

InfoSphere DataStage

Diagnosing a hung Information Server V8.1 – 11.3 DataStage parallel job on Windows

© 2014 IBM Corporation

This presentation discusses how to diagnose hung Information Server DataStage® parallel jobs on Windows®. This presentation is relevant for Information Server version 8.1 through version 11.3.

Objectives

- Rule out known issues
- Determine job is hung
- Environment variables to set
- Collecting information during a hang

The objectives of this presentation are to rule out known issues, show how to determine whether a parallel job is hung, and if so, it describes the environment variables that need to be set and what information to collect during the hang.

Known issues (1 of 13)

- Technote covers guidelines for configuring Windows 2003 and Windows 2008
 - Configuring a Microsoft® Windows Server to run InfoSphere® Information Server
<http://www.ibm.com/support/docview.wss?uid=swg21419242>

- Technotes cover guidelines for anti-virus software
 - Guidelines for anti-virus programs and security software for InfoSphere Information Server
<http://www.ibm.com/support/docview.wss?uid=swg21566611>
 - InfoSphere Information Server guidelines for McAfee HIPS
<http://www.ibm.com/support/docview.wss?uid=swg21576550>

The technotes that are displayed on this slide describe guidelines to configure your Windows server and anti-virus software. Anti-virus software can impact performance and interfere with Information Server. It is important to exclude Information Server directories from scans. See the technotes that are displayed on this slide for more information.

Known issues (2 of 13)

- DataStage 8 only
 - MKS Toolkit can be source of job hangs
 - Check for Nutcracker errors in Windows application event log
 - Global critical section Nut4SemIdC is being cleaned up with waiter count 1. [nutsv4.exe (.\globcs.c:767) PID=pppp TID=tttt]
 - Process pppp appears to have hung. Cleaning up global critical sections that are held by this process. [nutsv4.exe (sem.cpp:1160) PID=pppp TID=tttt]
 - Failed to create fork() child process. [forktest.exe (fork.c:719) PID=5780 TID=5088]
 - Upgrading to MKS Toolkit 9.4 FP1 + HotFix CFS31959


Level	Date and Time	Source	Event ID	Task Category
Error	9/10/2013 11:32:06 AM	NuTCRACKER 4	16000	NuTCCleanup
Error	9/10/2013 11:30:46 AM	NuTCRACKER 4	16004	NuTCCleanup
Error	9/10/2013 11:28:46 AM	NuTCRACKER 4	16005	NuTCCleanup
Error	9/10/2013 11:28:46 AM	NuTCRACKER 4	16004	NuTCCleanup
Error	9/10/2013 11:29:06 AM	NuTCRACKER 4	16004	NuTCCleanup
Error	9/10/2013 11:29:26 AM	NuTCRACKER 4	16004	NuTCCleanup
Error	9/10/2013 11:29:06 AM	NuTCRACKER 4	16005	NuTCCleanup
Error	9/10/2013 11:28:26 AM	NuTCRACKER 4	16005	NuTCCleanup

Event 16000, NuTCRACKER 4	
General	Details
Global critical section Nut4LockNo is being cleaned up with waiter count 1. [nutsv4.exe (.\globcs.c:767) PID=1196 TID=1268]	

At version 8 of DataStage, the MKS Toolkit can be the source of job hangs. You should upgrade any busy Windows system to the latest MKS Toolkit version, more so if the errors displayed on this slide are seen in the Windows application event log. Upgrading to MKS Toolkit 9.4 FP1 + HotFix CFS31959 resolves the errors that are displayed on this slide.

Known issues (3 of 13)

- Check version of MKS Toolkit on engine tier
 - Select Control Panel > Configure MKS Toolkit
 - Select Support Information tab
 - Click Category drop down list
 - Select File Versions
 - Find version for nutsys4.dll for 64 and 32 bit
 - Version after FP1 + HotFix CFS31959 is applied
 - nutsys4.dll 64 and 32 bit => 4.69.0104



If the exact version of the MKS Toolkit is not known, follow the steps that are displayed on this slide to check the version of MKS Toolkit installed on the engine tier. When FP1 + HotFix CFS31959 is applied, the nutsys4.dll is version 4.69.0104 for both 64 and 32 bit.

