then click **OK**.

## Setting a User Name for Scheduled Jobs

*Windows NT servers only.* DataStage uses the Windows NT Schedule service to schedule jobs. This means that by default the job runs under the user name of the Schedule service, which defaults to NT system authority. You may 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 follows:

1. From the Project Properties window, click the **Schedule** tab to move the **Schedule** page to the front, as shown in the following screen:



2. Enter the user name and password you want to use to run the scheduled jobs.
3. Click **Test** to test that the user name and password can be used successfully. This involves scheduling and running a command on the server, so the test may take some time to complete.
4. Click **OK** to save the user name and password.

## DataStage Project Administration

You must be careful when you delete or move DataStage projects. Always do so through the DataStage Administrator to ensure that licensing information is updated correctly.

## Deleting Projects

To delete a project from the server:

1. Ensure that no users are connected to the project you want to delete. DataStage returns an error message if you try to delete a project while another user is connected to it.
2. From the DataStage Project Administration dialog box, select the project you want to delete, then click **Delete**.
3. You are prompted to confirm the deletion. If you confirm, the project is deleted immediately.

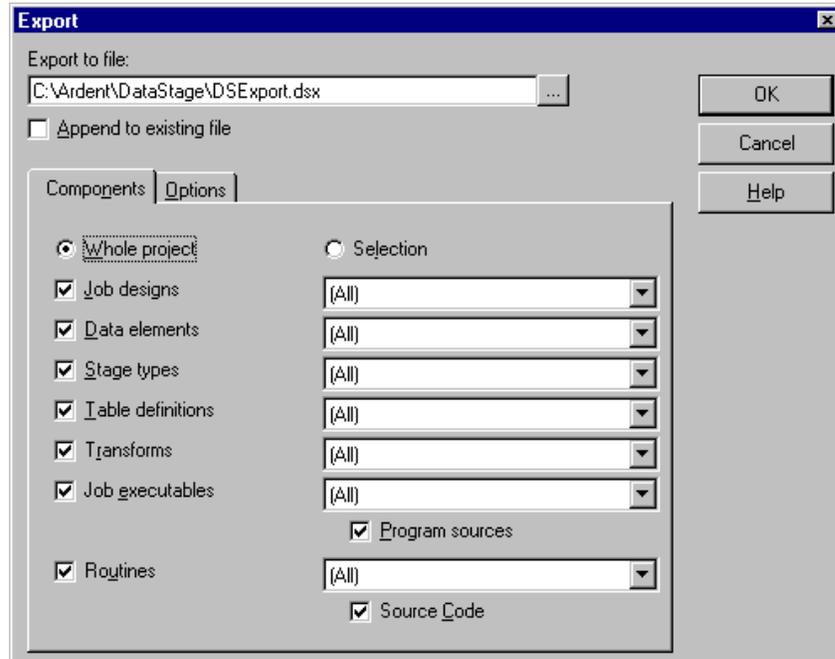
## Moving Projects

**Note:** This procedure can be used only on systems that have the DataStage Developer's Edition installed.

To move a DataStage project, you must export it to a file and then import it to the new location, as described in the following steps:

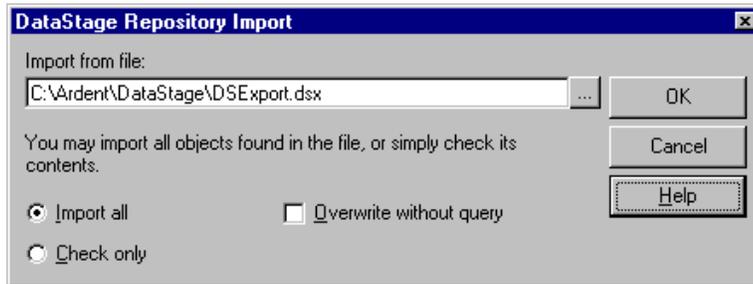
1. Start the DataStage Director in the project you want to move, and check that no jobs are running.
2. Choose **Tools ► Run Manager...** . The DataStage Manager window appears.
3. Close the DataStage Director to disconnect from the project.

4. Choose **Tools** ► **Export...** . The Export dialog box appears:



5. Enter a file name to hold the exported project.
6. Click the **Whole project** option button.
7. Click **OK**. The project is exported to the file.
8. Create a new project in the required location by running the DataStage installation program.  
**Note:** If you do not have enough licenses to create a new project, delete the old project from the DataStage Project Administration dialog box before you run the server installation program.
9. When the installation program finishes, open the DataStage Manager in the new project.

10. Choose **Tools ► Import ► DataStage Components...** . The DataStage Repository Import dialog box appears:



11. Enter the file name you used for the exported project in step 5, then click **OK**. The project is imported to the new location.

## Adding Projects

Each DataStage project must be licensed. If you did not create all the projects specified in your license during the installation process, you can repeat the server installation from the DataStage Master Setup screen and add new projects. Contact your DataStage supplier for an upgrade if you require more projects than your license allows.

## Cleaning Up Project Files

You can tidy the project by removing internal project files that are no longer required because:

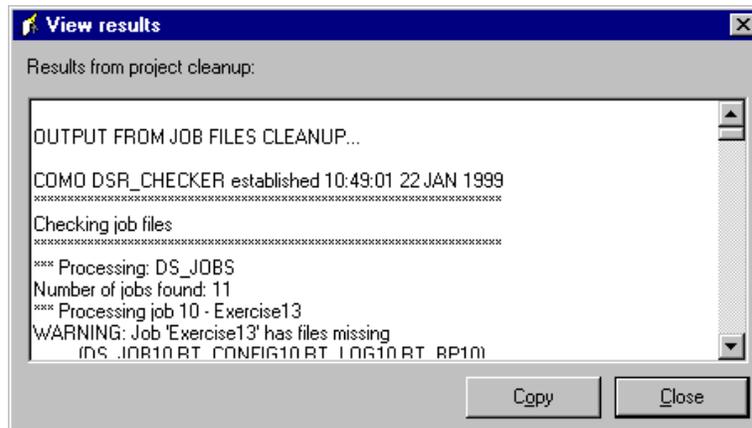
- Their job has been deleted
- They have an invalid job number

The cleanup process also checks the integrity of files in the DataStage Repository by recreating the secondary indexes for these files.

To clean up the project files:

1. Ensure that no other DataStage clients are logged on to the project you want to clean up, and that no jobs are running.
2. From the DataStage Project Administration dialog box (see page 1-2) select a project, then click **Cleanup**. You are prompted to confirm that you want to clean up the project.
3. Click **Yes** to proceed with the cleanup.

4. When the cleanup has ended, you are asked whether you want to view the results. Click **Yes** to view the report of the cleanup. The View results window appears:



5. To copy the entire report to the Clipboard, click **Copy**.  
Alternatively, to copy part of the report, select a block of text in the View results window, then click **Copy**.
6. Click **Close** to close the View results window.

## Job Log File Maintenance

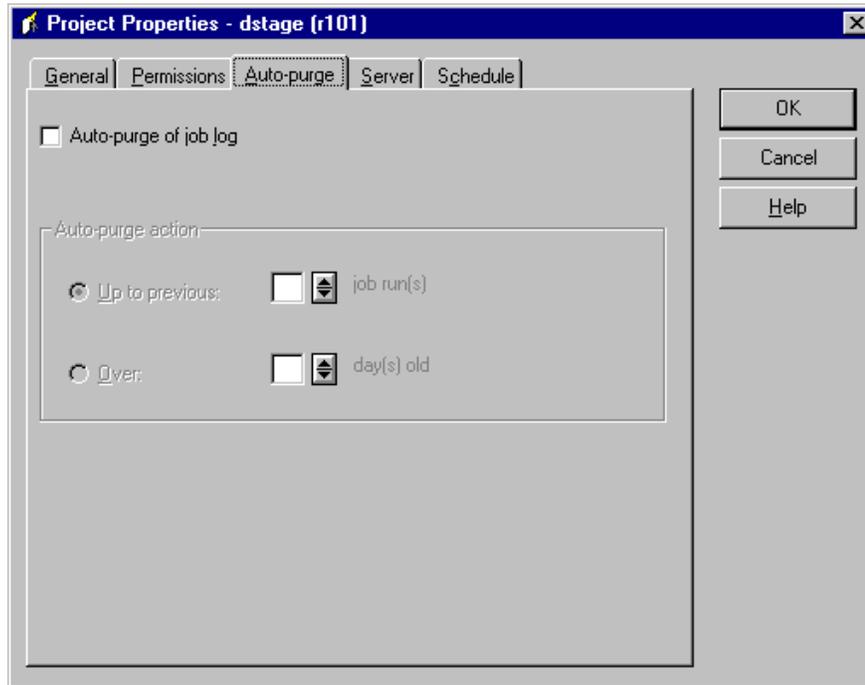
Every DataStage job has a log file, and every time you run a DataStage job, new entries are added to the log file. To prevent the files from becoming too large, they must be purged from time to time. You can set project-wide defaults for automatically purging job logs, or purge them manually. Default settings are applied to newly created jobs, not existing ones.

### Setting Auto-Purge Defaults

To set automatic purging:

1. From the DataStage Project Administration dialog box (see page 1-2), click **Properties**. The Project Properties window appears.

2. Click the **Auto-purge** tab to move the **Auto-purge** page to the front as shown in the following screen:



3. Select the **Auto-purge of job log** check box.
4. Select the Auto-purge action. You can purge jobs over the specified number of days old, or specify the number of jobs you wish to retain in the log. For example, if you specify 10 job runs, entries for the last 10 job runs are kept.
5. 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 dialog box. Choose **Job > Clear Log...** from the DataStage Director window to access this dialog box. For more information about purging log file entries, see *DataStage Operator's Guide*.

