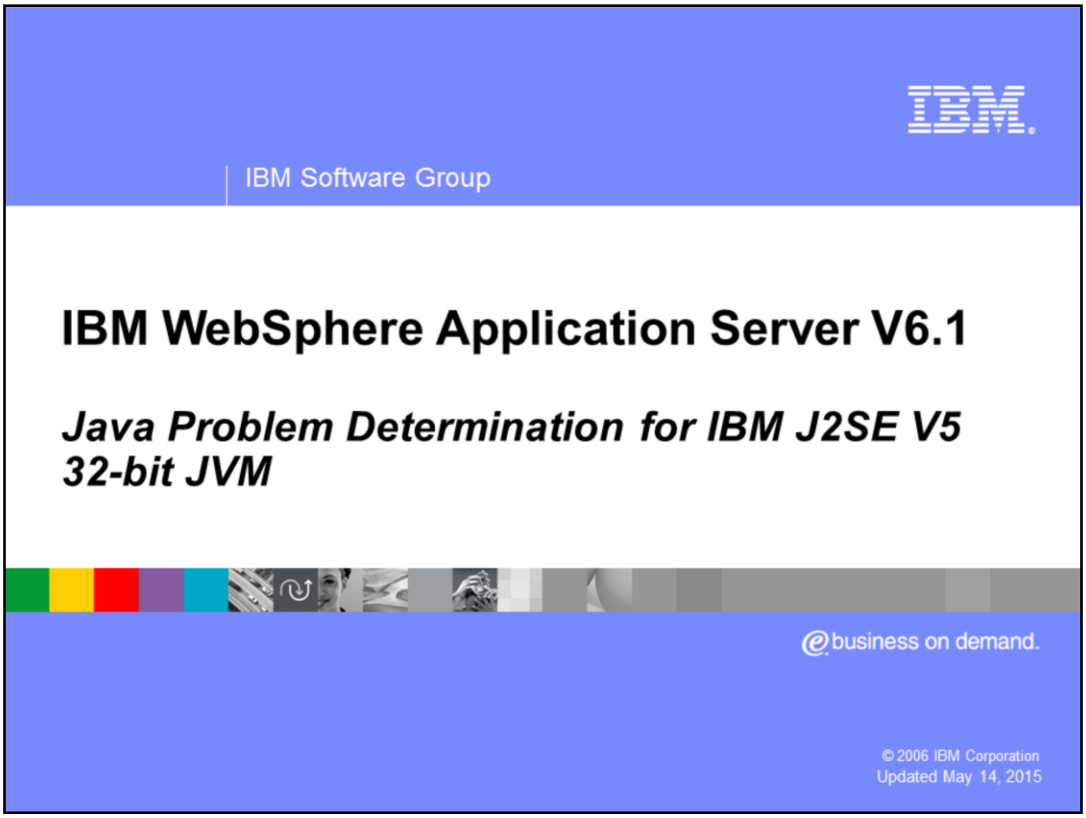

IBM Software Group

IBM WebSphere Application Server V6.1

Java Problem Determination for IBM J2SE V5 32-bit JVM

business on demand.

Updated May 14, 2015

This presentation focuses on Java™ problem determination for WebSphere® Application Server V6.1 when your application server has been enabled to use the IBM J2SE Version 5 32-bit Java Virtual Machine (or "JVM"). This JVM is also known as IBM Technology for Java. Starting with i5/OS release V5R4M0, two JVMs are now available for use with WebSphere Application Server V6.1 for i5/OS. The Classic iSeries Java Development Kit version 5 is the 64-bit JVM which is part of the i5/OS operating system and is a Sun-based JVM. The new IBM Technology for Java Version 5 JDK is a 32-bit JVM which runs in the Portable Application Solutions Enablement (or "PASE") environment.

How to enable your application server to use the IBM J2SE Version 5 32-bit JVM is discussed in the "i5/OS Java Virtual Machine Configuration" presentation.