Known issues (4 of 13)

- To upgrade to MKS Version 9.4 FP1 + Hotfix
- IS 9.1 and later
 - MKS Version 9.4 FP1 + Hotfix CFS31959 is included
- IS 8.7FP1
 - Where 8.7 media includes FP1 (Not IS 8.7 with FP1 applied) MKS Version 9.4 FP1 + Hotfix CFS31959 is included
- IS 8.5 and 8.7
 - JR41641: UPDATE MKS TOOLKIT TO VERSION 9.4 FP1 + Hotfix CFS31959
<http://www.ibm.com/support/docview.wss?uid=swg1JR41641>
 - Available on Fix Central
- IS 8.0.1 and 8.1
 - JR41654: UPGRADE INFORMATION SERVER 8.0.1 AND 8.1 TO MKS TOOLKIT 9.4FP1 WITH HOTFIX CFS31959
<http://www.ibm.com/support/docview.wss?uid=swg1JR41654>
 - Request from IBM Support

If you determined that you need to upgrade to MKS version 9.4 FP1 + Hotfix CFS31959, apply the appropriate patch for your releases of Information Server. This slide displays the fix that is required for the different versions of DataStage. The MKS upgrades use the MKS installer, not the suite patch installer. Follow directions carefully. Starting at Information Server 9.1, MKS is no longer used in the parallel engine layer, except where scripts are required. Instead, native Windows APIs are called. MKS is still installed since the suite relies on MKS shell services and system tools.

Known issues (5 of 13)

- Errors are normal for a busy system
 - Cleaning up process table entry 358 for process pppp, exit code 0x000000ff, name unknown. [nutsrv4.exe (.\process.c:347) PID=pppp TID=ttt]
 - SIGKILL signal has not caused process to die. Child watcher committing suicide. [osh.exe (.\ncchild.c:544) PID=pppp TID=ttt]

Level	Date and Time	Source	Event ID	Task Category
Error	9/10/2013 3:27:30 PM	NuTCRACKER 4	11033	Service
Error	9/10/2013 3:27:30 PM	NuTCRACKER 4	11033	Service
Error	9/10/2013 3:27:30 PM	NuTCRACKER 4	11033	Service
Error	9/10/2013 3:27:42 PM	NuTCRACKER 4	11033	Service
Error	9/10/2013 3:27:30 PM	NuTCRACKER 4	11033	Service
Error	9/10/2013 3:27:30 PM	NuTCRACKER 4	11033	Service
Error	9/10/2013 3:27:30 PM	NuTCRACKER 4	11033	Service
Error	9/10/2013 3:27:30 PM	NuTCRACKER 4	11033	Service

General	Details
Cleaning up process table entry 36 for process 264, exit code 0x00000001, name unknown. [nutsrv4.exe (.\process.c:347) PID=1196 TID=1268]	

It is important to note that not all errors in the Application log indicate a problem. This slide displays errors from the Windows application event log that are normal for a busy Windows system and do not indicate an issue.

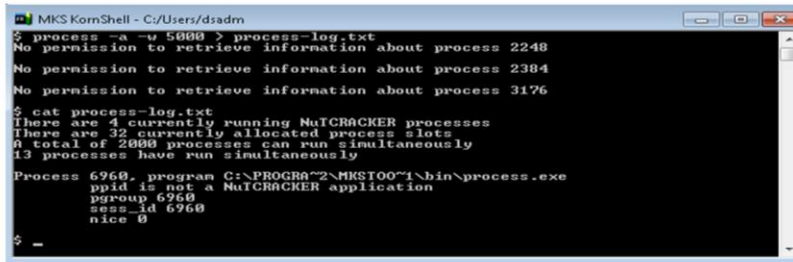
Known issues (6 of 13)

- MKS Toolkit Security ID setting causes Windows system lockup
 - DataStage parallel engine can lock up and you cannot start or compile parallel jobs
 - Existing parallel jobs keep running
 - You can start server jobs
 - Window functionality like Task Manager, Explorer still works
 - MKS shell commands hang
- Each process that is created by DataStage client on a Windows Server requires a Security ID token in MKS Toolkit
- On very busy systems, it is necessary to increase number of SSIDs above default

Another issue that you might encounter is the InfoSphere DataStage parallel engine on Microsoft Windows might 'lock up'. Windows functionality, for example Explorer, Task Manager, and so on, will still work. Existing parallel jobs, which are seen as OSH.exe in task manager, keep running, but you cannot start or compile any parallel jobs. However, you can start InfoSphere DataStage server jobs. Each process that is created by the DataStage client on a Windows server requires a Security ID token in the MKS Toolkit. On very busy systems, it is possible to run out of tokens, which causes the parallel engine to lock up. It is necessary to increase the number of SSIDs above the default to prevent lock up.