## Purging Log Files Manually

You can override automatic job log purging for an individual job by choosing **Job > Clear Log...** from the DataStage Director window. For more information, see *DataStage Operator's Guide*.

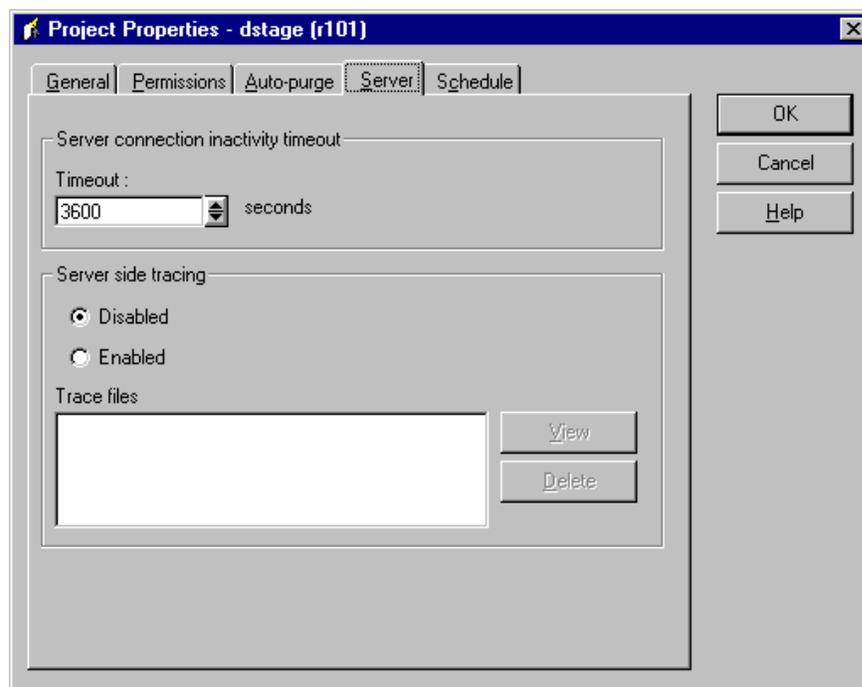
## Setting Server Properties

You can set the timeout on the connection between the client and server, and you can trace activity on the server as a help to diagnosing project problems.

### Setting the Server Timeout

The connection between the DataStage client and server times out after a specified period of inactivity. The default is 3600 seconds. To change the default:

1. From the DataStage Project Administration dialog box (see page 1-2), click **Properties**. The Project Properties window appears.
2. Click the **Server** tab. The **Server** page appears, as shown in the following screen:



3. Enter the time required, in seconds, in the **Timeout** field, then click **OK**.

## Enabling Tracing on the Server

The default is for server tracing to be disabled. When you enable it, information about server activity is recorded for any clients that subsequently 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 invoke a DataStage client.

**Note:** Server tracing is not compatible with the job administration option **Cleanup Resources**, which you also enable from the Administration client. See “DataStage Job Administration” on page 1-15 for more information.

To enable tracing on the server:

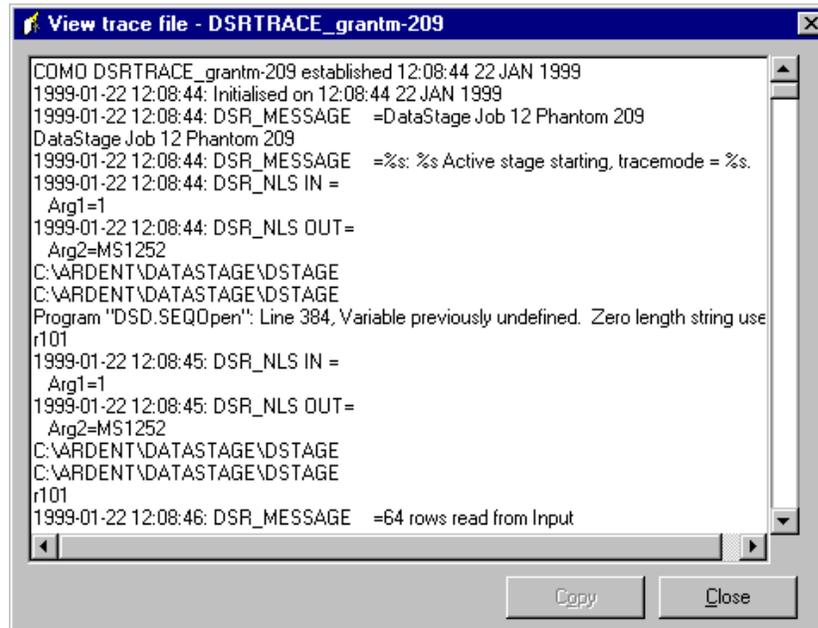
1. From the DataStage Project Administration dialog box (see page 1-2), click **Properties**. The Project Properties window appears.
2. Click the **Server** tab. The **Server** page appears, as shown in the previous section.
3. Click the **Enabled** option button in the **Server side tracing** area.
4. Click **OK** to activate server tracing.

Trace files are added to the **Trace files** list box on the **Server** page whenever a client attaches to the project.

To view a trace file:

1. 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:



2. To copy trace information to the Clipboard, select text in the View trace file window, then click **Copy**.
3. Click **Close** to close the View trace file window.

To delete a trace file, select one or more file names in the **Trace files** list box on the **Server** page, then click the **Delete** button.

## DataStage Job Administration

From the Administration client, the administrator can enable or disable job administration features in the DataStage Director. They let 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.

To enable job administration, select the **Enable job administration in Director** check box in the **General** page on the Project Properties window. This adds two options to the Director Job menu:

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

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

- View and end job processes
- View and release locks

**Note:** **Cleanup Resources** is not compatible with server tracing (see “Enabling Tracing on the Server” on page 1-13). If you enable server tracing and job administration, the DataStage 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.

**CAUTION:** These two options give the user considerable control over job resources, and should be used with care.

Refer to *DataStage Operator's Guide* for full information about how to use these options.

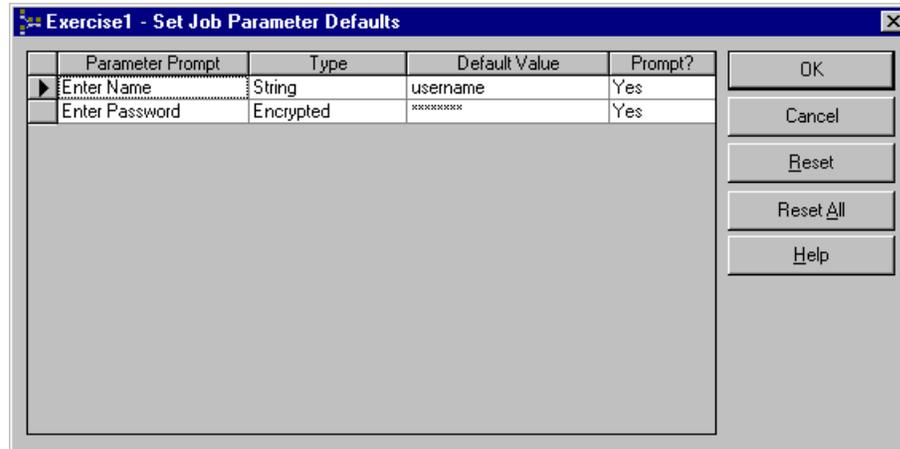
## Setting Job Parameter Defaults

DataStage jobs usually have parameters that are supplied to them for use at run time. Sometimes the parameters are variables such as file names, dates, or passwords that must be entered when the job is run. In other cases the parameters have default values that are specified by the job's designer. If the default values that were supplied with a DataStage job do not meet the needs of your system, you can change them from the Set Job Parameter Defaults dialog box.

You can reach the Set Job Parameter Defaults dialog box as follows:

1. Select the relevant job in the Job Status view of the DataStage Director window.

2. Do either of the following to display the Set Job Parameter Defaults dialog box:
  - Right-click to display the shortcut menu, then select **Set Defaults...**
  - Choose **Job > Set Defaults...** from the menu bar.



The columns in the dialog box are described in Table 1-1.

**Table 1-1. Columns in the Set Job Parameter Defaults Dialog Box**

| Column           | Description  | Editable   |
|------------------|--|------------|
| Parameter Prompt | The text used to prompt for this parameter when the job is run.  | No         |
| Type             | The type of parameter.   | No         |
| Default Value    | The default value for the parameter.   | Yes        |
| Prompt?          | The value is either Yes or No, which determines whether DataStage prompts for the parameter when the job is run. If you remove a default parameter value, the value changes to Yes and you will not be able to change it to No unless you enter a default value. (This does not apply to parameters with a type of String or Encrypted which can have an empty string as their default value, and No as the value in this column.) | Indirectly |

Click **Reset** to reset the value for the selected parameter to the job designer's default, or clear a value you have entered earlier.

Click **Reset All** to reset all parameter values to the job designer's defaults.

