



IBM Software Group

IBM WebSphere eXtreme Scale V7.0

Problem determination



@business on demand.

© 2009 IBM Corporation
Updated August 12, 2009

This presentation discusses problem determination for IBM WebSphere® eXtreme Scale V7.0.

Message prefixes

- **CWOBJ: eXtreme Scale core components**
 - ▶ Interactions between core components
 - ▶ Inter-server communications
 - ▶ Catalog server interactions
 - ▶ Administration
- **CWXQY: Query engine**
 - ▶ Query preparation and execution
- **CWPRJ: Entity projector**
 - ▶ Entity/Tuple data mapping
- **CWWSM: HTTP Session manager**
 - ▶ HTTP session filter operations and management

WebSphere eXtreme Scale produces four major classes of messages.

Messages beginning with the string “CWOBJ” are issued by core ObjectGrid components. These messages describe interactions between core components such as backing maps, grid servers, and catalog servers.

“CWXQY” messages come from the query engine. These messages relate to both entities and Java™ objects.

“CWPRJ” messages are issued by the Entity Manager projector components that convert data between Tuples and Entities.

“CWWSM” messages are issued by the WebSphere eXtreme Scale HTTP Session manager.

Log files

- Server log files are provided for review of catalog and container server messages
 - ▶ SystemOut.log – look for start confirmation messages, peer mode messages, replication messages and so on
 - ▶ SystemErr.log – look for error messages
 - ▶ Trace file – exists if trace was enabled with the traceSpec option
- For stand-alone installations, the *workingDirectory* property is the root for the logs and optional trace file
 - ▶ Default is <current_directory>/logs/<container_name>
- If eXtreme Scale was installed within the WebSphere environment, see WebSphere Application Server information center for more information about server log files

Server log files allow you to review messages issued during server start, runtime, and termination. All messages are logged into SystemOut.log, which you can use to look for confirmation messages from various components within eXtreme Scale. The SystemErr.log provides a summary of only error messages that have been reported.

For stand-alone installations, the location of the log directory is controlled by the “workingDirectory” property set in the server.properties file. If a working directory is not specified, the runtime will create a default *logs* directory under the “current” directory. The server’s logs are in a subdirectory with the same name as the server.

If trace is enabled for a grid server a trace file is also created in the server’s *logs* directory. You can separately direct the trace file to another location using the “traceFile” parameter.

If WebSphere eXtreme Scale was installed within a WebSphere environment, review the WebSphere Application Server information center for information about specifying log locations.

Methods to set trace

- Supply trace parameters in command invocation
- Server properties file
- WebSphere Application Server administrative console
- Programmatically at runtime
- Sample administrative utility - **xsadmin**



There are several ways to control tracing. You can supply parameters for tracing when servers are started using the `-traceSpec` parameter on the server start command. Another method uses a server properties file. If eXtreme Scale V7.0 is installed in the WebSphere environment, you can use the WebSphere administrative console to define the trace strings and trace file size, just like you trace other components within WebSphere. You can also programmatically start and stop tracing in your application during runtime. Finally, the sample administration utility, `xsadmin`, included with eXtreme Scale can set trace specifications and turn tracing on and off dynamically.

Trace parameters in command invocation

- Setting trace on a stand-alone server
 - ▶ **-traceSpec** – to specify trace specifications
 - ▶ **-traceFile** – to specify the path and trace file name

EXAMPLE:

```
startOgServer.sh <appropriate server parameters>  
-traceSpec ObjectGrid*=all=enabled  
-traceFile /home/user1/trace.log
```



This slide provides more detail about setting trace parameters during the start of a stand-alone server. You specify the *traceSpec* parameter to control which components are traced. More information about the available components is provided later in this presentation.

You can optionally provide the *traceFile* parameter to specify where you want the trace log to be written. If you do not specify *traceFile*, the trace file is written to the same directory as the *SystemOut.log* and *SystemErr.log* files.

The slide shows you an example of specifying both *traceSpec* and *traceFile* within the *startOgServer* command.

