

InfoSphere Information Server

Using dslogger for Information Server version 8 and 9

© 2014 IBM Corporation



This presentation will discuss how to use the dslogger utility to collect information for InfoSphere® Information Server.

Objectives

- Overview
- How to install
- How to run
- How to collect

The objectives of this presentation are to provide an overview about the dslogger utility, how to install it, how to run it, and how to collect information generated by the tool.

Overview

- Non-supported utility to collect data
- Provided "as is"
- Collects statistics in background
- Easy to install and run
- Minimal impact on performance
- Log files are circular

dslogger is a utility not officially supported by IBM. It is provided "as is" by IBM Support to collect DataStage® information in certain cases where problems are hard to reproduce at will. There is no guarantee that this tool will help to determine performance bottlenecks. This presentation assumes that you have already obtained this utility from IBM Support.

The dslogger utility is a program that runs in the background and collects statistics of the Engine Tier.

It is simple to install and execute. It only runs on UNIX® or Linux® platforms. It has an insignificant impact on the performance of the machine so it can be started safely. It creates circular logs, meaning that only the most recent logs are retained, while older log files are deleted. You can control the size of the log files, the number of logs you want to keep, and the frequency used to collect the statistics.

How to install

- Select a disk partition with free space
- Create a folder to place dslogger
 - Example

```
$ mkdir dslogger_dir
```
- Copy utility and uncompress if necessary
 - Example

```
$ cp dslogger.tar.Z dslogger_dir  
$ cd dslogger_dir  
$ uncompress dslogger.tar.Z  
$ tar -xvf dslogger.tar
```
- Create a sub-folder to collect logs
 - Example

```
$ mkdir logs
```

The dslogger should be placed in a partition with enough free space to collect logs. Create a folder in that partition to place the dslogger utility. Copy the file provided by Support and uncompress if necessary. Create a sub-folder to collect logs.

How to run (1 of 3)

- Login as dsadm
- Source dsenv
 - Example

```
$ cd $DSHOME
$ . ./dsenv
```
- Change location to dslogger folder and run
 - Example

```
$ cd dslogger_dir
$ nohup ./dslogger -logdir ./logs -interval 60 -iterations 60 -purge 12 -continuous
</dev/null >/dev/null &
```

The dslogger script should be run by the DataStage administrator, dsadm. Before running it, be sure to source the dsenv file. Change directory to the location of dslogger and then run it using the command displayed on this slide.

How to run (2 of 3)

▪ dslogger syntax

```
$ nohup ./dslogger -logdir <path> [-interval <number>] [-iterations <number>] [-purge <number>] [-continuous] </dev/null >/dev/null &
```

where

```
-logdir <path> : provides the location where log files will be placed  
-interval <number> : defines the sleep time (in seconds) between each  
iteration of data collection  
-iterations <number> : defines the number of iterations that will be stored  
per log file before a new log file is created  
-purge <number> : defines the number of maximum log files that will be  
kept before the oldest file is deleted  
-continuous : forces dslogger to continuously collect data
```

dslogger has several options that allow you to control the number of log files you will keep, the number of iterations each file will keep and the frequency of each iteration. It is recommended to run it using the command `nohup` so the process keeps running even if the terminal where it was issued is closed. It is also recommended to redirect the standard output to `/dev/null` to prevent unnecessary messages on the screen. The example on this screen also shows "`< /dev/null`". This will prevent the command from hanging in the event that one of the commands prompts for input. The parameter `logdir` specifies the location of the folder where the logs are created. The parameter `interval` defines the sleep time in seconds between each iteration of data. The parameter `iterations` defines the number of iterations that will be stored per log file before a new log file is created. The parameter `purge` defines the number of logs that will be kept before the oldest file is deleted. The parameter `continuous` should be used to force dslogger to continuously collect data.

How to run (3 of 3)

- **Example 1. Capture 12 hours worth of information, one hour per log file, and statistics every 60 seconds:**

```
$ nohup ./dslogger -logdir ./logs -interval 60 -iterations 60 -purge 12 -continuous </dev/null  
>/dev/null &
```
- **Example 2. Capture 24 hours worth of information, one hour per log file, and statistics every five minutes:**

```
$ nohup ./dslogger -logdir ./logs -interval 300 -iterations 12 -purge 24 -continuous </dev/null  
>/dev/null &
```
- **Example 3. Capture 5 hours worth of information, 20 minutes per log file, and statistics every 30 seconds:**

```
$ nohup ./dslogger -logdir ./logs -interval 30 -iterations 40 -purge 15 -continuous </dev/null  
>/dev/null &
```

This slide provides some examples on how to run dslogger. Example 1 shows how to capture 12 hours worth of information, with one hour of information per log file, and statistics collected every 60 seconds. Example 2 shows how to capture 24 hours worth of information, with one hour of information per log file, and statistics collected every five seconds. Example 3 shows how to capture five hours worth of information, with 20 minutes per log file and statistics collected every 30 seconds.

How to collect information

- Tar files under logs folder
 - Upload files to PMR
 - Stop dslogger if necessary
 - Example
- ```
$ ps -ef|grep dslogger
13222 34655 dsadm dslogger
$ kill 13222
```

Once the condition you are trying to reproduce has occurred, tar the files under the dslogger logs folder. The most recent log is the one called dsinfo.log. The older logs are saved in dsinfo.log.N files where N is a number based on the value given for the -purge option.

Upload the files to the PMR. If you want to stop dslogger, find its process ID and run kill. It may take several seconds to terminate.



## Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, DataStage, 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>

Linux is a registered trademark of Linus Torvalds 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 2014. All rights reserved.