## Changing Encrypted Parameters

If a parameter is specified as encrypted, its default value appears in the job parameters list as a string of asterisks. When you edit the value of an encrypted parameter, you see the Encrypted dialog box:

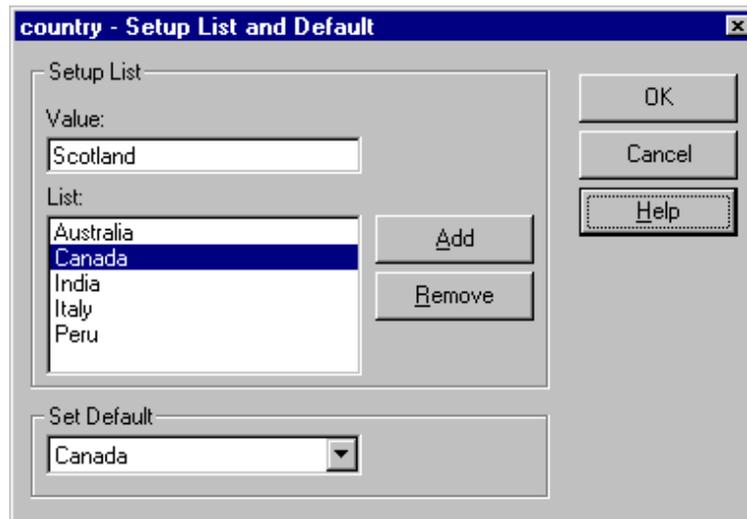


You specify the value in the same way as you would when you change a password.

1. Enter the value to be encrypted in the **Encrypted String** field. The value can be an empty string.
2. Reenter the value in the **Confirm Encrypted String** field to confirm it.
3. Click **OK**.

## Changing a Default List of Values

Some parameters have a list of values from which one is chosen when a job is run. You specify the contents of the list and the default value from the Setup List and Default dialog box. It appears whenever you edit an entry in the Default Value column in the Set Job Parameter Defaults dialog box that has a Type value of "List".



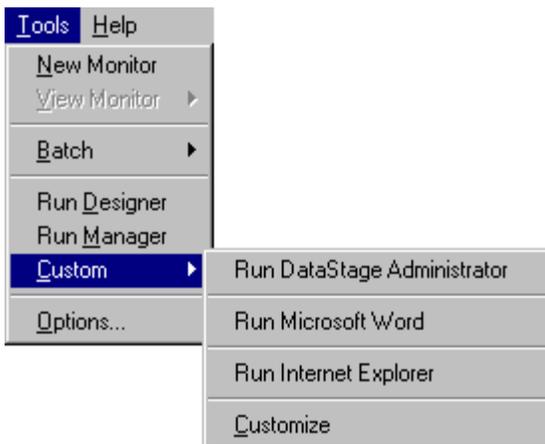
Enter a value to add to the list in the **Value** field, then click **Add**. The value appears in the **List** field.

To remove an entry, select it in the list and click **Remove**.

Choose one of the entries from the **Set Default** drop-down list to act as the default value. The current default value on this screen is Canada.

## Customizing the Tools Menu

The DataStage Designer, Manager, and Director windows each have a Tools menu that is used to call other DataStage components. You can customize the Tools menus to call other software directly from DataStage. For instance, you could add a word processor, an Internet browser, or the DataStage Administrator as shown in this example:



The Tools menu is customized for each user, and is saved as part of the user's DataStage environment. As well as adding and removing entries in the menu, you can specify the order in which they appear, and separate the entries with dividers. You can add up to 20 entries to the menu, including dividers.

### Adding a Tool

To add an entry to a Tools menu:

1. From the menu bar, choose **Tools** ► **Custom** ► **Customize**. The Customize dialog box appears.
2. Click **Add**. The Add Tool dialog box appears.
3. Enter the path of the program you want to add, or click ... to browse for the path. If the program you entered cannot be found, you are prompted to change the path. Click **No** to add the entry anyway if you intend to install the program later.
4. Click **OK**. The program's run file details appear in the Customize dialog box.

5. Edit the **Menu Text** field as required. (This field contains the text that appears in the Tools menu.)
6. Edit the **Arguments** field to specify any arguments required by the program. Click > for a list of variables you can use in this field to specify, for example, the current project, user name, or host.
7. Optionally add text in the **Status bar** field. This text appears in the status bar when you call the program from the Tools menu of the DataStage Designer.
8. Click **OK**. The entry is added to the Tools menu.

## Changing the Tools Menu Order

To change the order of entries on the Tools menu:

1. From the menu bar, choose **Tools > Custom > Customize**. The Customize dialog box appears.
2. Select the entry you want to move in the **Menu contents** field.
3. Click the **Up** or **Down** arrow to move the entry to the required position.
4. Click **OK** to save your new settings and close the Customize dialog box.

**Note:** You can only change the order of the **Custom** area of the Tools menu.

## Adding a Separator

To add separators between entries on the Tools menu:

1. From the menu bar, choose **Tools > Custom > Customize**. The Customize dialog box appears.
2. Click **Add...** . The Add Tool dialog box appears.
3. Select the **Add as menu separator** check box and click **OK**. The separator appears in the **Menu contents** field.
4. Click the **Up** or **Down** arrow to move the separator to the required position.
5. Click **OK** to save your new settings and close the Customize dialog box.

## Issuing UniVerse Commands

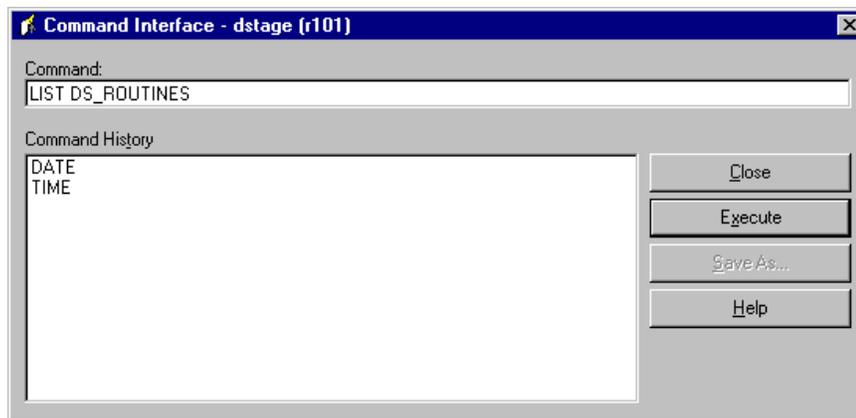
The Administration client lets you issue UniVerse commands directly from a selected project rather than having to use a Telnet session.

**Note:** You cannot issue the following UniVerse commands from the client: LOGOUT, LO, QUIT, Q, and OFF.

**CAUTION:** You should take care when issuing UniVerse commands that manipulate the NLS settings of your system. Incorrect use of these commands could disrupt both DataStage and the rest of your system.

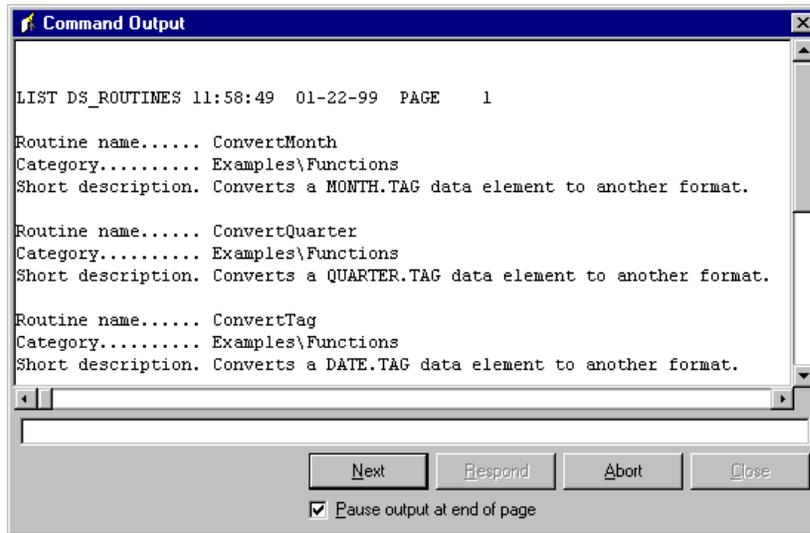
To issue a UniVerse command from the Administration client:

1. From the Ardent DataStage program folder, choose **DataStage Administrator**.
2. Use the Attach to DataStage dialog box to log on to the server. When you click **OK**, the DataStage Project Administration dialog box appears.
3. Select a project in the **Projects** list box, then click **Command**. The Command Interface dialog box appears:



4. Either type the UniVerse 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**. The Command Interface dialog box reappears.
8. If you want to save commands to the VOC file on the server, 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**.

**Note:** You cannot save a sentence or paragraph that contains a call to itself.

9. Click **Close** to close the Command Interface dialog box.



# 2

## Windows NT Server Administration

This chapter describes the administration tasks that need to be carried out when the DataStage server runs under Windows NT. These tasks include:

- Starting the Windows NT services used by DataStage
- Accessing ODBC databases and remote UniVerse systems
- Installing DataStage packages (including plug-ins)
- Installing MetaBrokers

### Starting the Schedule Service

The DataStage Director uses the Windows NT Schedule service to schedule jobs. If you try to schedule jobs when this service is not running, an error message appears. To correct this, start the Schedule service as follows:

1. Log on to the DataStage server. Choose **Start ► Settings ► Control Panel**. The Control Panel window appears.
2. Choose **Services**. The Services dialog box appears.
3. Scroll down and choose **Schedule**.
4. Click **Start**. The service is started.
5. Click **Startup...**, then click **Automatic**. This means the Schedule service restarts automatically when the system reboots.
6. Click **Close** to close the Services dialog box.

