



IBM Java Technology Center

IBM Java Serviceability Tools – A Practical View

Or: “How to get the tools to do all the heavy lifting”

Trent Gray-Donald IBM

Will Smith IBM

WUG UK | Sept 17, 2008

© 2008 IBM Corporation

Agenda

- History
- Problem Determination recap
- Strategic direction
- Tools dive
- Discussion / Questions

The IBM Java Trend...

- Customer adoption of Java continues to new heights...
 - Java is by many measures the dominant language in the enterprise.
- IBM is a huge consumer too!
 - All major brands in IBM have major releases built on Java
 - WebSphere Application Server
 - Lotus Notes
 - Rational Application Developer
 - (over 300 major IBM products, too many to list!)
- Industry Trends - complexity of Java environments is going up
 - Software: frameworks & middleware layers
 - Hardware: memory usage, multi-core CPUs
 - New paradigms: realtime and low latency environments
 - Production-time failure analysis is becoming the norm

So what?

Simplify.
Focus.
Execute.

History

- **Fragmented tooling story**
 - Different tools for different folks, find tools in various places
 - In the JDK itself
 - alphaWorks
 - developerWorks
 - From Java & WAS support
 - Tools JVM level specific
- **Substantial technology changes in underlying JVM implementation between 1.4.2 and 5.0**
 - Significant robustness improvements (better compaction / fragmentation support, enhanced FFDC)
 - Fundamental PD data produced in same format

Simplify.

IBM Support Assistant

- Hosting for Serviceability Tools across product families
- Automatic PD data gathering
- Assist with opening PMRs and working with IBM Support
- Documentation
 - Aggregated Search across sources
 - Regular updates to Diagnostics Guide

Welcome - IBM Support Assistant Workbench

File Administration Update Window Help

Support Assistant

Launch Activity Home

Welcome

Welcome to IBM Support Assistant

Find Information
Easily find the information you need including product specific information and search capabilities.

Analyze Problem
Diagnose and analyze problems through serviceability tools, collection of diagnostic artifacts, and guidance through problem determination.

Manage Service Request
Effectively submit, view and manage your service requests enhanced with automated collection of diagnostic data.

First Steps
Make first steps

Tutorials
Go through tutorials

Latest News
No News currently available

ISA - Search

Search Information - IBM Support Assistant Workbench

File Administration Update Window Help

Support Assistant

Launch Activity Home Find Information

Search Information Media Viewer Product Information

IBM Support Assistant Search

Search: Go! [Advanced Search](#)

Search Targets

- IBM Software Support Documents
- IBM developerWorks
- IBM Newsgroups and Forums
- Google

Product Information Centers

- IBM Developer Kit for Java 1.4
- IBM Developer Kit for Java 5.0
- IBM Developer Kit for Java 6.0

Welcome to IBM Support Assistant Search!

Concurrently search across multiple repositories within the Search Information activity.

The IBM Support Assistant Search Information activity may be leveraged to search across multiple repositories with a single search request. Start by ensuring that a search target is enabled. Search targets are enabled by checking the box next to the search target names. To configure advanced search settings for a search target click on the search target name or "Advanced Search".

After enabling and configuring the desired search targets enter a query and press Go!

Collect Data - IBM Support Assistant Workbench

File Administration Update Window Help

Support Assistant

Launch Activity Home Analyze Problem

Collect Data Guided Troubleshooter Tools

Select a case and incident: default Select Create Portable Collector

Select Collectors

Select a Collector

Collected system: My Computer

Select a product (or System Collector) and its problem:

- System Collector
 - Collect all information
 - Collect network information
 - Collect registry and installed software information
 - Collect system data information
 - Collect user environment information
 - Collect software inventory

Add>>

Collector Queue

System	Product	Problem

Collect All Cancel Remove

Collector Status

View Details

Tools - IBM Support Assistant Workbench

File Administration Update Window Help

Support Assistant

Launch Activity Home Analyze Problem

Collect Data Guided Troubleshooter Tools

Case/Incident

default [Select]

Tools Catalog

- IBM Monitoring and Diagnostic Tools for Java - Dump Analyzer v2.1.0.20080619115044
- IBM Monitoring and Diagnostic Tools for Java™ - Garbage Collection and Memory Visualizer v2.1.0.20080815**
- IBM Pattern Modeling and Analysis Tool for Java Garbage Collector (Tech Preview) v2.4.0.01
- Log Analyzer v4.4.1.200808271600
- Memory Dump Diagnostic for Java (MDD4J) Beta v2.0.0.beta-20080721212546
- Symptom Editor v4.4.1.200808081528
- ThreadAnalyzer (Tech Preview) v6.0.3.00
- Visual Configuration Explorer (Tech Preview) v1.0.13.200808150959