Known issues (7 of 13)

- Diagnosing issue
 - Obtain an MKS NuTCRACKER process report dump
 - Open Korn shell by clicking **Start->Run->ksh**
 - From Korn shell enter command: `process -a -w 5000 > process-log.txt`



```
MKS KornShell - C:/Users/dsadm
$ process -a -w 5000 > process-log.txt
No permission to retrieve information about process 2248
No permission to retrieve information about process 2384
No permission to retrieve information about process 3176
$ cat process-log.txt
There are 4 currently running NuTCRACKER processes
There are 32 currently allocated process slots
A total of 2000 processes can run simultaneously
13 processes have run simultaneously
Process 6960, program C:\PROGRAM~2\MKSTOO~1\bin\process.exe
ppid is not a NuTCRACKER application
pgroup 6960
sess_id 6960
nice 0
$
```

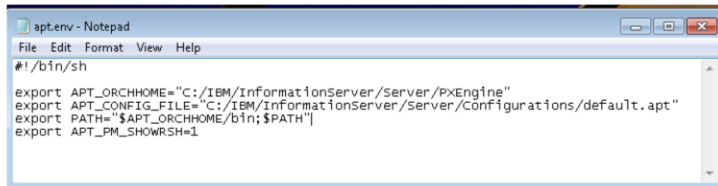
Note: It is possible this command might hang and not produce results

The process command reports how many NuTCRACKER platform processes are running and displays information on active NuTCRACKER platform processes. In the example that is displayed on this slide, there are four NuTCRACKER processes currently running. It also shows how many processes can run simultaneously.

It is possible the process command might hang and not produce any results. If this happens, continue to the next test.

Known issues (8 of 13)

- Check that you can successfully run a parallel job
 - Create an environment file in PXEngine directory
 - Use an editor such as notepad
 - Save file as apt.env



```
apt.env - Notepad
File Edit Format View Help
#!/bin/sh
export APT_ORCHHOME="C:/IBM/InformationServer/Server/PXEngine"
export APT_CONFIG_FILE="C:/IBM/InformationServer/Server/configurations/default.apt"
export PATH="$APT_ORCHHOME/bin;$PATH"
export APT_PM_SHOWRSH=1
```

The next test is to run a simple parallel job from a command line. First, create an environment file in the PXEngine directory. By default this is C:\IBM\InformationServer\Server\PXEngine. Create the file using an editor such as Windows Notepad and add the lines that are displayed on this slide. Then, save the file as apt.env. Note, if you are using Windows Notepad be sure to set the “**Save as type**” to “**All Files**” and enter **apt.env** as the File name. Do not use a .TXT file name extension.


Known issues (9 of 13)

- Test parallel engine by executing a simple parallel job
 - Invoke Korn shell select Start> Run> ksh
 - Change to PXEngine directory


```
cd C:/IBM/InformationServer/Server/PXEngine
```
- Source apt.env environment file


```
./apt.env
```
- Run test


```
osh "generator -schema record(a:int32) | peek"
```



11

Diagnosing a hung Information Server V8.1 - 9.1 DataStage parallel job on Windows

© 2014 IBM Corporation

After creating the environment file, test the parallel engine by executing the simple parallel job from the shell or command level. To do so, invoke the Korn shell by selecting Start> Run> ksh. Then, change the directory to the PXEngine directory. By default C:\IBM\InformationServer\Server\PXEngine is the default directory. Note from the Korn shell prompt, you need to use forward slashes instead of backward slashes. Also, adjust the path if you installed to a different directory or drive letter. Now source the apt.env file you created by typing “. ./apt.env”. The last step is to run the test parallel job using the osh command that is displayed on this slide.

