


IBM WebSphere eXtreme Scale V8.6

High Performance Extensible Logging and IBM Support Assistant Data Collector



© 2013 IBM Corporation

In this short presentation, you will learn how to use High Performance Extensible Logging and the IBM Support Assistant Data Collector, which are new features in IBM WebSphere® eXtreme Scale version 8.6.

Agenda

- High Performance Extensible Logging (HPEL)
 - Overview
 - Procedure
 - logViewer command
- IBM Support Assistant Data Collector (ISA DC)
 - Overview
 - Procedure

For each of the two features, there is an overview and a discussion of the steps needed to use that feature.

High Performance Extensible Logging (HPEL)

This section discusses High Performance Extensible Logging, which is referred to as “HPEL”.

Overview

- Available in both stand-alone and WebSphere Application Server managed eXtreme Scale
- Alternative to the basic log and trace facility (SystemOut.log/SystemErr.log, trace.log and activity.log)
- Provides a log data repository, a trace data repository and a text log file
- Enabled using the server properties file for container and catalog servers
- Disabled by default
- There are two scenarios in which you might want to use the new HPEL feature:
 - To take advantage of the HPEL logViewer command which allows you to filter existing logs for specific details
 - To increase log performance by storing only a binary representation of the systemout systemerr and trace logs

High Performance Extensible Logging is a new log and trace facility available in both the stand-alone and WebSphere Application Server-managed types of eXtreme Scale. HPEL provides a convenient mechanism for storing and accessing log, trace, System.err, and System.out information produced by the application server or applications. It is an alternative to the basic log and trace facility, which provides the JVM logs, diagnostic trace, and service log files commonly named SystemOut.log/SystemErr.log, trace.log, and activity.log. HPEL provides a log data repository, a trace data repository, and a text log file.

Rather than replacing the existing logging facility, HPEL is provided as an alternative that can be enabled using the server properties file for container and catalog servers. HPEL is not enabled by default. Once enabled, HPEL will handle all WebSphere eXtreme Scale logging. The resulting log files are placed in the HPEL repository location. This location is specified in the same server properties files for container and catalog servers.

Procedure

- Edit the server properties file for each catalog and container server with HPEL properties (full list is in the information center: <http://pic.dhe.ibm.com/infocenter/wxsinfo/v8r6/index.jsp>)
 - **hpelEnable** (Default: false)
 - **hpelRepositoryLocation** (Default is the runtime location)
 - **hpelEnablePurgeBySize** (Default: true)
 - **hpelMaxRepositorySize** (Default: 50 MB)
 - **hpelOutOfSpaceAction** (Default: PurgeOld)
 - **hpelOutputFormat** (Default: Basic)
 - **hpelIncludeTrace** (Default: False)
- Restart the servers. Previous logging information is now replaced by equivalent HPEL information
- Use the HPEL command-line log viewer to view your log files
 - From a command line, go to <WXS_install_root>/ObjectGrid/bin
 - Run logViewer.sh or logViewer.bat
 - Get help by running logViewer -help

Enabling and using HPEL is simple. First, edit the server properties file for each catalog and container server. The value for the **hpelEnable** property needs to be changed from false to true. You also need to set the values of some other properties that affect the output, such as location, maximum size, format, and tracing. A sample list of properties is shown here. You can find the complete list of properties in the eXtreme Scale version 8.6 information center. A link is provided.

After editing the server properties file, you need to restart the server. After this restart, any previous log information is now replaced with HPEL information. You can look at your log files by running the logViewer script in the ObjectGrid/bin directory.

Common logViewer commands

- **`./logViewer.sh -outLog ../logs/legacyFormat.log -minLevel INFO -maxLevel SEVERE`**
 - Create a legacy-format log file (legacyFormat.log) that contains only log records (INFO, WARNING, and SEVERE)
- **`./logViewer.sh -thread 0`**
 - View only the log records for thread 0
- **`./logViewer.sh -thread 0 >> messagesFromThread0.log`**
 - Store messages for thread 0 in local file named messagesFromThread0.log
- **`./logViewer.sh -level WARNING`**
 - View only WARNING messages
- **`./logViewer.sh -excludeLoggers com.ibm.*`**
 - Retrieve all log records not from loggers beginning with com.ibm
- **`logViewer -minLevel WARNING -maxLevel SEVERE -extractToNewRepository ../logs/newHPELRepository`**
 - Extract a repository of just WARNING and SEVERE messages and save the resulting file in a new directory
- **`./logViewer.sh -repositoryDir ../logs/newHPELRepository -outLog ../logs/newFormat.log`**
 - Export the contents of the repository to a text format log file

This slide shows the syntax for some frequently used logViewer commands. You can filter log messages, view records for specific threads, and export the log repository to a text format log file. A complete list of commands can be found in the information center.