Restrictions

None

Description

The IBM Monitoring and Diagnostic Tools for Java™ - Garbage Collection and Memory Visualizer is a verbose GC data visualizer. The GC And Memory Visualizer parses and plots all IBM verbose GC logs and -Xtgc output (and is extensible to parse and plot other forms of input). It provides graphical

Associations

Tool is not associated with any products

[Launch] [Help]

Case Manager

Filter: []

Private	Name
<input type="checkbox"/>	Results-A
<input type="checkbox"/>	mdd4j.log

[Update] [Undo]

History:

Date	Message
3/13/08 12:55 PM	Case Manager: ..
3/13/08 12:55 PM	Case Manager: ..

Focus.

Common Support Concerns

- **OutOfMemoryError / Heap Size Tuning**
 - It's hard to tune the right GC parameters, and figure out where memory leaks come from.
- **Deadlocks / hangs / spins**
 - Need ability to introspect on a running JVM to determine what's happening at the moment – in a report based way.
- **General analysis tools**
 - Need ability to examine JVM data – classloaders, threads, monitors, etc.. to do general PD tasks.

2008 Themes for Strategy

- **Unify**
 - Bring currently separate tools together
 - Common view of all facets of Java execution, live or after the fact
- **Tools must be usable everywhere**
 - GUI mode for interactive use
 - Report generation for headless environments
- **Recommend**
 - Tools mature from visualization of data to recommendations on potential issues, recommending future actions
 - Point from tools to documentation and vice versa.
- **Refine**
 - Feedback key to continual improvement, both from internal (L2, L3, other developers) and external customers
 - Agile development model central

Java Problem Determination Strategy - 2008

- **Improved Documentation and Education**
 - Revamp and streamline documentation
 - Unlock existing expert knowledge via recorded presentations
 - Expand IBM Guided Activity Assistant coverage for base Java issues.
- **Improved FFDC and ‘must gather’ streamlining**
 - Capture the right data first time, reducing recreates
 - Streamline process of sending information to IBM
- **Runtime performance analysis**
 - Facilitate determining the cause of application bottlenecks
 - Very low overhead monitoring
- **Post Mortem tools**
 - Continue good work begun in 2007 on ‘IBM Dump Analyzer for Java’ and ‘Garbage Collector and Memory Visualizer’
 - Tools being deployed very regularly – looking for customer feedback.

Execute.

Java Tools

- If you have a problem with your Java-based application there are tools to help you get in control of it
- The tools can help diagnose problems in Java applications as well as defects in the Java runtime
- IBM Support Assistant can boost the productivity of time you spend on support calls
- It's easy to get familiar with the tools, the starting point is:
 - <http://www.ibm.com/developerworks/java/jdk/tools/index.html>
- So what do the tools work on?...

Java Diagnostics – Java Dump, Heap Dump

- **Java Dump**
 - Human readable, text file
 - Small and quick to write
 - Default trigger events: GPF, quit or abort signal, OutOfMemoryError
 - Contains version, java command line, heap info, lock info, thread stacks, loaded class list
 - User can control location, with a chain of fallback locations
- **Heap Dump**
 - A summary of the objects on the Java heap
 - Quicker and smaller than writing the whole heap to disk
 - Default trigger event: OutOfMemoryError

Java Diagnostics – System Dump, Snap Dump

- System Dump
 - Operating System dump
 - Contains full address space image
 - Large and time consuming to write
 - Default trigger events: GPF, abort signal
 - IBM Java support need the output of the jextract post processor in preference to the original dump
- Snap Dump
 - JVM trace records from a memory buffer
 - Default trigger events: GPF, abort signal, OutOfMemoryError
 - Formatted into readable text by `com.ibm.jvm.format.TraceFormat` postprocessor

Java Diagnostics – others

- Verbose Garbage Collection log
 - Summary of JVM Garbage Collection events with heap statistics at each GC
 - Performance overhead is low, but this is not on by default
 - Output can go to stderr or a file
- Platform diagnostics, process monitoring, netstat
 - Typically requested by a support engineer to diagnose a specific problem
- Product logs
- Application level logging is important too
- All these logs can take a lot of work to collect repeatedly ...