## Starting and Stopping the Server Engine

You may need to restart or stop the server engine (UniVerse), for example, if you are recommended to do so during an upgrade.

To start UniVerse:

1. Log on to the DataStage server. Choose **Start ► Settings ► Control Panel ► Services**. The Services dialog box appears.
2. Select **UniVerse Resource Service**, then click **Start**.
3. Click **Startup...**, then click **Automatic**. This ensures that UniVerse is started automatically whenever the server is rebooted.
4. Repeat steps 2 and 3, selecting **Uni RPC Service** and **UniVerse Telnet Service**.

To stop UniVerse, select **UniVerse Resource Service**, then click **Stop**. This also stops the Uni RPC service and the UniVerse Telnet service.

## Accessing ODBC Databases

If you want to access an ODBC database from DataStage, you must first install ODBC drivers for the ODBC data source on the DataStage server. You can install your own preferred drivers or the OEM drivers distributed with DataStage. If you choose to install the DataStage ODBC drivers, read *DataStage Installation Instructions*.

**Note:** Do not use trial versions of ODBC drivers on an active DataStage system. When the trial version expires, these drivers tend to display reminder screens that need user input. The DataStage server cannot handle these screens and the job or meta data import hangs.

When you have installed a suitable ODBC driver:

- Add the path of the directory containing the driver DLL to the Path system environment variable.
- Configure the ODBC data source as a System Data Source, not a User Data Source.

Otherwise DataStage cannot connect to the ODBC data source.

## Accessing Remote UniVerse Systems

To access files in remote UniVerse systems from DataStage, you must define the files in the *uvodbc.config* file in the relevant DataStage project directory on the server. The *uvodbc.config* file is a text file containing comments that describe how to add additional entries. The default *uvodbc.config* file looks like this:

```
*** Allows access to ODBC sources through BCI
***
*** To get to any ODBC source other than UniVerse, you need entries
*** that look as follows (the data source must also be configured
*** via the operating system's own mechanisms):
***
*** <data source name>
*** DBMSTYPE = ODBC
***
*** UniVerse on the local server is available via the data source
*** name "localuv" as defined below - please do not alter this
*** entry!
***
*** To access UniVerse as a data source, other than through localuv,
*** you need another entry exactly similar to that for localuv but
*** with a remote host name in place of "localhost".
***
*** Note that the spaces around the " = " signs are required, and
*** the data source name must be enclosed in angle brackets "<>".
[ODBC DATA SOURCES]
<localuv>
DBMSTYPE = UNIVERSE
network = TCP/IP
service = uvserver
host = localhost
```

You must not edit or delete the default DataStage entry in the file. To add an entry for a remote UniVerse system, position your cursor at the bottom of the file and add the following lines, with the italicized variables changed to suit your system:

```
<name>
DBMSTYPE = UNIVERSE
network = TCP/IP
service = uvserver
host = hostname
```

*name* identifies the remote UniVerse system in any way you find useful and must be enclosed in angle brackets, as shown.

*hostname* specifies the host where UniVerse is installed and must be recognized as a node name by your TCP/IP system.

Save the file when you have added all the host names you require.

## Installing DataStage Packages

Your DataStage supplier may provide packaged jobs and components to add to DataStage. DataStage plug-ins are provided as packages. To install a package into DataStage:

1. Either load the media containing the package, or copy the package files or directories into a temporary directory on your hard disk, as described in the instructions supplied with the package.
2. In the Ardent DataStage program folder, click the **DataStage Package Installer** icon to start the installation program.
3. Follow the instructions on the screen to install the package into DataStage.
4. Delete the files from the temporary directory, if necessary.

Consult the *Read Me* file in the plug-in directory (*Packages*) on the DataStage installation CD for details of plug-in packages provided.

The DataStage Package Installer also accepts command line options, allowing you to call it as part of another installation program.

### Package Installer Command Line Options

To run the Package Installer from the command line, use the following command:

```
%DSPackagePath%\setup.exe package=dirpath  
project=proj_name | allprojects  
[logfile=pathname] [appendlog]
```

**DSPackagePath** is a preset environment variable that specifies the location of the Package Installer setup program. The setup program takes the following options:

| Option                    | Description   |
|---------------------------|---|
| package= <i>dirpath</i>   | <b>package</b> is mandatory. <i>dirpath</i> is the directory in which the package is located. If the specified path does not exist or does not point to a directory, an entry is added to the log file (see the <b>logfile</b> option description) and setup exits. |
| project= <i>proj_name</i> | The package is installed into the project named <i>proj_name</i> . You can specify only a single project. The command line must contain either <b>project</b> or <b>allprojects</b> . If it includes both, <b>allprojects</b> overrides <b>project</b> .            |

| Option                   | Description  |
|--------------------------|--|
| allprojects              | Installs the package into all the projects on the current server. This option overrides <b>project</b> . You must include either <b>project</b> or <b>allprojects</b> on the command line.   |
| logfile= <i>pathname</i> | Identifies the file that logs information about the progress of the package installation. The pathname is optional. If you do not specify a directory or the specified directory does not exist, it defaults to <i>uvhome</i> . If you do not include <b>logfile</b> on the command line, the default is to use the file <i>install.log</i> in the directory <i>uvhome</i> . |
| appendlog                | Specifies that information should be appended to the log file, rather than overwrite the existing information.   |

For example, the command:

```
%DSPackagePath%\setup.exe package=c:\ikt\package1
allprojects appendlog
```

installs the package in directory *c:\ikt\package1* into all the projects on the current server. The default log file will be used (*uvhome/install.log*) and new log information appended to the file. (*uvhome* is the UV account directory, which is normally *C:\Ardent\DataStage\ServerEngine*.)

## Installing MetaBrokers

To use MetaBrokers in DataStage, you need to install the MetaBrokers for each data warehousing tool with which you want to exchange meta data.

MetaBrokers are provided on the DataStage installation CD under the *MetaBrkr* directory. The *Read Me* file in *MetaBrkr* lists the available MetaBrokers. To install a MetaBroker:

1. Navigate to the directory on the CD containing the MetaBroker you want to install.
2. Run the *setup.exe* program in this directory.
3. An installation wizard appears and guides you through the installation process. This includes supplying license information for the MetaBroker.

Documentation is supplied for each individual MetaBroker in the installation directory.

**CAUTION:** If you previously had MetaBrokers installed with DataStage 3.5, you must uninstall these before installing the DataStage 3.6 MetaBrokers (otherwise the install will fail).

# 3

## UNIX Server Administration

This chapter describes the administration tasks that need to be carried out when the DataStage server runs under UNIX. These tasks include:

- Checking installation requirements
- Tuning UNIX for DataStage
- Configuring the DataStage engine
- Configuring the RPC daemon environment
- Starting and stopping the server engine (UniVerse)
- Accessing remote UniVerse systems
- Accessing ODBC databases
- Installing DataStage packages (including plug-ins)
- Observing DataStage UNIX server restrictions
- Accessing remote files via NFS

You should read this chapter in conjunction with the *Read Me* file on your installation CD. Print the *Read Me* file and have it on hand as you read this chapter. This chapter describes what you need to do, and the *Read Me* file gives information specific to your UNIX platform.

### Installation Requirements

The DataStage UNIX server is installed from the DataStage CD. Before installing, you should consider the following points:

- If you are installing on a system that has never had UniVerse or DataStage installed, then you may need to alter some of your kernel parameters. See “Tuning UNIX for DataStage” on page 3-2.

- If you are installing on a system that already has UniVerse installed, but not GCI, then you will need to install GCI before installing DataStage. See “System with UniVerse Installed” on page 3-3.
- If you are installing on a system where you already use GCI, then you will have to perform some manual relinking after DataStage installation. See “System with UniVerse and GCI Installed” on page 3-3.
- If you are upgrading (or reinstalling) DataStage, you need to set your shared library search path to include the *lib* subdirectory of the DataStage server home directory *uvhome* (specified in the *./uvhome* file). See “Environment Requirements for unirpcd” on page 3-4.
- If you require access to ODBC data sources, you will need to install appropriate ODBC driver software. See “Accessing ODBC Databases” on page 3-6.
- If you intend to install any plug-in additions to DataStage, you may need to configure the RPC daemon environment. See “Environment Requirements for unirpcd” on page 3-4.

To install the DataStage server, run the *install.sh* program from the CD.

When you upgrade to DataStage 3.6 from an earlier release, all projects are automatically upgraded.

## Tuning UNIX for DataStage

The DataStage Server runs jobs and other activities as UNIX processes. The demands the DataStage Server makes on the operating system depend on:

- The particular DataStage activity being performed
- The number of DataStage processes active on the system

To get the best out of your system, you will need to tune UNIX kernel parameters. The parameters you may need to change are:

- MSGMAX
- MSGMNB
- SHMSEG
- SHMMAX
- SEMMNS
- SEMMSL
- SEMMNI

The required values for these parameters vary according to platform. Consult the *Read Me* file for the minimum settings for your system. These values may need to be further adjusted to meet your requirements. Use the information provided by error messages, or performance measuring tools such as *sar*, to find the optimum values for your system.

If necessary, rebuild the kernel with changes to these parameters. For more information on rebuilding the kernel, see the UNIX manuals supplied with your system.

## DataStage Engine Configuration

### System with UniVerse Installed

Before installing DataStage on a machine that already has UniVerse, you must install the UniVerse layered product, GCI. GCI (General Calling Interface) allows C code to be called from UniVerse BASIC.

GCI can be installed from your UniVerse CD. Use the UniVerse system administrator account and install it in *uvhome* (specified in the *./uvhome* file).