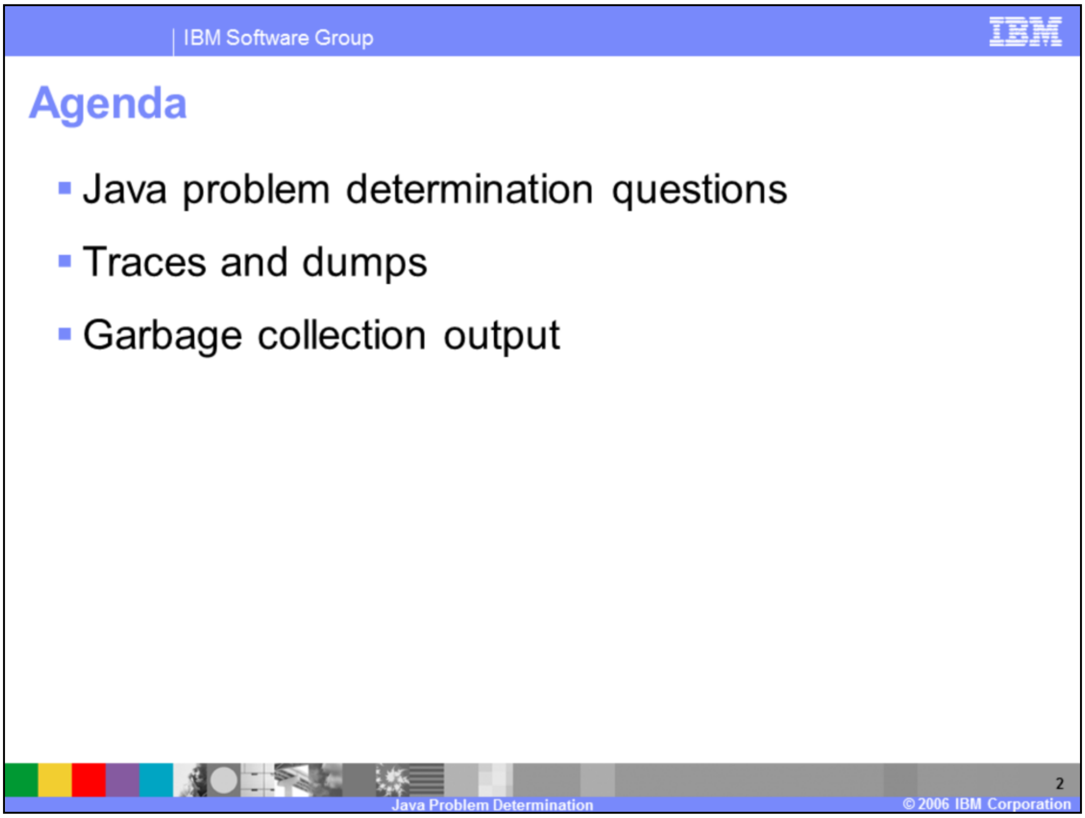

IBM Software Group

Agenda

- Java problem determination questions
- Traces and dumps
- Garbage collection output

2

Java Problem Determination

© 2006 IBM Corporation

This presentation will look at some basic Java problem determination steps, including questions to guide your information gathering, suggestions for working with dumps and traces, and information on garbage collection output.

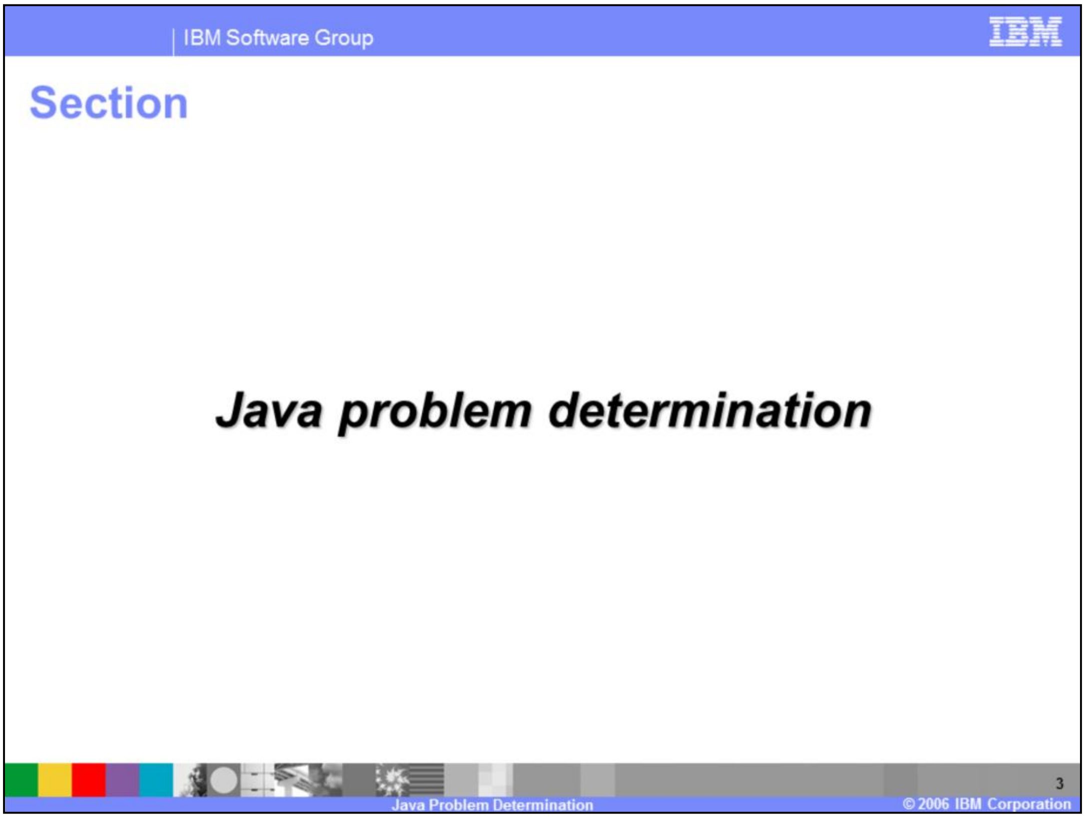

Section

Java problem determination

This section offers some suggestions for debugging Java problems.

