

InfoSphere Information Server

Manually deleting a job in Information Server DataStage



© 2011 IBM Corporation

This presentation describes how to manually delete a DataStage® job for all versions of Information Server DataStage.

Objectives

- Remove job from local repository
- Remove job from xmeta (version 8 only)

Sometimes in DataStage, problems arise where a job cannot be deleted with the DataStage client. All versions of DataStage have a local repository located on the DataStage server. At version 8 of DataStage, there is also a database repository, named xmeta by default, that is located in a separate DB2®, Oracle, or SQL Server database. This presentation will show how to manually remove the job from both the local repository and from xmeta.

Local repository (1 of 3)

- Start uvsh shell (UNIX®/Linux®)
cd \$DSHOME
./dsenv
bin/uvsh

```
$ cd /opt/IBM/InformationServer/Server/DSEngine
$ ./dsenv
$ bin/uvsh
DataStage Command Language 8.1 Licensed Materials - Property of IBM
(c) Copyright IBM Corp. 1997, 2007 All Rights Reserved.
DSEngine logged on: Monday, February 07, 2011 11:40
>
```

- Start uvsh shell (Windows®)
cd \\IBM\InformationServer\Server\DSEngine
bin\uvsh

The next few steps will require that you work from a DataStage engine prompt. cd to your \$DSHOME directory where \$DSHOME is the path to the DSEngine directory. Next, on a UNIX or Linux system, you will need to source the dsenv file to set up your DataStage environment by typing:

```
./dsenv
```

Next, start a DataStage engine shell by typing: bin/uvsh. You must be in the DSEngine directory to do this. Do not attempt to run uvsh from within the bin directory. If you have run everything correctly, you will see a ">" prompt as in the example on this slide.

Local repository (2 of 3)

- Remove job from DS_JOBS file

LOGTO <project>

LIST DS_JOBS <jobname> ****Note Job Number****

DELETE DS_JOBS <jobname>

```

DataStage Command Language 8.1 Licensed Materials - Property of IBM
(c) Copyright IBM Corp. 1997, 2007 All Rights Reserved.
DSEngine logged on: Thursday, March 10, 2011 10:52

LOGTO support_us
LIST DS_JOBS delete_job
LIST DS_JOBS delete_job 10:56:55 03-10-11 PAGE 1

Job name.... delete_job
Description..
No..... 210
Category.... \JOBS

1 records listed.
DELETE DS_JOBS delete_job

1 records DELETED.

```

4

Manually deleting a job in Information ServerDataStage

© 2011 IBM Corporation

Once at a uvsh prompt, you will need to log in to the Project the job is located in. The engine commands are the same for UNIX, Linux, and Windows. Note that the uvsh prompt will type everything in uppercase automatically so you will need to hold the shift key down to do lowercase letters. The project and job names are case sensitive.

First, type:

LOGTO projectname

Next you will need to delete the job record from the DS_JOBS file. Check to see if the job record exists in DS_JOBS by typing:

LIST DS_JOBS jobname

Be sure that the case of your job name is correct. If you find the job in DS_JOBS, then you will need to delete it. Be sure to note the job number returned from the LIST command.

To delete the job from DS_JOBS, type:

DELETE DS_JOBS jobname

You should get a message returned that one record was deleted.

Local repository (3 of 3)

- Remove job files
 - For files RT_LOGxxx, RT_CONFIGxxx, RT_STATUSxxx
 - DELETE.FILE **DATA** file name
 - DELETE VOC file name
 - For file DS_TEMPxxx
 - DELETE.FILE **DATA** file name
 - xxx is job number
- Must use DATA key word
- Verify directories removed from project directory
- Version 7 and below – Delete is now complete

Next you will need to remove the job files. Each job file will have a RT_LOGxxx, RT_CONFIGxxx, RT_STATUSxxx, and DS_TEMPxxx file where xxx is the job number of the job to be deleted. You obtained this job number in the previous step in this presentation.

For each file, execute:

```
DELETE.FILE DATA file name
```

You **MUST** use the DATA key word or you will delete the shared dictionary file that is shared by all jobs and you will not be able to open the files for any jobs.

Next, for the files that begin with RT, execute:

```
DELETE VOC file name
```

Once this is complete or if the command errors that it can't find the file, verify that the files do not exist in the project directory on disk. These files open as directories at the operating system level. If you are at version 7 or lower, your job deletion is complete. If you are at version 8, continue to the next slides.

Xmeta database – Version 8 only

- DStageWrapper
 - Need to get from support for 8.0 through 8.1 FP1
 - Shipped with 8.1 FP2 and 8.5
 - Located in InformationServer/ASBNode/bin
 - UNIX/Linux – DStageWrapper.sh
 - Windows – DStageWrapper.bat
- At Version 8.5 and later the following changes are required to the DStageWrapper.sh script

Original:

```
eval exec "$JAVA" $J2EE_OPTS -Xmx384m -classpath "$CLASSPATH"  
com.ascential.dstage.proxies.CommandLine "$@"
```

Updated:

```
eval exec "$JAVA" $J2EE_OPTS -Xmx384m -classpath "$CLASSPATH"  
com.ascential.dstage.proxies.CommandLine "$@"
```

For version 8.0 and higher, you will need the DStageWrapper script in order to manually remove a job from the xmeta database. In version 8.0 through 8.1 fix pack 1, you will need to call Information Server Technical support to get a copy of this script and install instructions. Version 8.1 fix pack 2 and 8.5, both ship with a copy of the DStageWrapper. It is located in the InformationServer/ASBNode/bin directory for UNIX, Linux and Windows. For simplicity, this section of the presentation uses the UNIX and Linux syntax for DStageWrapper. The Windows syntax is exactly the same except the DStageWrapper command ends in a .bat instead of a .sh.