The job should start normally and you should not experience any hang or error message dialog box. If an error dialog opens with the error, “Security initialization failed” when trying to run the osh test, you need to increase the number of MKS NuTcracker security IDs as described in the next slides.

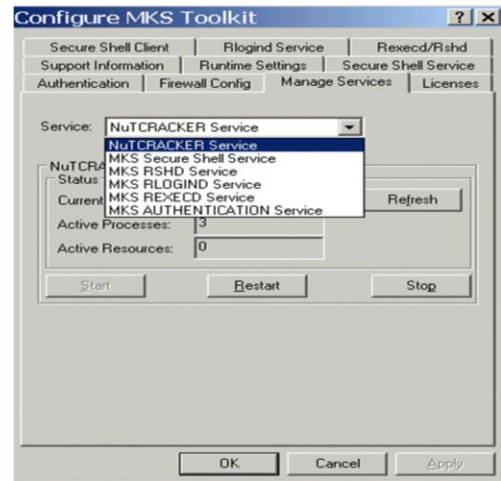
Known issues (10 of 13)

- Resolving issue
- Problem can be avoided by changing a parameter in MKS Control Panel
- To change parameter
 - Ensure that there are no uvsh.exe or osh.exe processes running on server
 - This prevents changed settings from taking effect
 - Stop all DataStage services
 - InfoSphere Engine Resource Service
 - DataStage Telnet Service
 - DSRPC Service

The issue can be resolved by changing the Max Number of Security ID parameter in the MKS control panel. To do this, you must be sure that there are no uvsh.exe or osh.exe processes running. If these processes are running, it prevents the changes from taking effect. Next, you need to stop all of the DataStage services.

Known issues (11 of 13)

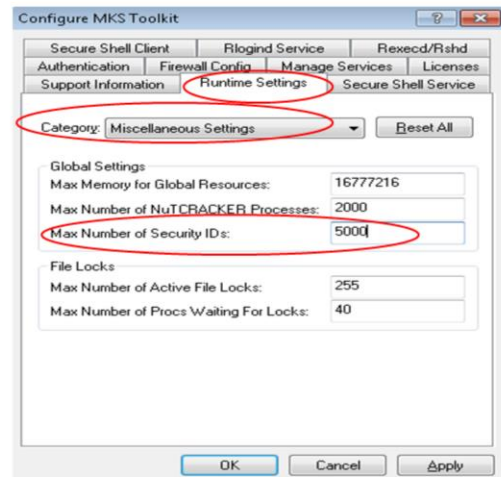
- Resolving issue continued
 - Open Microsoft Windows control panel
 - Click Configure MKS Toolkit
 - Select 'Manage Services' tab
 - Starting with bottom service shown in dropdown, stop each MKS Service
 - Click 'Refresh' button
 - 'Active Processes' box should display zero



Next, you need to open the Microsoft Windows control panel and click Configure MKS Toolkit. Next, click the Manage Services tab. Click the Services drop down. Starting at the bottom of the list and working backwards, stop each MKS service that is listed in the drop-down. Click the refresh button after all the services have been stopped. The Active Processes box should now display 0.

Known issues (12 of 13)

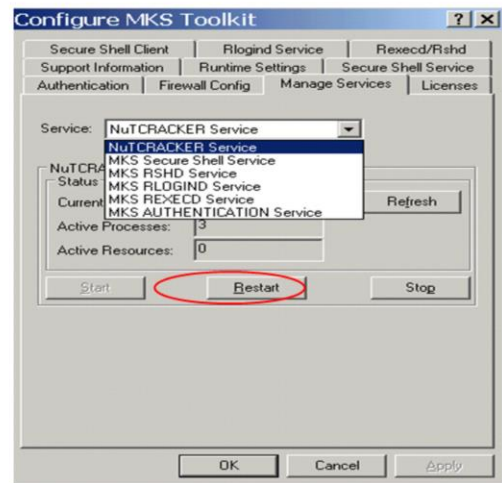
- Resolving issue continued
 - Select 'Runtime Settings' tab
 - Select 'Miscellaneous Settings' from Category dropdown
 - Change 'Max Number of Security IDs' to 5000



Next, click the Runtime Settings tab. Click the Category dropdown and select Miscellaneous Settings. Change the Max Number of Security IDs to 5000. The default is 2500.