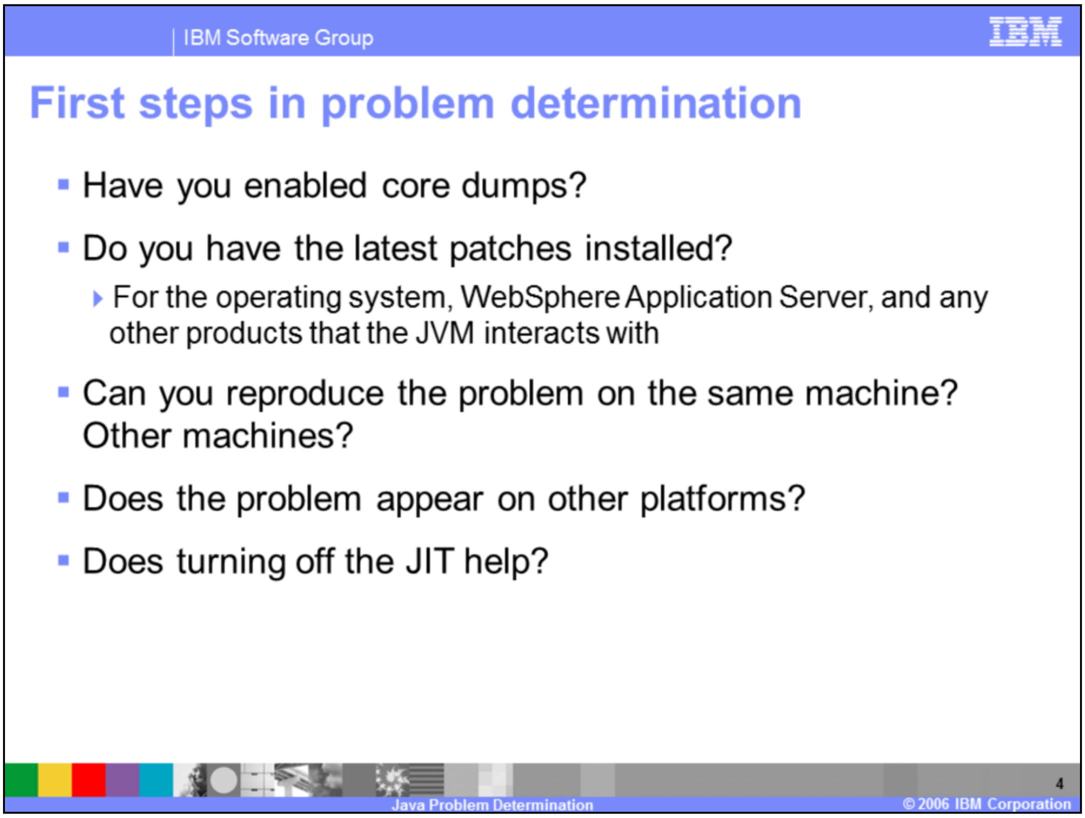

IBM Software Group

First steps in problem determination

- Have you enabled core dumps?
- Do you have the latest patches installed?
  - For the operating system, WebSphere Application Server, and any other products that the JVM interacts with
- Can you reproduce the problem on the same machine? Other machines?
- Does the problem appear on other platforms?
- Does turning off the JIT help?

4

Java Problem Determination

© 2006 IBM Corporation

There are several questions that you should ask yourself and initial troubleshooting tasks that you can perform when trying to diagnose a possible Java problem. Core dumps are large files that contain information about the Java Virtual Machine environment and the status of the JVM at the time of the dump. These dumps are enabled by default on WebSphere Application Server V6.1. The information in these dump files can be very helpful to IBM service in diagnosing problems. Verifying patch levels, knowing which JVM you are using, knowing if a problem is reproducible, and being able to provide a reproduction scenario are all helpful pieces of information to gather before contacting IBM service. For example, it is important to know if a problem can be reproduced on multiple systems and different hardware architectures – perhaps the problem only occurs on systems with multiple processors. If turning off the JIT fixes the issue, then the problem is most likely in the JIT. The JIT can be disabled from the administrative console for your application server. In general, when reporting a possible problem, try to provide as much information as you can about your environment and how to reproduce the problem that you are experiencing. It is very important for IBM Service to know which JVM you are using when reporting a Java problem. Additional problem determination tips can be found in the Java 5.0 Diagnostics Guide, found in the reference section of this presentation.