At version 8.5 and later, check the last line of the DStageWrapper script to ensure your script contains the proper syntax. After `com.ascential.dstage.proxies.CommandLine`, the syntax should read single quote, double quote, \$ sign, @ symbol, double quote, single quote. The fix as described in this slide corresponds to APAR JR40597 and is included in 8.7 fix pack 1 and future releases.

Xmeta database – List job RIDs – Version 8 only (1 of 2)

- List DataStage Job Definition:

Version 8.0:

```
./DStageWrapper.sh -query "select rid(x),x.DSNameSpace from x in
DataStageX::DSJobDef where x.name='<name_of_job>'"
```

Version 8.1 and 8.5:

```
./DStageWrapper.sh -query "select rid(x),x.DSNameSpace from x in
DataStageX::DSJobDef where x.name='<name_of_job>'" -domain <domain name>
-user <username> -password <password>
```

- Results returned:

c2e76d84.43058877.1e4e1e4e.79bc2379-2e7c-4cfb.bfe2.3d903e3de296 ← RID

ServerName: Project_Name

1 rows returned

DoQuery (select rid(x),x.DSNameSpace from x in DataStageX::DSJobDef where
x.name='delete_job') 37

You will need to find the unique RID for the job in the Job Definition table and in the Item table. RIDs are strings of characters used by Information Server to uniquely identify objects in the xmeta repository. For version 8.0, use the DStageWrapper shown in this slide to get the RID for the job. For version 8.1 and higher, you will need to add the `-user` and `-password` arguments. If you are not running the command on the domain server, you will also need to run the `-domain` argument.

If more than one row is returned, use the column after the unique RID to correctly identify the row. This column contains `server_name>:<project_name>`.

You will see multiple rows returned if you have a job with the same name in multiple projects. Make sure you get the RID for the job in the correct project. Copy and paste the RID into notepad because you will need it in the next two steps.

Xmeta database – List job RIDs – Version 8 only (2 of 2)

- List DataStage Job item:

Version 8.0:

```
./DStageWrapper.sh -query "select rid(x) from x in DataStageX::DSItem where  
x.ReposId='<RID_from_last_step>'"
```

Version 8.1 and 8.5:

```
./DStageWrapper.sh -query "select rid(x) from x in DataStageX::DSItem where  
x.ReposId='<RID_from_last_step>' -domain <domain name> -user <suite admin user>  
-password <suite admin password>
```

- Results returned:

c2e76d84.78be1842.39083908.62c23f62-27aa-4a43.9e5e.dcf403dc5e9b

RID

ServerName: Project_Name

1 rows returned

```
DoQuery (select rid(x) from x in DataStageX::DSItem where  
x.ReposId='c2e76d84.43058877.1e4e1e4e.79bc2379-2e7c-4cfb.bfe2.3d903e3de296') 51
```

Using the RID from the last query, you will now start a second query to gather the DSItem RID. For version 8.0, use the DStageWrapper command in this slide to obtain the RID. In this command, RID_from_last_step is the RID found in the last query.

For version 8.1 and higher, you will need to add the –user and –password arguments. If you are not running the command on the domain server, you will also need to run the -domain argument. Copy and paste the RID into notepad because you will need it in the next step.

Xmeta database – Delete job – Version 8 only

- Delete Job definition and item

Version 8.0

```
./DStageWrapper.sh –delete DSJobDef <first_RID_obtained>
```

```
./DStageWrapper.sh –delete DSItem <second_RID_obtained>
```

Version 8.1 and 8.5

```
./DStageWrapper.sh –delete DSJobDef <first_RID_obtained> -domain <domain  
name> -user <suite admin user> -password <suite admin password>
```

```
./DStageWrapper.sh –delete DSItem <second_RID_obtained> -domain <domain  
name> -user <suite admin user> -password <suite admin password>
```

- Job should now be deleted

Now you can issue repository deletes using the two RIDs obtained in the earlier steps. For version 8.0, the DStageWrapper command is:

```
./DStageWrapper.sh –delete DSJobDef first_RID
```

```
./DStageWrapper.sh –delete DSItem second_RID
```

For version 8.1 and higher, you will need to add the –user and –password arguments. If you are not running the command on the domain server, you will also need to run the -domain argument. You should see a confirmation that both the DataStage Item and DataStage Job Definition were deleted.

The job should now be deleted.

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, DataStage, DB2, and InfoSphere 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 other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. Windows, and the Windows logo are registered 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.

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

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2011. All rights reserved.