Server properties file and trace

- Create the property file with trace specifications

```
EXAMPLE: test.properties
traceSpec="ObjectGrid*=all=enabled"
systemStreamToFileEnabled=true
traceFile=C:/temp/trace3.log
```

Another way to set trace is to use a property file. First, create a property file and specify the trace specification with the *traceSpec* parameter. The property file must contain the required parameter "systemStreamToFileEnabled=true." You can optionally specify where you want the trace file to be written using the *traceFile* parameter.

How eXtreme Scale finds the property file

- Reference the property file in the command **-serverProps**
- Auto-discovery of well-named property file anywhere in the classpath
 - ▶ **objectGridServer.properties**
- System property
 - ▶ **objectgrid.server.props=<property file name>**
- Programmatically
 - ▶ **ServerFactory.getServerProperties()**
- Sample property files:
 - ▶ <extreme_scale_installation_root>/ObjectGrid/properties
 - ▶ <websphere_installation_root>/optionalLibraries/ObjectGrid/properties

8

Problem determination

© 2009 IBM Corporation

You can choose how to have the grid server recognize your property file. You can modify the script or command you use to start the server, specifying the *serverProps* to reference the property file you created. Or you can use a “well-named property file” name, *objectGridServer.properties*, which the system can locate if it is within the classpath. Finally you can define a system property, *objectgrid.server.props*, to point to your property file. The Extreme Scale information center also discusses a programmatic way to locate a property file using the *ServerFactory* object.

Sample property files are provided within the stand-alone eXtreme Scale installation in the <extreme_scale_installation_root>/ObjectGrid/properties directory. For an eXtreme Scale installation in a WebSphere environment, the sample property files are located in <websphere_installation_root>/optionalLibraries/ObjectGrid/properties.

Property file in a command invocation

- Example of specify a property file within a command invocation

```
startOgServer server1 <server start parameters> -serverProps  
../properties/test.properties -jvmArgs -cp <args to pass to JVM>
```

test.properties

```
traceSpec="ObjectGrid*=all=enabled"  
systemStreamToFileEnabled=true  
traceFile=C:/temp/trace3.log
```

Here is an example of specifying a property file within a command invocation. Notice the JVM arguments must be the last parameters within the command invocation.

Additional ways to set trace

- Use WebSphere Application Server administrative console
 - ▶ For eXtreme Scale installations in WebSphere environments
 - ▶ Use same way as other WebSphere component tracing
 - ▶ See the WebSphere Network Deployment information center for information about working with trace
- Use programmatic invocation during runtime
 - ▶ Set trace during runtime on an ObjectGridManager
 - ▶ See WebSphere eXtreme Scale Information Center for more information about starting a trace programmatically
- Administrative utility - **xsadmin**



If eXtreme Scale is installed within a WebSphere environment, use the WebSphere Application Server administrative console to set tracing for eXtreme Scale and to control the size and location of log files. See the WebSphere Application Server Network Deployment information center for details about how to perform this task.

You can programmatically set tracing within your application at runtime. This is covered in more detail in the WebSphere eXtreme Scale information center.

Finally, the sample administration utility, `xsadmin`, included with eXtreme Scale can set trace specifications and turn tracing on and off dynamically. This utility is discussed in a separate presentation and in the eXtreme Scale information center.

Trace string format

- **<ObjectGridComponent>=<level>=enabled | disabled**
- You can concatenate trace strings; use “:” as separator

Projector=all=enabled:ObjectGridEntityManager=all=enabled

- Use the * symbol to specify a wild card

ObjectGrid*=all=enabled



The trace string specification has three basic parts: the component to be traced, the trace level, and the “enabled” or “disabled” control string. You can specify more than one trace string specification by separating the strings with a colon. You can use the asterisk as a wild card to enable more than one trace component in one specification.

Trace components

- eXtreme Scale ObjectGrid components for tracing

Component	Description
ObjectGrid	<i>General core cache engine</i>
ObjectGridBalancer	<i>Dynamic deployment shard balance configuration</i>
ObjectGridCatalogServer	<i>General catalog service</i>
ObjectGridChannel	<i>Static deployment topology communications</i>
ObjectGridCORBA	<i>Dynamic deployment topology communications</i>
ObjectGridDataGrid	<i>AgentManager API</i>
ObjectGridDynaCache	<i>eXtreme Scale cache provider</i>
ObjectGridEntityManager	<i>EntityManager API</i>
ObjectGridEvictors	<i>ObjectGrid built-in evictors</i>
ObjectGridForward	<i>Client/Server request forwarding details</i>



The table on this slide and the next show the eXtreme Scale V7.0 components that can be traced and a brief description of the component.