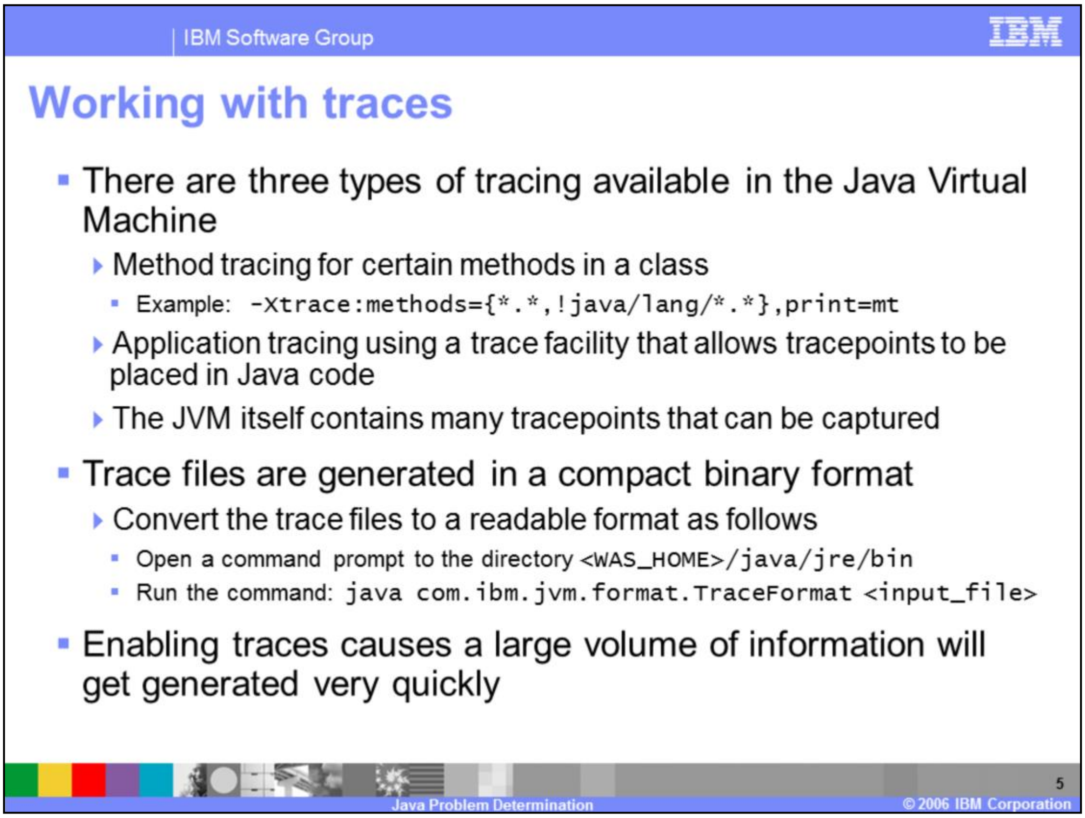

IBM Software Group

Working with traces

- There are three types of tracing available in the Java Virtual Machine
  - Method tracing for certain methods in a class
    - Example: `-Xtrace:methods={*.*,!java/lang/*.*},print=mt`
  - Application tracing using a trace facility that allows tracepoints to be placed in Java code
  - The JVM itself contains many tracepoints that can be captured
- Trace files are generated in a compact binary format
  - Convert the trace files to a readable format as follows
    - Open a command prompt to the directory `<WAS_HOME>/java/jre/bin`
    - Run the command: `java com.ibm.jvm.format.TraceFormat <input_file>`
- Enabling traces causes a large volume of information will get generated very quickly

Java Problem Determination

5

© 2006 IBM Corporation

A variety of different trace information can be gathered using Xtrace JVM options. You can trace Java methods, Java applications, and activity going on within the JVM itself. The JVM has been instrumented with several tracepoints, and you need to use the tracepoint APIs to instrument your Java applications for tracing. Java applications that contain tracepoints will not run on other Java Virtual Machines. The Xtrace example shown will trace all methods in all classes except those that start with java/lang. Notice that, when specifying what you would like to trace, you can use compound conditions and negation. Other examples of using Xtrace can be found in the Java 5.0 Diagnostics Guide.

Trace files are generated in a compressed format and need to be converted to a readable format using a special tool. Enabling traces will cause a very large amount of data to be generated, so be careful to only trace what is needed.