Usage examples

- Goal - You want your JVM log files to be recycled once they get to 100 megabytes in size.
 - Update the **hpelMaxRepository** size to be 100 in the server properties file used by the referenced JVM
 - Update the **hpelEnablePurgeByTime** to be false and the **hpelEnablePurgeBySize** to be true
 - Restart the JVM
- Goal - If log files have gotten to the point of exhausting the allotted hard-drive space you want the JVMs to stop running.
 - Access the server properties file used by the JVM
 - Update the **hpelOutOfSpaceAction** to be StopServer
 - Restart the JVM
- Goal - You like the efficiency that binary files produce, but also want trace file data to be sent to the more human-readable text format
 - Access the server properties file used by the JVM
 - Update the **hpelIncludeTrace** setting to be true
 - Restart the JVM

This slide explains how to accomplish some goals you might have for your log files. In the first example, if your goal is for the log files to automatically recycle when they get to a certain size, you set the **hpelMaxRepository** size in the relevant server properties file. You also need to be sure that **hpelEnablePurgeByTime** is set to false and **hpelEnablePurgeBySize** is set to true.

The second example explains how you tell the server what to do if hard-drive space is running low. The third example explains how to enable tracing.

Remember that the server must be restarted in order for your changes to the HPEL property values to take effect.

IBM Support Assistant Data Collector

This next section covers IBM Support Assistant Data Collector, which is referred to as “ISA DC”.

Overview

- ISA DC is used for log gathering for the purpose of problem determination
- In previous versions of eXtreme Scale, the IBM Support Assistant Lite (ISA Lite) tool was used. ISA Lite is still shipped in the `wxs_home/isalite_wxs` directory.
- In eXtreme Scale V8.6, ISA DC is also installed: located in `wxs_home/isalite_dc`
- ISA DC improves ease of use for collecting data by remembering various inputs, reducing repetitive typing during console input
- ISA DC installs with both eXtreme Scale stand-alone and WebSphere Application Server-managed eXtreme Scale

ISA DC is new in eXtreme Scale version 8.6. IBM Support Assistant Lite, which is available in earlier versions of eXtreme Scale, is also available in version 8.6. The ISA Lite tool is located in the `isalite_wxs` directory. The new ISA DC tool is located in the `isalite_dc` directory. ISA DC is easier to use because it remembers inputs and requires less typing.

Procedure

- Start the tool. The tool runs in console mode by starting the launch script from the command line
 - Navigate to *wxs_home/isalite_dc/*
 - Run *isadc.sh* or *isadc.bat*
- You are prompted with a series of system configuration questions. Have this information ready
 - File name for saving the collected data
 - Java Home directory
 - ObjectGrid Home directory
 - Working directory used by eXtreme Scale
 - Location of additional script files used to start servers
- Now this eXtreme Scale environment related information should be bundled in the <file name supplied for saving the collected data>.zip
 - Log files
 - eXtreme Scale version information
 - Java version information
 - Information on /ObjectGrid directory structure (including what files are currently stored in various directories - does not mean saving the actual files)

You run the ISA DC tool from a command line by running the *isadc* script found in the *isalite_dc* directory. It is important to have your system configuration information ready so you can answer the prompts. You need to provide a file name for saving the collected data, the Java Home directory, and the ObjectGrid Home directory. You also need to enter the working directory used by eXtreme Scale and the location of any script files used to start servers. After the tool runs, the resulting file contains the log files, eXtreme Scale version, and information showing the ObjectGrid directory structure. Note that this will tell you the files that are in various directories; but it will not contain the actual files.

Run it on each system

- ISA DC gathers JVM configuration and logs on a per-system basis
- You have to run the command on each JVM that resides in the ObjectGrid topology (each catalog and each container server separately)

You have to run the ISA DC tool on each catalog and container server separately, in eXtreme Scale version 8.6.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_XS86_HPEL_ISA.ppt

This module is also available in PDF format at: ../XS86_HPEL_ISA.pdf

You can help improve the quality of IBM Education Assistant content by providing feedback.



Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, and WebSphere 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>

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 2013. All rights reserved.