Known issues (13 of 13)

- Resolving issue continued
 - Select 'ManageServices' tab
 - Restart each MKS service in order that is shown in dropdown, starting at top
 - Restart DataStage Services



After the parameter is changed, restart the MKS services and then restart the DataStage services. Select the ManageServices tab and select each MKS service in the list starting at the top and working down the list. For each service, click the service name in the drop-down and then click the Restart button. Click the OK button when finished and restart all of the DataStage services.

Determining the parallel job is hung (1 of 6)

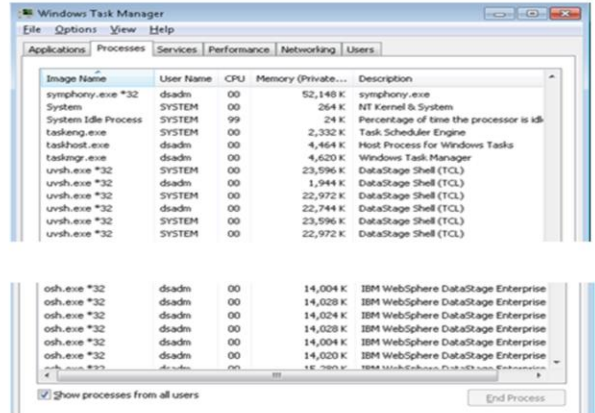
- Status in DataStage Director Client is running
- Check job progress
 - Select Tools > New Monitor
 - Check whether Num rows are increasing

Stage/Link name	Link type	Status	Num rows	Started at	Elap
Row_Generator_0		Running	0	3:53:07 PM	
Transformer_2		Running	0	3:53:08 PM	

The next topic this presentation discusses is how to determine whether a parallel job is hung. If a job is hung, the DataStage Director shows that the job is running but the job monitor will not show any progress. First, you need to open the DataStage Director, click the Tools menu, and then click New Monitor. If rows are still being processed, even if it is very slow, the job is not hung.

Determining the parallel job is hung (2 of 6)

- Check for processes at operating system level
- Open Task Manager on engine tier and look for
 - osh.exe
 - uvsh.exe
- Aborted job
 - No osh processes
 - Status file was not updated
 - Clear status file from Director
- Difficult to associate osh processes with a particular running job
 - Show pids for jobs
 - Set APT_PM_SHOW_PIDS



If the job appears to be hung, the next step is to check for job processes at the operating system level. Open Task Manager on the engine tier and look for uvsh.exe and osh.exe processes. If no processes are seen, then most likely the job failed, but was not able to update the status of the job before terminating. This is not considered a hung job. Clear the status file from DataStage Director.

If other jobs are running on the server, it is difficult to distinguish the osh processes associated with the hung job. The next section describes how to set the environment variable APT_PM_SHOW_PIDS. Setting this environment variable causes the pids to be written to the job log, which can then be used to find the osh processes at the operating system level.

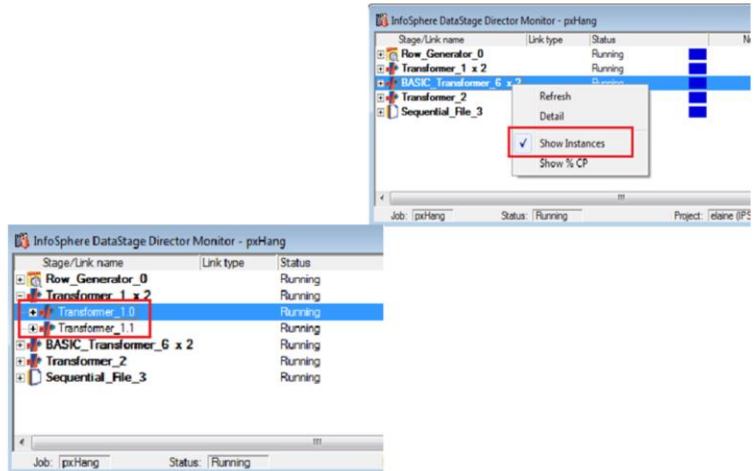
Determining the parallel job is hung (3 of 6)