### System with UniVerse and GCI Installed

If you already have UniVerse and GCI installed and in use, then you will need to relink applications that use GCI after installation of DataStage.

Installing the DataStage server replaces *uvsh* and *uvapi\_slave* with new versions that have been linked with DataStage GCI routines. Old copies of these programs are saved as *uvsh.ds.saved* and *uvapi\_slave.ds.saved*.

Before reusing your existing applications, you will have to manually relink *uvsh* and *uvapi\_slave* to include both your routines and the DataStage GCI routines. There are two DataStage GCI files supplied with the server: *dsdgc.i.o* and *dsradmin.o*.

**Note:** These are located in *gcidir* in the UV account directory. You should have a thorough understanding of GCI administration before attempting to relink the programs.

## The RPC Daemon (*unirpcd*)

The RPC daemon for the server engine is called *unirpcd*. It must be running to allow connection of DataStage clients.

By default it is started when the server is installed, and should start whenever you restart your machine. If you cannot connect from a DataStage client, check that *unirpcd* is running.

If you need to start *unirpcd* manually, use the following command:

```
# uvhome/bin/rpc_start.sh
```

*uvhome* is specified in the file */.uvhome*.

If your site is already using UniVerse, you can use the UniVerse Package Administration menu to start and stop *unirpcd*. If you do so, bear in mind that *unirpcd* inherits the user's environment, so you should ensure that the environment is sufficient to run DataStage.

## Environment Requirements for *unirpcd*

ODBC drivers and some plug-ins require that certain directories are included in the shared library environment variable setting for *unirpcd*. Consult the *Read Me* file for the environment variable for each platform type.

Some databases also require specific environment variables (for example, Oracle requires ORACLE\_HOME and Sybase requires SYBASE). To access these databases from DataStage, *unirpcd* must have these environment variables set to appropriate values.

DataStage provides two scripts that start the *unirpcd* process: *rpc\_start.sh* to manually restart and *uv.rc* which is run when the machine is rebooted. The location and exact name of *uv.rc* is platform-dependent. Consult the *Read Me* file for details. If you are using ODBC drivers or plug-ins that require amendment to the shared library environment variable, then you must edit these two scripts to add the appropriate directories. Find the string USER MODIFIABLE PATH and modify the library path that follows it. In addition, you must add any database-specific environment variable settings to the script.

After editing these scripts, you should restart *unirpcd* by killing the process and running the *rpc\_start.sh* script.

## Starting and Stopping the Server Engine

From time to time you may need to stop or restart the DataStage server engine (UniVerse). You should always stop the server engine before adding additional DataStage projects or upgrading to a new release, and whenever you wish to shut down the server.

A script is provided for these purposes, which is generically known as *uv.rc*. The exact name and location are platform-dependent. See the *Read Me* file for details.

To stop the server engine, use:

```
# uv.rc stop
```

This shuts down the server engine and frees any resources held by server engine processes.

To restart the server engine, use:

```
# uv.rc start
```

This ensures that all the server engine processes are started correctly.

## Accessing Remote UniVerse Systems

To access files in remote UniVerse systems from DataStage, you must define the files in the *uvodbc.config* file in the relevant DataStage project directory on the server. The *uvodbc.config* file is a text file containing comments that describe how to add additional entries. The default *uvodbc.config* file looks like this:

```
*** Allows access to ODBC sources through BCI
***
*** To get to any ODBC source other than UniVerse, you need entries
*** that look as follows (the data source must also be configured
*** via the operating system's own mechanisms):
***
*** <data source name>
*** DBMSTYPE = ODBC
***
*** UniVerse on the local server is available via the data source
*** name "localuv" as defined below - please do not alter this
*** entry!
***
*** To access UniVerse as a data source, other than through localuv,
*** you need another entry exactly similar to that for localuv but
*** with a remote host name in place of "localhost".
***
*** Note that the spaces around the " = " signs are required, and
*** the data source name must be enclosed in angle brackets "<>".
```

```
[ODBC DATA SOURCES]
<localuv>
DBMSTYPE = UNIVERSE
network = TCP/IP
service = uvserver
host = localhost
```

You must not edit or delete the default DataStage entry in the file. To add an entry for a remote UniVerse system, position your cursor at the bottom of the file and add the following lines, with the italicized variables changed to suit your system:

```
<name>
DBMSTYPE = UNIVERSE
network = TCP/IP
service = uvserver
host = hostname
```

*name* identifies the remote UniVerse system in any way you find useful and must be enclosed in angle brackets, as shown.

*hostname* specifies the host where UniVerse is installed and must be recognized as a node name by your TCP/IP system.

Save the file when you have added all the host names you require.

## Accessing ODBC Databases

If you want to access an ODBC database from DataStage, you must first install the appropriate ODBC drivers. To do so, take the following steps:

1. Install and configure an ODBC driver manager and one or more ODBC drivers.
2. Configure ODBC data sources according to the instructions supplied by the driver vendor. Test access to these sources using the tools supplied with the driver. Do not try to access these sources from DataStage until you have successfully tested access.
3. Set up the environment variables required by the RPC daemon (*unirpcd*) and any further environment variables required by the particular ODBC driver or DBMS. See “Environment Requirements for unirpcd” on page 3-4.
4. For each ODBC data source that you want to access from DataStage, add an entry to the *uvodbc.config* file. See “Adding Data Sources to uvodbc.config” on page 3-7.
5. Run the *relink.uvlibs* shell script to access the ODBC driver manager that you installed in step 1. See “Using relink.uvlibs to Maintain Links” on page 3-8.

An OEM version of the Intersolv ODBC pack is supplied with DataStage.

## Installing Intersolv Drivers

To install the Intersolv ODBC drivers supplied with DataStage:

1. Copy the file *ivodbc.tar* from the top-level directory on the DataStage CD to a location on your system with at least 100 MB of free space.
2. Enter the following command in order to extract the archived file:  

```
# tar - xf ivodbc.tar
```
3. Read the *install.txt* file in the extracted directory for further instructions on installing the drivers.
4. Read the *odbcref.pdf* file in the extracted directory for instructions on configuring the drivers.

## Adding Data Sources to *uvodbc.config*

You must add an entry to the *uvodbc.config* file for every ODBC data source that you want to access from DataStage.

The *uvodbc.config* file is normally in the UV account directory (*uvhome*), but you can also have different *uvodbc.config* files in each project directory to configure different data sources. This is useful where you configure a data source that is known to some projects but not others. By default, DataStage searches the current project directory for a *uvodbc.config* file and, if it finds one, uses this in preference to a file in *uvhome*.

Add an entry to *uvodbc.config* in the form:

```
<name>  
DBMSTYPE = ODBC
```

*name* is the ODBC data source name, enclosed in angle brackets as shown. Also note that the spaces either side of the equal sign are required.

See “Accessing Remote UniVerse Systems” on page 3-5 for more information about editing the *uvodbc.config* file. For more information about configuring ODBC connections, see “ODBC Connection Problems on UNIX” on page 5-4.

## Configuring the ODBC Driver Manager

The procedure for configuring the ODBC driver manager for UNIX may vary according to the type of ODBC driver you are using. If in any doubt, consult the

vendor of your ODBC driver software. Considerations for different UNIX platforms and different drivers are in the *Read Me* file.

Provided your ODBC driver manager is a shared object (*.so* or *.sl* file), then the configuration procedure is largely as follows:

1. Install the ODBC driver manager according to the vendor's instructions.
2. Configure environment variables for *unirpcd*.
3. Determine where the ODBC shared library *libodbc.xx* resides. For example, for the Intersolv driver, this is *\$ODBCHOME/libs*.
4. Shut down the DataStage server using the *uv.rc* script for your platform. See "Starting and Stopping the Server Engine" on page 3-5.
5. Include the ODBC shared library directory in the library path environment variable setting for the RPC daemon. See "Environment Requirements for unirpcd" on page 3-4.
6. Execute the *relink.uvlibs* script to link to the ODBC shared library (omit this step if you are using Intersolv on an AIX platform). See "Using *relink.uvlibs* to Maintain Links."
7. Restart the DataStage server using the *uv.rc* script for your platform. See "Starting and Stopping the Server Engine" on page 3-5.

## Using *relink.uvlibs* to Maintain Links

The *relink.uvlibs* script is executed to maintain the correct links for the ODBC shared library.

DataStage initially installs a stub version of the ODBC shared manager library in the *uvdlls* directory of the UV account directory (*uvhome/uvdlls*). The library has the name *libodbc.xx*, where *xx* is supplied by the system you are running on. The DataStage installation program creates a symbolic link from */.uvlibs* to the *uvdlls* directory and both *uvsh* and *uvsrvd* look for their shared libraries in */.uvlibs*, so it is vital that the symbolic link is not broken.

To use an ODBC driver you need to reset the symbolic link to use the driver manager supplied by your ODBC driver vendor (you may also need to reset the link if the symbolic link is broken or if the shared libraries are moved to another directory).

To reset the link, use the *relink.uvlibs* shell script, which lives in *uvhome/bin*. It has the following syntax:

```
relink.uvlibs pathname
```

*pathname* is the full pathname of the directory containing the shared libraries, for example:

```
# relink.uvlibs $ODBCHOME/libs
```

**CAUTION:** If you are using an Intersolv driver on an AIX platform, you should not run *relink.uvlibs* (because the driver manager is not a shared object). To use an ODBC driver other than Intersolv on an AIX platform, contact your Ardent representative for assistance in configuring DataStage to work with the driver.