Trace components

Component	Description
ObjectGridJPA	<i>XS-specific JPA; JPA loaders</i>
ObjectGridLocking	<i>ObjectGrid cache entry lock manager</i>
ObjectGridPlacement	<i>Catalog server shard placement service</i>
ObjectGridReplication	<i>Replication service</i>
ObjectGridRouting	<i>Client/Server routing details</i>
ObjectGridServerSessionCache	<i>ObjectGrid session cache</i>
ObjectGridStats	<i>ObjectGrid statistics</i>
ObjectGridStreamQuery	<i>Stream Query API</i>
QueryEngine	<i>Query engine for the Object Query API and EntityManager Query API</i>
QueryEnginePlan	<i>Query plan diagnostics</i>

* New in V7



This slide shows more of the trace components and associated descriptions.

Trace components

Component		Description
ObjectGridJPACache	*	<i>JPA cache plug-ins</i>
ObjectGridMBean	*	<i>Management beans</i>
OpenJPA	*	<i>Open Java Persistence API (JPA)</i>

* New for V7



This slide shows the trace components and associated descriptions that are new for eXtreme Scale V7.0.

The OpenJPA component is also available in the WebSphere Application Server EJB3 Feature Pack environment or in the WebSphere Network Deployment V7 environment.

Trace levels

- Common trace levels
 - ▶ **all** – all trace records
 - ▶ **debug** – debug trace records
 - ▶ **entryExit** – method entry and exit trace records
 - ▶ **event** – event trace records
- <ObjectGridComponent>=<**level**>=enabled | disabled
- Example

ObjectGridPlacement=entryExit=enabled

- If unsure of level, use “all”



The trace levels provide you some control of the amount of detail you want traced. If you do not know what trace level to specify when providing diagnostic traces to IBM support, use “all” to ensure that the necessary data is included in the trace log.

FFDC files

- FFDC files are for IBM support to assist in debug
- Might be requested by IBM support for problem diagnosis
- Appear in a directory labeled **ffdc** and contain files that resemble:
 - ▶ server2_exception.log
 - ▶ server2_200802080_07.03.05_10.52.18_0.txt
- FFDC's are written for servers only
- Properties directory must be in classpath for stand-alone eXtreme Scale installations for FFDC's

17

Problem determination

© 2009 IBM Corporation

WebSphere eXtreme Scale provides “First-Failure Data Capture” data, also called FFDC data. This data is primarily intended to be used by IBM support. FFDC data files are written to an “FFDC” subdirectory within the logging path of the eXtreme Scale installation. The FFDC directory will contain an exception summary record for each server (shown in the example as “server2_exception.log”) and one or more records for each server with file names composed of the server name and a date/time stamp. The contents of the FFDC directory should be sent to IBM support any time you supply SystemOut.log or SystemErr.log data.

FFDC's are written for server processes only and not produced for client processes. For a WebSphere eXtreme Scale stand-alone installation, the installation properties directory must be in the classpath for FFDC's to be produced.

Summary

- Messages
- Log files
- Trace
- FFDC



In summary, WebSphere eXtreme Scale V7.0 uses specific message prefixes which indicate the component within the product that issued the message. You access the log files to review the catalog and container server messages, which can help you determine if problem are occurring during server start or during server runtime. The trace facility allows you to turn on trace, specify particular components for tracing, and allows you to specify particular trace levels to control the amount of detail in the trace log. First Failure Data Capture data is available to provide to IBM support when you report problems to IBM.

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 e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_WXS70_Logs_Traces.ppt

This module is also available in PDF format at: [../WXS70_Logs_Traces.pdf](..WXS70_Logs_Traces.pdf)



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

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both: WebSphere

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. 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>

Java, JVM, and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2009. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.