Diagnostics Collector – what it is

- IBM Monitoring and Diagnostic Tools for Java - Diagnostics Collector
- Runs as a separate process when the JVM detects a 'dump event'
 - GPF
 - Java heap OutOfMemoryError
 - Unexpected signal received
 - (optionally) JVM start, JVM stop
- Dumps can end up in multiple locations
- Diagnostics Collector knows about all the possible dump locations and searches them to gather all dumps into a single zip file
- If a System Dump is found, Diagnostics Collector runs jextract to remove a previously manual step
- What it isn't: an ISA Data Collector – it runs immediately when a problem occurs
- Working with ISA to integrate functionality into the ISA model

Diagnostics Collector – how to get it

- Early in development - currently users download the tool and configure it for their Java installation
- Unzip download package, edit a config file, add a `-Xoptionsfile` parameter to the Java command line or append to an existing options file
- Linked from the Java tools page
 - <http://www.ibm.com/developerworks/java/jdk/tools/index.html>
- Direct link to download page
 - <http://www-01.ibm.com/support/docview.wss?rs=3068&context=SSNVBF&dc=D400&uid=swg24019419>

Demo: Diagnostics Collector

- sample program that exhausts the Java heap (with system dump for good measure)

Troubleshooting with IBM Guided Activity Assistant (IGAA)

- Product Add-ons supply Troubleshooting content for their product
- IGAA Framework is part of the base IBM Software Assistant (ISA) install
- IGAA guides you through troubleshooting steps
- Each step gives instructions for gathering data and interpreting it to move to the next step
- You can pause at any stage and work more later
- IGAA uses the Case Manager to organise the diagnostics files for a problem

Demo: Guided Troubleshooter

- what does ISA Guided Troubleshooter do for our OutOfMemoryError?

GCMV

- IBM Monitoring and Diagnostic Tools for Java - Garbage Collection and Memory Visualizer
- Tool to analyze IBM Java verbose GC log
- Graphs show Garbage Collection and Java heap statistics over time
- It's not just for Out Of Memory Errors, good for performance tuning in the right hands
- Tuning recommendations use heuristics to guide you towards issues that may be limiting the performance of the application

Demo: GCMV

- Analyze the verbose GC log from the test program using GCMV

Demo: MAT

- Analyse the Java Heap from the test program using MAT

Dump Analyzer

- IBM Monitoring and Diagnostic Tools for Java - Dump Analyzer
- Analyzes a System dump that has been post processed by jextract
- Heuristics used to detect problem conditions captured in the dump
- Various analysis modules are provided, for different views on the Java process in the dump

Demo: Dump Analyzer

- Analyze the System Dump from the OutOfMemoryError test program using Dump Analyzer
- Demonstrate deadlock detection using Dump Analyzer

Strategic Enabler – Diagnostic Tool Framework for Java

- Technology within the IBM JDK to facilitate analysis and diagnosis of problems in Java applications
 - Read RAS artifacts from a JVM (e.g. a core file) and extract all kinds of useful information from that dump
 - Not just one tool: an extensible framework for building many different tools
- Accelerates implementation and robustness of Java problem determination tools
 - Provides definitive parsers for diagnostics files
 - Relieves tools developers from reinventing low value components and lets them focus on the important part: data analysis and visualization.
- Now being discussed in public as JSR 326

Simplify

- IBM Support Assistant can make the time you spend on support cases more productive
- The Java troubleshooting support in ISA walks you through data gathering and using tools to diagnose a problem
- Diagnostics Collector automates file collection so that a system operator needs fewer files to document a problem
- The Eclipse framework automatically detects when a newer version of an installed tool is available for upgrade

Focus

- Guided troubleshooter gives directions towards useful steps to diagnose a problem
- GCMV helps to visualize the performance of your Java application over time
- Dump Analyzer helps to quickly understand the issues in a System dump for a Java process

Execute

- A good starting point is:
 - <http://www.ibm.com/developerworks/java/jdk/tools/index.html>
- Feedback is very welcome – the development team know that we need to reach out and work with customers
- Feature requests and problem reports from tool users gets developers' attention

Email Contact

Will Smith, software engineer, Java Consumability project

– will.Smith@uk.ibm.com

■ Java Tools support email id

– javatool@uk.ibm.com

Simplify.
Focus.
Execute.



IBM Java Technology Center

Questions & Answers

© IBM Corporation 2008. All Rights Reserved.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. 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. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. 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 the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including 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 results similar to those stated here.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM trademarks, see www.ibm.com/legal/copytrade.shtml
AIX, CICS, CICSplex, DB2, DB2 Universal Database, i5/OS, IBM, the IBM logo, IMS, iSeries, Lotus, OMEGAMON, OS/390, Parallel Sysplex, pureXML, Rational, RCAF, Redbooks, Sametime, Smart SOA, System i, System i5, System z , Tivoli, WebSphere, and z/OS.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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