## Trial Version of ODBC Drivers

You should not use trial versions of ODBC drivers. The DataStage server cannot handle the screens highlighting the evaluation status of the driver and may hang.

## Installing DataStage Packages

Your DataStage supplier may provide packaged jobs, components, or plug-in stage types to add to DataStage. To install the packages on the DataStage server:

1. Either load the media containing the package, or copy the package files or directories into a temporary directory on your hard disk.
2. Log on to the DataStage server as superuser.
3. At the UNIX prompt, enter the command:

```
# uvhome/bin/dspackinst
```

*uvhome* is the installation directory of the DataStage server engine (UniVerse) and is specified in the file */.uvhome*.

4. Follow the instructions on the screen to install the package into DataStage.
5. Delete the files from the temporary directory, if necessary.

Consult the *Read Me* file in the plug-in directory (*Packages*) on the DataStage UNIX installation CD for details of plug-in packages provided.

Also see the *Read Me* file for platform-specific considerations about the BCPLoad plug-in.

## BCPLoad Plug-In Restrictions

Various restrictions apply when using BCPLoad stages with a UNIX server.

The BCPLoad plug-in does not support MSDBLIB as a client library on UNIX.

If you use the BCPLoad plug-in stage to access Sybase, the *unirpcd* process must have the SYBASE environment variable set and possibly the DSQUERY variable. (SYBASE points to the directory where the Sybase clients are installed.) In addition, the *unirpcd* process must have the Sybase library directory (*SYBASE/lib*) included in the environment variable used to locate shared libraries.

Consult your Sybase documentation for information about these variables. See also the section “Environment Requirements for unirpcd” on page 3-4.

## Server Restrictions

### Connecting to the Server

The **Omit** option on the DataStage client Attach to Project dialog boxes is applicable only for systems with Windows NT servers. It should not be used when connecting to UNIX servers, and you should make all your DataStage users aware of this.

By default, only *root* can view the pull-down menu of installed projects from the client Attach to Project dialog boxes. Users other than *root* can access these projects, but have to explicitly enter the project name into the **Project** field.

You can enable ordinary users to view the **Project** drop-down list by creating an empty file called *.developer.adm* in the DataStage server home directory *uvhome* (specified in the file */.uvhome*). Note that this only gives users permission to view the projects, not permission to use them.

If you need further restrictions, you can add a list of UNIX user groups to the *.developer.adm* file (one per line). Only users in these groups will be allowed to access the drop-down list.

### Scheduling Jobs

When using the DataStage Director to schedule a job to run on a UNIX server, the **Every** option is available only with a single day of the week or month (it is not possible to use it with multiple days within a week or month).

The DataStage server uses a combination of the UNIX *at* and *cron* system commands to provide flexible job scheduling. A scheduled job runs under the user name of the connected Director client session that scheduled the job. You must ensure that any DataStage operators who need to schedule jobs have adequate permissions to invoke the *at* and *cron* commands. See your UNIX documentation for details of how to do this.

The *cron* command does not set environment variables for the jobs that it runs. Therefore some DataStage jobs, particularly those whose components use third-party software, may not run. If this is the case, you may need to manually edit the *crontab* entry to set the required environment variables.

## Accessing Remote Files via NFS

If you do not have UV/Net (the UniVerse network support facility) it is possible to configure your DataStage server to allow access to remote files across the Network Filing System (NFS).

To do so, you must reconfigure the ALLOWNFS configurable parameter in the *uvconfig* file in the UV account directory *uvhome*:

1. Edit *uvconfig* to change ALLOWNFS from 0 to 1.
2. Execute *uvregen* in the *bin* directory of the UV account directory (*uvhome/uv/bin/uvregen*).
3. Reboot the system.

When using NFS, you will have to use fully qualified UNIX pathnames in the format *machine/path* rather than the Windows NT UNC-type pathnames.

**CAUTION:** Using NFS for remote file access in this way may not use the locking mechanisms of the database or file system on the remote machine. Consequently, you should exercise extreme caution when using this method to avoid data integrity violations.



# 4

## NLS Configuration

DataStage has built-in National Language Support (NLS). This means 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, the DataStage server engine holds data in UNICODE format. This is an international standard character set that contains many of the characters used in languages around the world. DataStage maps data to or from UNICODE format as required.

Each 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, that the project should use. The DataStage client and server components also have maps assigned to them during installation to ensure that data is transferred in the correct format.

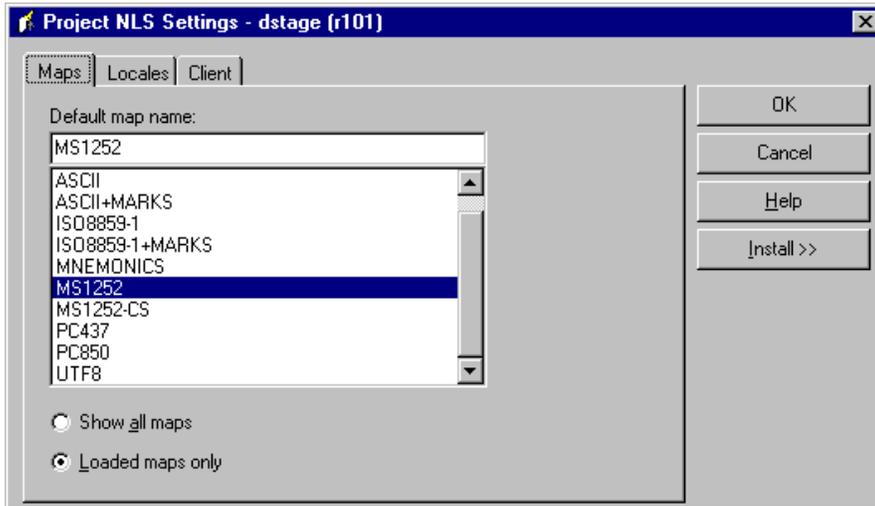
From the DataStage Project Administration dialog box, you can check which maps and locales were assigned during installation and change them as required.

### Changing Project Maps

To view or change a project map, select the project in the DataStage Project Administration dialog box, then click **NLS**.

**Note:** If the **NLS** option is not active, you do not have NLS installed. You must install NLS from the DataStage Master Setup screen as described in *DataStage Installation Instructions* before you can access the windows described in the rest of this chapter.

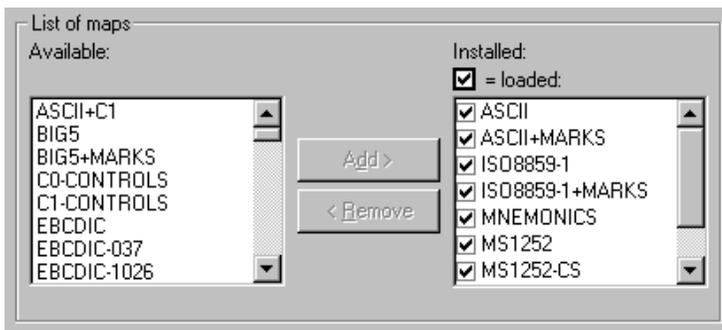
The Project NLS Settings window appears showing the **Maps** page:



The **Default map name** field shows the current map that is used for the project. By default, the list shows only the maps that are loaded and ready to use in DataStage. You can examine the complete list of maps that are supplied with 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 DataStage, click **Install>>** to see additional options on the **Maps** page:

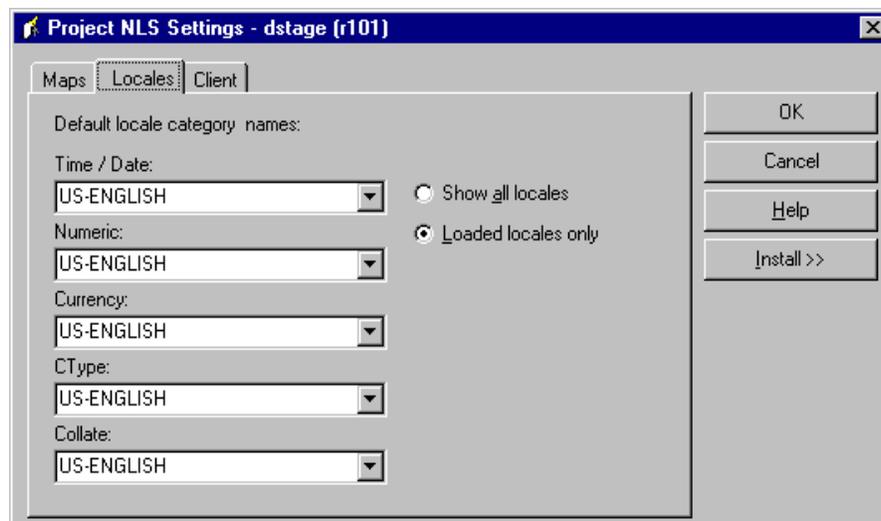


The **Available** list shows all the character set maps that are supplied with DataStage. The **Installed/loaded** list shows the maps that are currently installed. To install a map, select it from the **Available** list and click **Add>**. The map is loaded into DataStage ready for use the next time the server is restarted. If you want to use the map immediately, you can restart the server engine. For more information, see “Starting and Stopping the Server Engine” on page 2-2 for Windows NT servers or “Starting and Stopping the Server Engine” on page 3-5 for UNIX servers.

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.