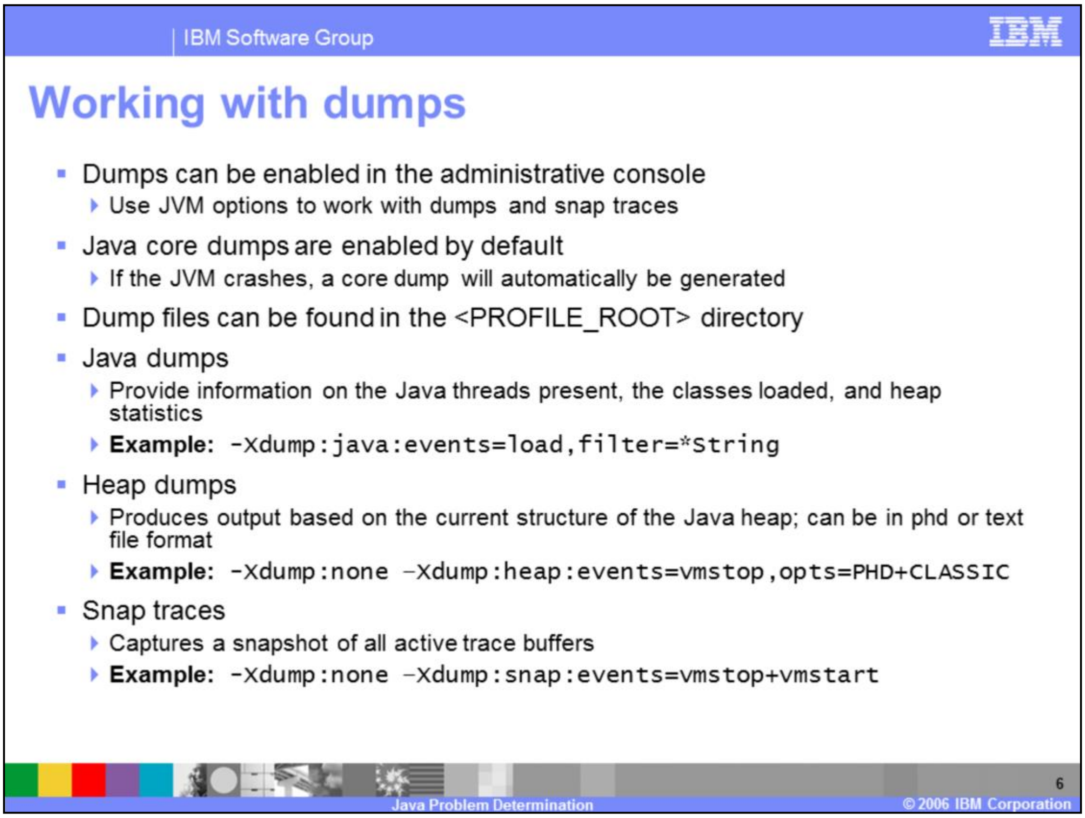

IBM Software Group

IBM

Working with dumps

- Dumps can be enabled in the administrative console
  - Use JVM options to work with dumps and snap traces
- Java core dumps are enabled by default
  - If the JVM crashes, a core dump will automatically be generated
- Dump files can be found in the <PROFILE_ROOT> directory
- Java dumps
  - Provide information on the Java threads present, the classes loaded, and heap statistics
  - **Example:** `-Xdump:java:events=load,filter=*String`
- Heap dumps
  - Produces output based on the current structure of the Java heap; can be in phd or text file format
  - **Example:** `-Xdump:none –Xdump:heap:events=vmstop,opts=PHD+CLASSIC`
- Snap traces
  - Captures a snapshot of all active trace buffers
  - **Example:** `-Xdump:none –Xdump:snap:events=vmstop+vmstart`

6

Java Problem Determination

© 2006 IBM Corporation

Java core dumps are enabled in WebSphere Application Server V6.1 by default, and a core dump will be produced whenever a JVM crashes. Specific dump behavior can be configured in the administrative console, including the ability to capture dumps based on certain criteria. The dump files are stored in the corresponding profile’s root directory.

Different types of dumps can be captured, based on the options that you specify. The first Java dump example will force a core dump to occur whenever any class is loaded whose name ends with String. The second example will force a heap dump when the JVM is stopping. Two dumps will be created: one in phd format and one in classic, or text file, format. The final example will capture a snap trace of all of the active trace buffers when the JVM is starting and stopping.

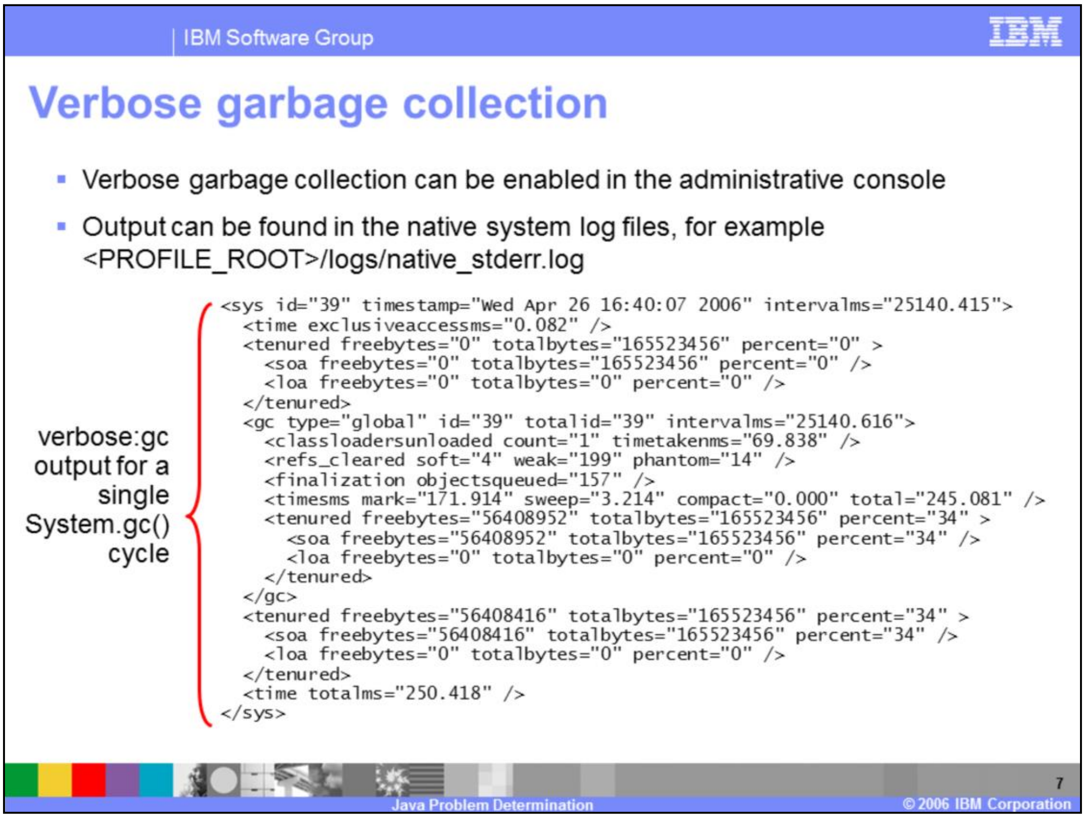

IBM Software Group

Verbose garbage collection

- Verbose garbage collection can be enabled in the administrative console
- Output can be found in the native system log files, for example <PROFILE_ROOT>/logs/native_stderr.log