- If failure can be reproduced at 9.1 and above, use new feature to generate stack trace
- Add user-defined environment variables
 - **APT_DUMP_STACK** - Setting this to 1 enables basic stack trace dump
 - **APT_DUMP_STACK_DIRECTORY** - Dump files are created in specified directory
- After setting **APT_DUMP_STACK** , feature is automatically invoked when an unrecoverable exception occurs
- A new feature to generate stack traces for Parallel jobs at version 9.1 of DataStage
- Technote:
A new feature to generate stack traces for Parallel jobs at version 9.1 of DataStage
<http://www.ibm.com/support/docview.wss?uid=swg21639558>

If the failure can be reproduced, starting at version 9.1 of DataStage, there is a new facility to generate stack traces and capture other valuable information for parallel jobs. The feature can be invoked by adding the user-defined environment variables **APT_DUMP_STACK** and **APT_DUMP_STACK_DIRECTORY**. Set **APT_DUMP_STACK** to 1 to enable basic stack trace dump. Set **APT_DUMP_STACK_DIRECTORY** to a valid path where files are written; if undefined or not set to a valid path then the dump files default to %TEMP%. If the job is successful, a dump is not created, therefore, you can leave this set to capture a dump for an intermittent issue.


Determining the parallel job is hung (4 of 6)

- Find pids using Monitor
 - Right-click
 - Select Show instances

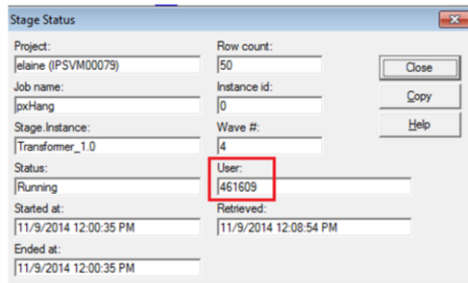


As mentioned previously, if other jobs are running on the server, it is difficult to distinguish the osh processes associated with the hung job but the pids can be found using the monitor window. First, right-click and select show instances if not already checked. If running with a multi-node configuration file, you see each of the instances. The example that is displayed on this slide shows two nodes.

Determining the parallel job is hung (5 of 6)



- To find pid
 - Right-click stage instance
 - Select Detail
 - pid is listed under User
 - Repeat this for each instance
- Use Task Manager on Windows to see processes



Next, to find the pid, select the stage instance, right-click, and select Detail. In the Stage Status dialog box, the pid is listed under User. Repeat this for each stage and instance. You can then use Task Manager to look for the processes.

Determining the parallel job is hung (6 of 6)

- Other reasons for leftover osh processes
 - DataStage engine stopped while parallel jobs running
 - Information Services Director job not undeployed
 - Parallel job stopped from DataStage Director when still in startup/handshake phase
 - Clean up left over processes using Task Manager

There are other reasons that leftover osh processes might be seen even though the job is not hung. One example is when the DataStage engine is stopped while parallel jobs are running or an Information Service Director™ job is not undeployed before stopping the DataStage engine. To prevent this, check for running jobs and undeploy any Information Service Director jobs before stopping the DataStage engine.

Another reason osh processes might be left over is when a parallel job is stopped from the DataStage Director during the startup phase. This is the phase where the conductor communicates with section leaders, the section leader communicates with players, or players communicate with players.

In both of these cases the leftover processes can be cleaned up using Task Manager.

Environment variables to set (1 of 2)

- Set environment variables at job or project level
 - APT_PM_SHOW_PIDS=True
 - APT_DUMP_SCORE=True
- Variables that are listed under Parallel > Reporting section

Environment Variables

The following categorized environment variables are defined in this project. Either set a default value for an existing env or add a new environment variable to the user defined category. When you export or import environment variables, values are preserved.

