



IBM Software Group

IBM® WebSphere® Application Server V6

Generating an IBM Heap Dump



@business on demand.

© 2005 IBM Corporation
Updated March 2, 2005

This presentation will focus on generating an IBM heap dump in WebSphere Application Server V6.

Goals

- Explain how to generate an IBM heap dump in WebSphere Application Server V6



The goal of this presentation is to explain how to generate an IBM heap dump using WebSphere Application Server V6. The behavior is slightly different than in previous versions of WebSphere Application Server.

Platform Support

- IBM heap dumps can only be generated by Java™ Virtual Machines (JVMs) that are built for WebSphere Application Server
 - ▶ IBM JVM on Linux®, AIX®, and Windows®
- Heap information for platforms using the Sun JVM can be obtained using the built-in HPROF profiler
 - ▶ The Heap Analysis Tool (HAT) from Sun can be used to analyze the heap information



IBM heap dumps can only be generated by JVMs that are built by IBM for use with WebSphere Application Server. This includes the IBM JVM running on Linux, AIX, and Windows. The information in this presentation does not apply to any other platforms.

IBM Heap Dump

- Prints a record of all objects in the Java heap to a text file
- Displays size, address, and references for each object
- Garbage collection is forced before a dump
 - ▶ This is a recent enhancement, added in JVM 1.4.1 SR1
- Compare heap dumps over time to determine if memory is leaking



Heap dump capability is a feature of the IBM JVM that prints a record of all objects in the Java heap to a text file. The size and address of each object are recorded, as well as the addresses of all objects that it references. This information can help you understand which objects are responsible for taking up large amounts of memory.

Recently, the JVM was enhanced to trigger a garbage collection before dumping the heap information. This makes the dump more accurate, because only “live” objects (that is, objects that still have references to them) will appear in the dump.

By taking multiple heap dumps and comparing object sizes over time you can determine if a particular object is not being dereferenced properly, which could result in a memory leak.

Note that Heap dumps are a feature of the IBM Java Virtual Machine. To diagnose memory problems on platforms that do not use the IBM JVM, you will need to use a Java Virtual Machine Profiling Interface (JVMPPI) profiler, such as Hprof.

Default Heap Dump Behavior

- IBM heap dump is created in <PROFILE_HOME> when an OutOfMemoryException is thrown
- Can be disabled with an environment variable
 - ▶ IBM_HEAPDUMP_OUTOFMEMORY=false



By default, WebSphere Application Server will generate a heap dump in your profile's home directory whenever an OutOfMemoryException is thrown. To disable this behavior, create an environment variable called "IBM_HEAPDUMP_OUTOFMEMORY" with a value of "false".

Generating an IBM Heap Dump Manually

- **IBM_HEAPDUMP=true**
 - ▶ Enables heap dump
 - ▶ Heap dump is in .phd (Portable Heap Dump) format
 - .phd is a more compact format used by some of the newer memory tools
- **IBM_JAVA_HEAPDUMP_TEXT=true**
 - ▶ Enables text heap dumps (classic heap dump format)
- **SIGQUIT (kill -3) generates heap dump in <PROFILE_HOME>**
 - ▶ Can also be triggered with wsadmin

```
> set jvm [$AdminControl completeObjectName type=JVM,process=servername,*]  
> $AdminControl invoke $jvm dumpThreads
```



As in previous versions of WebSphere Application Server, IBM-built JVMs can be configured to dump the contents of the Java memory heap on command by setting the “IBM_HEAPDUMP” environment variable to “true”. Unlike previous versions, the generated file will be in a more compact binary format known as a Portable Heap Dump file. With this environment variable set, a SIGQUIT sent to the JVM process will generate a heap dump file in your profile’s root directory. Heap dumps can also be generated using wsadmin, using the dumpThreads command as shown here. To enable heap dumps in the classic text format, you also need to set the variable “IBM_JAVA_HEAPDUMP_TEXT” to “true”.

Portable Heap Dump Format

- New versions of memory tools support .phd
 - ▶ IBM Rational® Application Developer V6
 - ▶ HeapRoots
 - Version 2.0.5 supports .phd
 - <http://www.alphaworks.ibm.com/tech/heaproots>
 - ▶ HeapAnalyzer
 - <http://www.alphaworks.ibm.com/tech/heapanalyzer>



Older versions of memory analysis tools do not support the new Portable Heap Dump format. The new tools that support the PHR format are listed here. HeapRoots, which also supports z/OS® svcdumps, is available for free from IBM AlphaWorks. Versions 2.0.5 and newer support the Portable Heap Dump format. HeapAnalyzer is also freely available from Alphaworks.

Summary

- Generation of an IBM heap dump is slightly changed in WebSphere Application Server V6
 - ▶ Default heap dump format is Portable Heap Dump
 - ▶ Text heap dumps can still be enabled



In summary, this presentation has focused on the slight modifications to the IBM heap dump capability in WebSphere Application Server V6. While text heap dumps can still be enabled, the default format is a new binary format called a Portable Heap Dump.

Trademarks, Copyrights, and Disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	CICS	IMS	MQSeries	Tivoli
IBM (logo)	Cloudscape	Informix	OS/390	WebSphere
e (logo)/business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

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

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

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel 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.

Linux is a registered trademark of Linus Torvalds.

Other company, product and 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 and/or changes in the product(s) and/or program(s) 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.

Information is provided "AS IS" without warranty of any kind. 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 (e.g., 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 2004. 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.