verbose:gc output for a single System.gc() cycle

```
<sys id="39" timestamp="Wed Apr 26 16:40:07 2006" intervalms="25140.415">
  <time exclusiveaccessms="0.082" />
  <tenured freebytes="0" totalbytes="165523456" percent="0" >
    <soa freebytes="0" totalbytes="165523456" percent="0" />
    <loa freebytes="0" totalbytes="0" percent="0" />
  </tenured>
  <gc type="global" id="39" totalid="39" intervalms="25140.616">
    <classloadersunloaded count="1" timetakenms="69.838" />
    <refs_cleared soft="4" weak="199" phantom="14" />
    <finalization objectsqueued="157" />
    <timesms mark="171.914" sweep="3.214" compact="0.000" total="245.081" />
    <tenured freebytes="56408952" totalbytes="165523456" percent="34" >
      <soa freebytes="56408952" totalbytes="165523456" percent="34" />
      <loa freebytes="0" totalbytes="0" percent="0" />
    </tenured>
  </gc>
  <tenured freebytes="56408416" totalbytes="165523456" percent="34" >
    <soa freebytes="56408416" totalbytes="165523456" percent="34" />
    <loa freebytes="0" totalbytes="0" percent="0" />
  </tenured>
  <time totalms="250.418" />
</sys>
```

7

Java Problem Determination

© 2006 IBM Corporation

Verbose garbage collection information allows you to see how your heap space is being used and how much time garbage collection cycles take. Garbage collection information can be very useful when trying to diagnose performance problems. Verbose garbage collection produces output in XML format that gets directed to the native system log files; it can be enabled from the administrative console.

The sys tag indicates that a System.gc() call has occurred. The information between the tenured tags describes the occupancy levels of the different sections of the heap – the large object area and the small object area. The time tag shows the total amount of time, in milliseconds, that it took to handle the System.gc() call.

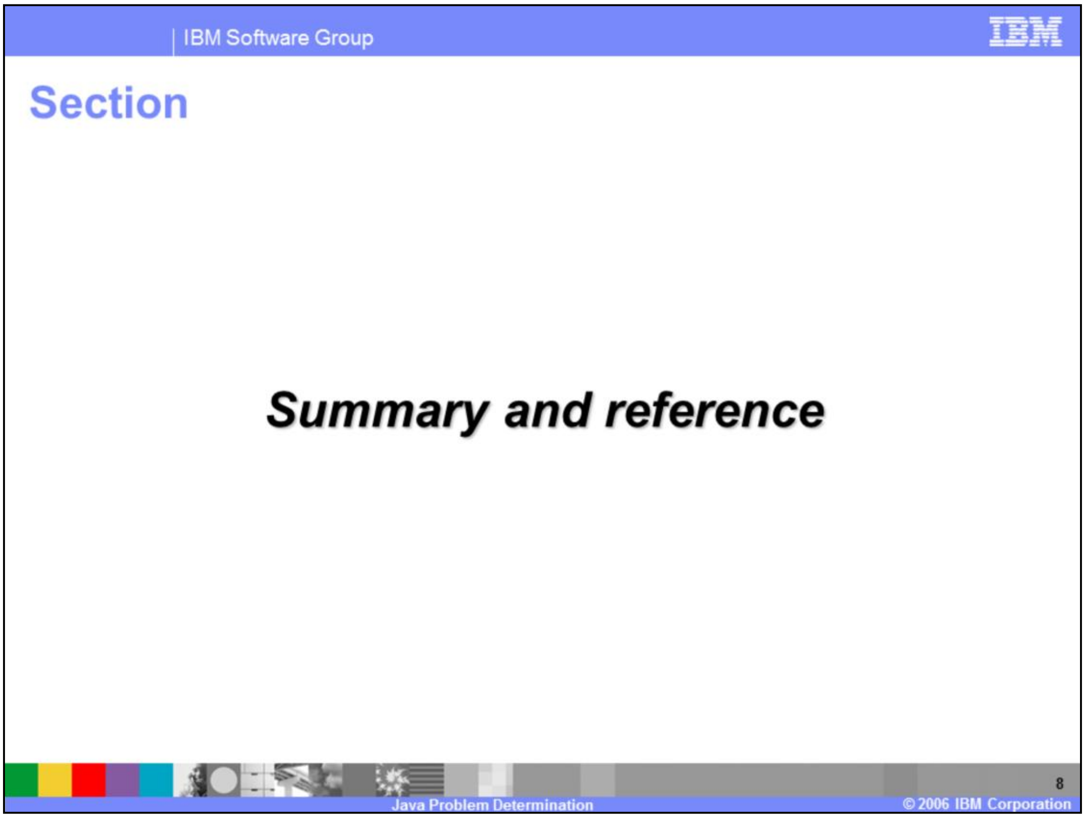

IBM Software Group

Section

Summary and reference

8

Java Problem Determination

© 2006 IBM Corporation

The last portion of the presentation contains a summary and references.