Categories:	Name	Prompt	Value
General			
Customize			
Parallel			
Operator Specific			
Reporting	APT_DUMP_SCORE	Report score	False
Reporting	APT_MSG_FILELINE	Extra logging information	False
Reporting	APT_NO_JOBMON	Disable job monitor	False
Reporting	APT_PERFORMANCE_DATA	Performance data directo	False
Reporting	APT_PM_PLAYER_MEMORY	Report player memory alloc	False
Reporting	APT_PM_PLAYER_TIMING	Report player calls	False
Reporting	APT_PM_SHOWRSH	Show RSH commands	False
Reporting	APT_PM_SHOW_PIDS	Show internal PIDs	False
Reporting	APT_RECORD_COUNTS	Report record counts	False
Reporting	APT_SHOW_COMPONENT_CALLS	User-overloadable functio	False
Reporting	APT_STARTUP_STATUS	Extra startup messages	False
Reporting	OSH_DUMP	Report step description	False
Reporting	OSH_ECHO	Report step specification	False
Reporting	OSH_EXPLAIN	Report terse step descript	False
Reporting	OSH_PRINT_SCHEMAS	Report schemas	False
Compiler			
User Defined			

Diagnosing a hung Information Server V8.1 - 9.1 DataStage parallel job on Windows

22

© 2014 IBM Corporation

Once it is determined that the job is hung, setting the environment variables APT_PM_SHOW_PIDS and APT_DUMP_SCORE to True will provide the information needed the next time that the job hangs. These environment variables can be set at the job or project level.

Environment variables to set (2 of 2)

- Create user-defined environment variable by way of DataStage Administrator
 - On non-production environments create
 - DS_PXDEBUG
 - Leave default value blank at project level
 - Set default value to 1 at job level
 - On Production environment
 - If able to compile a job, set DS_PXDEBUG to 1 at job level

Next, create a user-defined environment variable that is called DS_PXDEBUG if the job is not running in a production environment, and set the default value to 1 at the job level. If it is a production environment and there is the ability to compile a job, set DS_PXDEBUG to 1 at the job level. DS_PXDEBUG should not be set at the project level because it will greatly impact the performance of jobs and add a significant amount of debug information to all the job logs.

Collecting information during a hang (1 of 3)

- Send export of detailed job log
- Send export of job design
 - Example - *.dsx, *.isx, *.xml
- With DS_PXDEBUG set
 - zip and send Debugging/<job_name> directory that is found under project directory
- Send an export of Windows application and system event log

Once the environment variables are set and the hang is reproduced, the next step is to collect the log information. It is important to send the detailed job log, an export of the job design of the hanging job, an export of the Windows application and system event log. The export of the application and system event logs need to include the time frame during the job hang.

Collecting information during a hang (2 of 3)

- Send ISALite Basic System Summary
 - Can be done at any time before or after hang
 - Download the ISALite for InfoSphere Information Server tool
<http://www.ibm.com/support/docview.wss?uid=swg24022700>
- **Alternately** at IS 8.5 and later, send ISALite Job Log Collection
 - Include
 - Job logs and job design in *.isx format
 - Version.xml, .odbc.ini, dsenv, uvconfig
 - DSParams file
 - Run from Engine tier
 - Does not require Designer or Director client

Collect and send an ISALite Basic System Summary. This can be done at any time before or after the hang.

Alternately, at Information Server version 8.5 and later, send an ISALite Job Log Collection. This collects all the job-related information in one step and it includes additional information such as Version.xml, .odbc.ini, dsenv, uvconfig, DSParams, and more. It also includes a collection report. An export of the Windows application and system event log is still required in addition to the ISALite Job Log Collection.

Collecting information during a hang (3 of 3)

- Collect stack trace on each process ID (pid) seen in job log
 - APT_PM_SHOW_PIDS writes pids to job log
- One parallel job can have large number of osh processes
 - Script to collect stack trace on all pids available for Windows 64 bit
 - Script assumes that all osh processes are on same machine
 - Script requires windbg utility from Microsoft be installed on engine tier
- Download PXOSHCallStackCollectorWin64.zip from IBM Education Assistant website

If the job processes are hung, the call stack from each osh process in the job provides information on the state of the osh process at the time of the hang. Collecting the stack traces is critical for debugging the issue. A stack trace is needed for each pid that is listed in the job log.

There can be many pids that need stack traces. A script is available to automate this process on Windows 64 bit. The script assumes that all osh processes are on the same machine. The script is available for download from this IBM Education Assistant site. The file name is PXOSHCallStackCollectorWin64.zip and it includes a README.txt with instructions on how to run the script.

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, DataStage, InfoSphere, and Service Director 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](#)" at <http://www.ibm.com/legal/copytrade.shtml>

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

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 2014. All rights reserved.