## Changing Project Locales

To view or change default project locales, click the **Locales** tab. The **Locales** page appears at the front of the Project NLS Settings window:



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

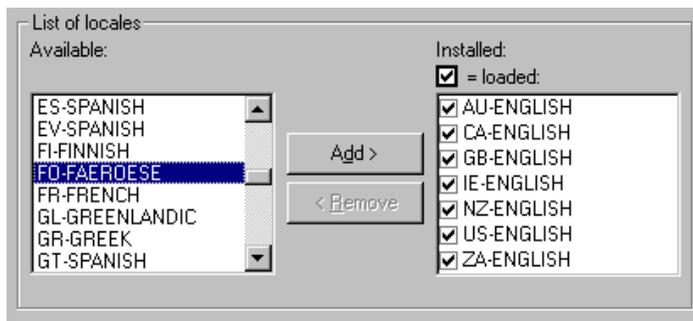
- **Time/Date** – The format for dates and times, for example, 31 Dec 1999 or 12/31/99 are two ways of expressing the same date that may be used in different locales.
- **Numeric** – The format used for numbers, including the thousands separator and radix (decimal) delimiter.

- **Currency** – The format for monetary strings, including the type and position of the currency sign (\$, £, F, DM, 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 DataStage by clicking **Show all locales**. These locales must be installed and loaded into DataStage before you can use them.

## Installing and Loading Locales

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



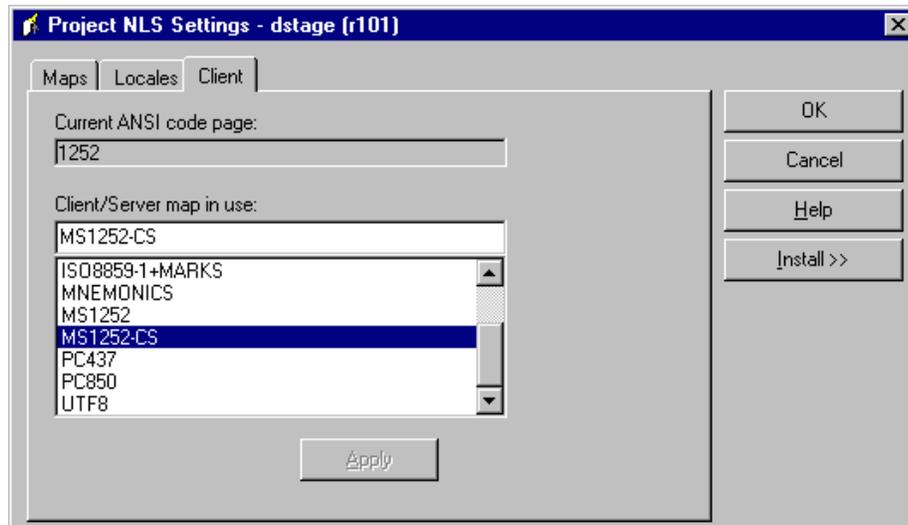
The **Available** list shows all the locales that are supplied with DataStage. The **Installed/loaded** list shows the locales that are currently installed. To install a locale, select it from the **Available** list and click **Add>**. The locale is loaded into DataStage ready for use the next time the server is restarted. If you want to use the locale immediately, you can restart the server engine. For more information, see “Starting and Stopping the Server Engine” on page 2-2 for Windows NT servers or “Starting and Stopping the Server Engine” on page 3-5 for UNIX servers.

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

## Client/Server Maps

When you installed the DataStage server, you specified the language that you want DataStage to support. DataStage automatically sets the language supported on the DataStage clients to match what you specified for the server. But if you access the DataStage server from a different client, data may not be mapped correctly between the client and the server.

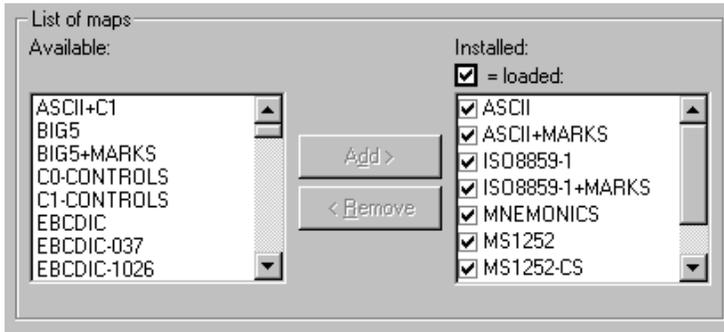
To prevent this from happening, you must change the client maps. To view the current mapping, click the **Client** tab on the Project NLS Settings window. The **Client** page appears:



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 server. The **Client/Server map in use** field shows the name of the map being used on the server for the current client session. The list shows all loaded maps.

If you select a map and click **Apply**, 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 may be rejected if it is not appropriate.

To install further maps into DataStage, click **Install>>** to display further options on the **Client** page:



DataStage uses special maps for client/server communication, with names ending in “-CS” (for Client Server). You should always choose one of these maps for this purpose.

The **Available** list shows all the character set maps that are supplied with 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 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. For more information, see “Starting and Stopping the Server Engine” on page 2-2 for Windows NT servers or “Starting and Stopping the Server Engine” on page 3-5 for UNIX servers.

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.

# 5

## Troubleshooting

This chapter gives some ideas for troubleshooting problems in DataStage, including:

- Error messages
- Problems with scheduled jobs
- ODBC connection problems on UNIX
- DataStage client to UNIX server connections

### Error Messages

The first indication of a problem may be an error message. If you see an error message in DataStage, it is important to remember that there are potentially several sources for the error:

- The DataStage client
- The DataStage server
- The network connection between the client and the server
- The DataStage server engine (UniVerse)

Some message boxes include a **Copy** button that allows you to copy the error message to the Clipboard if you want to save the information that it contains. This is always advisable if you need to contact Ardent technical support about the problem.

### Scheduled Jobs

A DataStage user can schedule jobs to run at convenient times when the system is less busy. (This task is carried out from the Job Schedule view in the DataStage Director window.) DataStage does not have its own separate scheduling program. Instead, whenever a DataStage user schedules a job, the control of that job is

handed over to the underlying operating system on the server. This means that if scheduled jobs do not run correctly, the problem usually lies with the operating system configuration on the DataStage server.

The following sections offer some advice on troubleshooting the scheduler on both Windows NT and UNIX servers.

## **Windows NT Scheduling Problems**

On Windows NT servers, job scheduling is carried out by the Schedule service. If your scheduled job did not run, there are a number of things you can do to identify the cause.

### **Check That the Schedule Service Is Running**

First check whether the Schedule service is running, and if it is not, start it. See “Starting the Schedule Service” on page 2-1.

### **Look at the Schedule Log**

If the scheduled job still does not run, examine the schedule log for more clues. The schedule log is a text file called *dsr\_sched.log* in the project directory on the server. This file records any problems that occurred before control was handed over from the scheduler to DataStage. (After that point, messages are written to the appropriate job log file.) This file would contain a message if, for example, the server password you specified has expired.

### **Test the User Name and Password**

If you specified a user name and password to run all scheduled jobs in a project, make sure that you test the user name and password. This process is described in “Setting a User Name for Scheduled Jobs” on page 1-6.

- If the test works correctly but scheduled jobs still do not run, check that the user name you specified has permission to read and write to the project directory.
- If the test fails, there may be a problem with the user rights for the user name you specified. In which case, try the procedure in the next section.

## Check the User Rights

If the Windows NT Schedule service on the DataStage server is not running under the default user name, try this procedure to ensure that the Schedule service has the correct user rights:

1. From the Windows NT Start menu on the server, choose **Programs ► Administrative Tools ► User Manager**.
2. Choose **Policies ► User Rights...** . The User Rights Policy dialog box appears.
3. Select the **Show Advanced User Rights** check box.
4. From the **Right** list, select **Act as part of the operating system**.
5. If the user name for the renamed Scheduled service is not in the **Grant To** list:
  - a. Click **Add...** . The Add Users and Groups dialog box appears.
  - b. Click **Show Users** to add user names to the **Names** list.
  - c. Select the user name required and click **Add**.
  - d. Click **OK** to close the Add Users and Groups dialog box.
6. From the **Right** list, select **Replace a process level token**. Repeat the procedure in step 5, then continue at step 7.
7. From the **Right** list, select **Increase quotas**. Repeat the procedure in step 5, then continue at step 8.
8. Click **OK** to save the changes.

## UNIX Scheduling Problems

On UNIX servers, the scheduling of DataStage jobs is handled by the *at* and *cron* commands. These offer slightly different options than the Windows NT Schedule service, and the problems associated with them also vary.

### Viewing Scheduled Jobs

On UNIX servers, jobs are scheduled for each user. That means a user can only view jobs that are scheduled for the current user name. For a DataStage administrator, this means it is not possible to get a quick overall view of all the DataStage jobs that are scheduled to run over a particular period. The only way to find out which jobs are scheduled is to examine the files in the *cron* directory for each user ID. The naming and location of these files varies from system to system. For more information, see the reference page for the *cron* command.

### **If a Job Does Not Run**

If a scheduled job does not run, check that the user who scheduled the job has permission to use the *cron* command. To do this, examine the *cron.allow* and *cron.deny* files which contain lists of users who can and cannot use the command. The location of these files varies from system to system. For more information, see the reference page for the *cron* command.

## **Job Termination Problems**

If you experience delays in the termination of a DataStage job when it is run, empty the &PH& directory. There is a &PH& directory in each DataStage project directory, which contains information about active stages that is used for diagnostic purposes. The &PH& directory is added to every time a job is run, and needs periodic cleaning out.