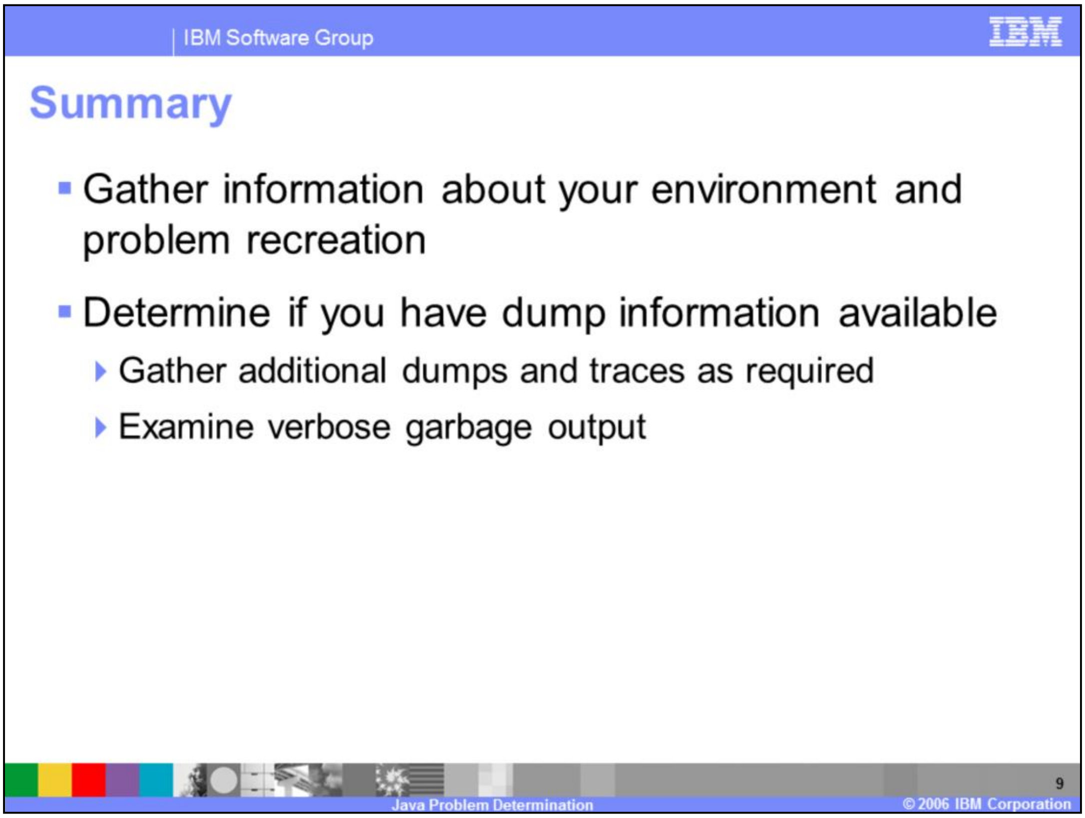

Summary

- Gather information about your environment and problem recreation
- Determine if you have dump information available
  - Gather additional dumps and traces as required
  - Examine verbose garbage output

There are a variety of techniques that you can use to help you debug possible Java problems in WebSphere Application Server V6.1 when you are using the IBM J2SE Version 32-bit JVM. First, you should determine as much information about your environment as possible. If you are able to recreate the problem, providing a detailed reproduction scenario will be very helpful to IBM service. A variety of dump and trace information – including core dumps, heap dumps, traces, and verbose garbage collection output – can be gathered to help diagnose problems.

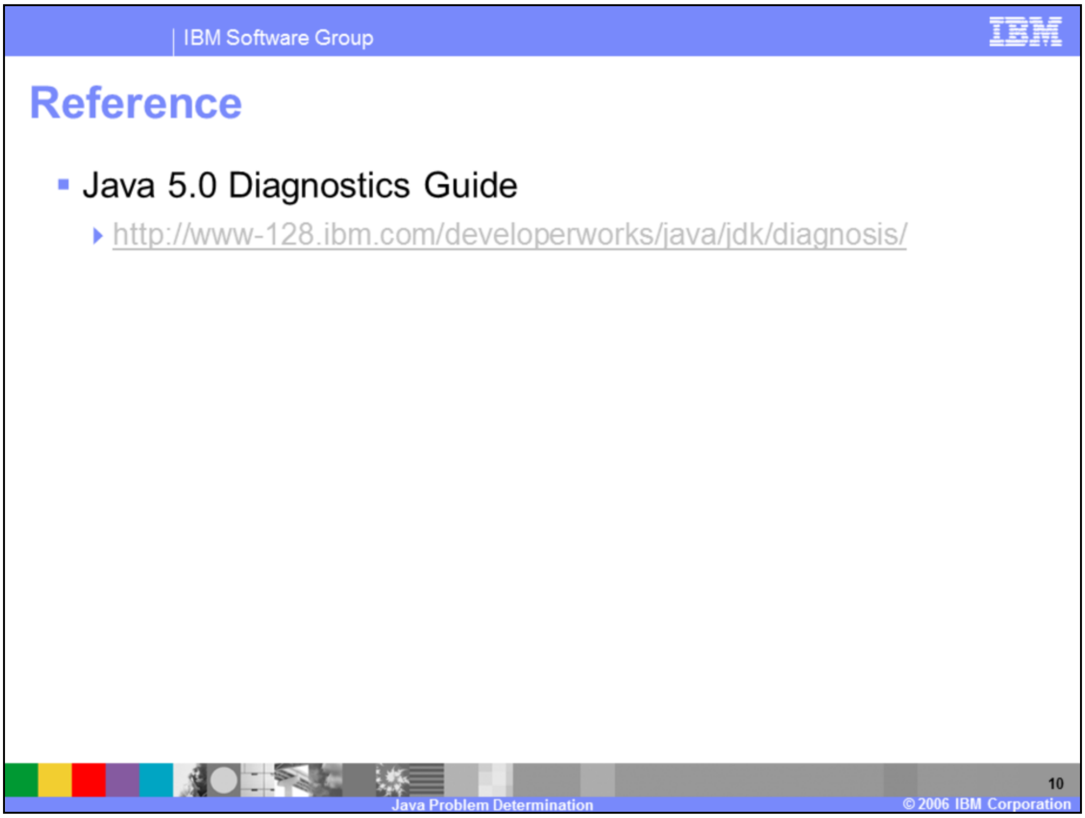

IBM Software Group

Reference

- Java 5.0 Diagnostics Guide
  - http://www-128.ibm.com/developerworks/java/jdk/diagnosis/

10

Java Problem Determination

The Java 5.0 Diagnostics Guide contains detailed information about the structure of the IBM Virtual Machine for Java, steps for troubleshooting problems, and how to go about reporting problems to IBM service.

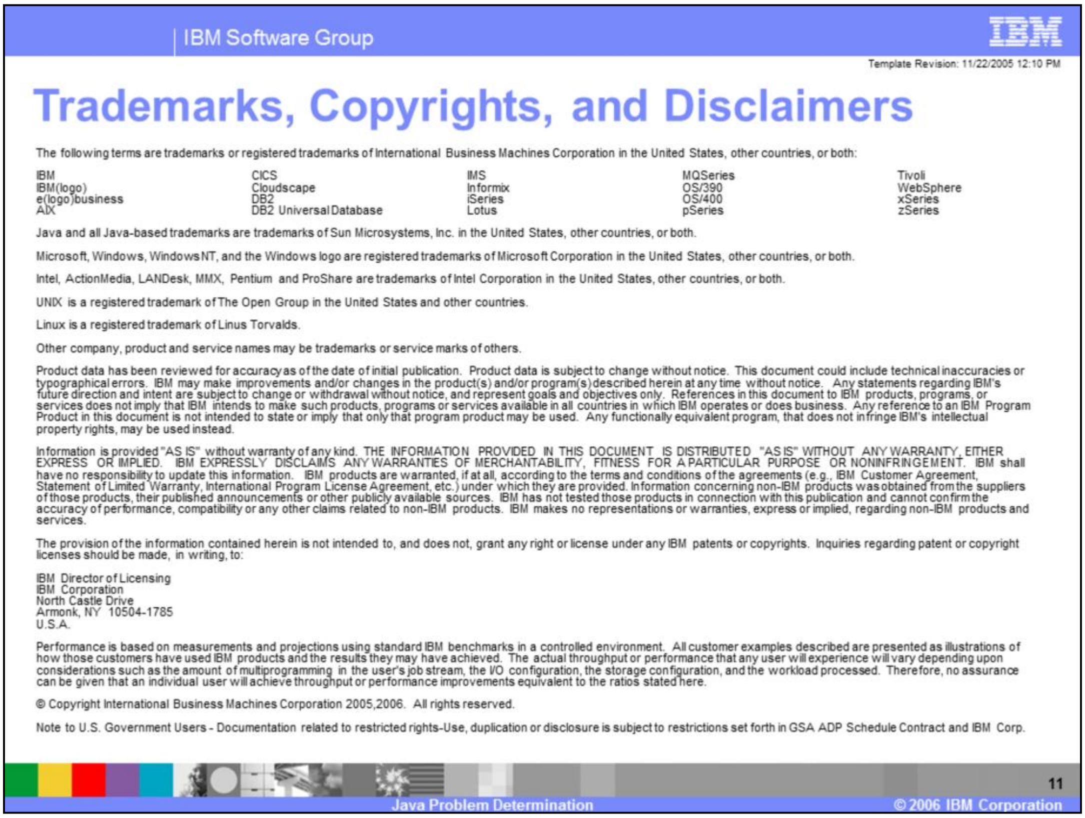

IBM Software Group

Template Revision: 11/22/2005 12:10 PM

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

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.

11

Java Problem Determination

© 2006 IBM Corporation