## **ODBC Connection Problems on UNIX**

DataStage relies on third-party ODBC drivers to connect to ODBC data sources. This means that it is not always easy to track down problems in an ODBC connection. The following sections contain some ideas for troubleshooting ODBC connection problems.

### **Check the Shared Library Environment**

If you see a message similar to this one:

```
ld.so.1: uvsh: fatal: libxxxx: can't open file: errno=2
```

check that the ODBC driver's shared library has been added to the environment variable used to locate shared libraries (see "Accessing ODBC Databases" on page 3-6).

### **Check Symbolic Links**

If you have moved shared libraries to a new directory or have installed a new ODBC driver manager, you may have broken symbolic links used by the server engine.

To reset the symbolic links to a new directory, run this command at the UNIX prompt:

```
# uvhome/bin/dspackinst relink.uvlibs pathname
```

*uvhome* is the installation directory of the DataStage server engine (UniVerse).

*pathname* is the full pathname of the directory containing the shared libraries.

To reset links for a new ODBC driver manager:

1. Install the ODBC driver manager according to the vendor's instructions.
2. Determine where the ODBC shared library *libodbc.xx* resides. For example, the library for the Intersolv driver lives in *\$ODBCHOME/dlls*, and the library for the Visigenics driver lives in *\$ODBCHOME/libs*.
3. Exit any DataStage client applications.
4. Run the *relink.uvlibs* command as described above.
5. Restart any DataStage client applications.

## DataStage Client to UNIX Server Connections

If you cannot connect from a DataStage client to a UNIX server, check that the RPC daemon is running. The RPC daemon is started when the DataStage server is installed, and should start automatically when you reboot. If the daemon has stopped for some reason, restart it with the following command:

```
uvhome/bin/rpc_start.sh
```

*uvhome* is the DataStage server engine installation directory.

## Connecting to UniData Data Sources

If you encounter an error when connecting to a UniData data source, such as:

```
UniData Client error: call to UniOpenPos returned 45 - Client  
version (11) and server version (12) are incompatible
```

you need to edit the *UNIAPL.INI* file in the Windows directory on your client machine to change the value of the *PROTOCOL* variable to match the server version. So, in the case of the example error, you would need to change the variable value from 11 to 12:

```
PROTOCOL = 12
```



# Index

## A

- accessing
  - ODBC databases from UNIX servers 3-6
  - remote files via NFS from UNIX servers 3-11
  - remote UniVerse systems from UNIX servers 3-5
- Add Tool dialog box 1-20
- adding
  - data sources to *uvodbc.config* 3-7
  - entries to Tools menu 1-20
  - projects 1-10
  - separators in Tools menu 1-21
- Administrators group 1-2
- Ardent DataStage program folder, Administrator 1-2
- at* command 5-3
- authorization details 1-4
- auto-purge defaults 1-11

## B

- BCPLoad plug-in restrictions on UNIX servers 3-10

## C

- Change License Details dialog box 1-4
- changing
  - license details 1-4
  - locales 4-3
  - maps 4-1
  - Tools menu order 1-21
  - user group assignments 1-5
- character types 4-4

- cleaning up project files 1-10
- Cleanup Resources option 1-16
- Clear Status File option 1-16
- client to server connection 5-5
- client/server maps 4-5
- code page 4-5
- Command Interface dialog box 1-22
- Command Output window 1-23
- commands
  - at* 5-3
  - cron* 5-3
  - DSPackagePath 2-4
  - relink.uvlibs* 5-4
  - sar* 3-3
  - UniVerse 1-22
- Copy button 5-1
- cron* command 5-3
- cron.allow* file 5-4
- cron.deny* file 5-4
- currency signs 4-4
- Customize dialog box 1-20
- customizing the Tools menu 1-20

## D

- data sources
  - adding to *uvodbc.config* on UNIX servers 3-7
  - connecting to UniData 5-5
  - defining ODBC 2-2, 3-6
- DataStage Director window 5-1

- DataStage engine configuration on UNIX servers
  - starting and stopping 3-5
  - with UniVerse and GCI installed 3-3
  - with UniVerse installed 3-3
- DataStage installation program 1-9
- DataStage licenses 1-4
- DataStage Project Administration dialog box 1-2, 1-5, 1-8, 1-11, 4-1
- DataStage Repository Import dialog box 1-10
- DataStage users, setting up 1-5
- dates 4-3
- defining ODBC data sources 2-2, 3-6
- deleting projects 1-8
- documentation conventions vii
- DSPackagePath command 2-4
- dsr\_sched.log* file 5-2

**E**

- editing *uvodbc.config* file 2-3, 3-5
- environment requirements for *unirpcd* on UNIX servers 3-4
- errors in DataStage 5-1
- Export dialog box 1-9

**F**

- files
  - cleaning up 1-15
  - cron.allow* 5-4
  - cron.deny* 5-4
  - dsr\_sched.log* 5-2
  - lib* 5-4
  - log 1-11, 1-12
  - remote 3-11
  - trace 1-14
  - uvodbc.config* 2-3, 3-5, 3-7

**G**

- GCI 3-3

**I**

- installation requirements 3-1
- installing
  - DataStage packages 2-4, 3-9
  - Intersolv drivers on UNIX servers 3-7
  - locales 4-4
  - maps 4-6
  - MetaBrokers 2-5
  - plug-ins 2-4
- issuing UniVerse commands 1-22

**J**

- job administration 1-15
- job parameters
  - default lists 1-19
  - encrypted 1-18
  - setting defaults 1-16
- Job Schedule view 5-1
- jobs
  - cleaning up files 1-15
  - clearing status file 1-16
  - scheduling 1-6, 2-1, 3-10, 5-1
  - termination problems 5-4

**L**

- lib* files 5-4
- licenses 1-4
- locales 4-3
  - changing 4-3
  - installing 4-4
- log files
  - automatic purging 1-11
  - manual purging 1-12

## M

- maps 4-1
  - changing 4-1
  - client/server 4-5
  - installing 4-6
- meta data 2-5
- MetaBrokers, installing 2-5
- moving projects 1-8

## N

- National Language Support, *see* NLS
- NFS, accessing remote files via (UNIX servers) 3-11
- NLS (National Language Support)
  - client/server maps 4-5
  - locales 4-3
  - maps 4-1
  - overview 4-1
- NLS option 4-1
- NT system authority 1-6
- numbers 4-3

## O

- ODBC configuration file 2-3, 3-5
- ODBC connection problems 5-4
- ODBC data sources, defining 2-2, 3-6
- ODBC drivers 3-6, 5-4

## P

- Package Installer 2-4, 3-9
  - command line options 2-4
- Permissions page 1-5
- plug-ins, installing 2-4
- primary domain server user groups 1-6
- Project NLS Settings window 4-2
- Project Properties window 1-3

- projects
  - adding 1-10
  - cleaning up files 1-10
  - deleting 1-8
  - moving 1-8
- purging, automatic 1-11

## R

- radix 4-3
- rebuilding secondary indexes 1-10
- relicensing DataStage 1-4
- relink.uvlibs* command 5-4
- remote files 3-11
- remote UniVerse systems 2-3
- requirements
  - installation on UNIX server 3-1
  - UNIX environment for *unirpcd* 3-4
- resetting parameter defaults 1-18
- restrictions on UNIX servers
  - Omit option 3-10
  - scheduling jobs 3-10
- root* user 1-2
- RPC daemon 5-5
- rpc\_start.sh* script 5-5

## S

- sar* performance testing command 3-3
- schedule log 5-2
- Schedule page of Project Properties window 1-7
- scheduled jobs, troubleshooting 5-1
- scheduling jobs 1-6, 2-1, 3-10, 5-1
- secondary indexes, rebuilding 1-10
- server engine
  - starting and stopping from UNIX 3-5
  - starting and stopping from Windows NT 2-2
- server timeout, setting 1-13
- server tracing 1-14

- Set Job Parameter Defaults dialog box 1-16
- setting
  - auto-purge defaults 1-11
  - job parameter defaults 1-16
  - server timeout 1-13
- setting up DataStage users 1-4
- shared libraries 5-4
- sort order 4-4
- starting and stopping the server engine
  - from UNIX 3-5
  - from Windows NT 2-2
- starting the Windows NT Schedule service 2-1
- superuser 1-2
- symbolic links 5-4

## T

- testing user name for scheduled jobs 1-7
- thousands separator 4-3
- times 4-3
- Tools menu
  - adding a separator 1-21
  - adding an entry 1-20
  - changing order of entries 1-21
  - customizing 1-20
- trace files, viewing 1-14
- tracing server activity 1-14
- troubleshooting scheduled jobs 5-1
- tuning UNIX parameters 3-2

## U

- UNICODE 4-1
- UniData data sources, connecting 5-5
- unirpcd* daemon 3-4, 5-5

- UniVerse
  - accessing remote systems 3-5
  - issuing commands 1-22
  - remote systems 2-3
  - starting and stopping 2-2
- UniVerse Resource Service 2-2
- UNIX kernel parameters 3-2
- UNIX servers
  - and scheduling 5-3
  - restrictions 3-10
- user categories 1-5
  - assigning user groups 1-5
- User Rights Policy dialog box 5-3
- uvodbc.config* file
  - adding data sources 3-7
  - editing 2-3, 3-5

## V

- View trace file window 1-15
- viewing
  - locales 4-3
  - maps 4-1
  - trace files 1-14
- views, job schedule 5-1

## W

- Windows NT Schedule service 1-6, 2-1, 5-2
- Windows NT servers, and scheduling 5-2