



JAVAPOLIS

11 - 15 DECEMBER ■ ANTWERP ■ BELGIUM






Why So Slow?

Debunking Speculative Tuning

Heinz Kabutz
Java Specialist
Maximum Solutions
<http://www.cretesoft.com>

Kirk Pepperdine
Performance Specialist
Kodewerk Ltd.
<http://www.kodewerk.com>




Our Typical Customer


- ☞ Customer JoGoSlo Ltd calls us in desperation
 - ☞ Millions of \$\$\$\$ invested
 - ☞ Users complain about poor performance
 - Customers are starting to abandon the project
- ☞ Developers in a panic
 - ☞ 6 man months already spent “tuning” with no results
 - ☞ Can almost reproduce the problem
 - ☞ Still have some ideas of what to do
 - ☞ But, management has lost confidence
- ☞ We have 5 days to diagnose problem and propose fix



Tuning Tool for Managers



Tuning Tool for Engineers – “The Box”



Heinz Kabutz

- Author of *The Java Specialists' Newsletter*
- Sun Java Champion
- <http://www.cretesoft.com>
- Lives in Greece
- Consults and trains companies about Java



Kirk Pepperdine



- Engaged in performance tuning world wide.
- Co-author of www.javaperformancetuning.com
- Editor www.theserverside.com
- Sun Java Champion
- Speaks frequently about performance tuning
- <http://www.kodewerk.com>




Topics

- Dynamic nature of systems
- Measure don't guess
- People
- Hardware/OS
- JVM
- Application
- External systems
- Putting it all together





Time to Setup



Time to Setup TipsDB


- Download from <http://www.createsoft.com/outgoing/javapolis.zip>
- Set path to your JDK in the setenv.bat
- Go into tipsdb directory
- Call startDB.bat
- Call createDB.bat
- Call appserverStart.bat
- Connect to <http://localhost:8080/tips/wildcard>
- Connect to <http://localhost:8080/tips/keyword>



Dynamic Nature of Systems

Knowing what to measure and how to measure it makes a complex world much less so

Steven D. Levitt
Stephen J. Dubner
Authors of Freakonomics



Dynamic Nature of Systems

- 🔄 Performance tuning is a complex task
 - ➔ Need to reverse engineer complex systems
 - ➔ Need right view of the system
 - Most useful view comes from measurements
- 🔄 We will take introductory look at
 - ➔ What to measure
 - ➔ How to measure
 - ➔ How to understand the measurements



Importance of the Environment

- 🔄 Need to understand all elements in the environment
- 🔄 Changing elements of a system can change the dynamics of that system
 - ➔ E.g. different users, CPUs, network



Importance of Tooling

- 🔗 Tooling allows us to see what is otherwise invisible




Importance of process

- 🔄 Process or ways of investigating the problem can change or hide the problem
- 🔄 Systematic investigation

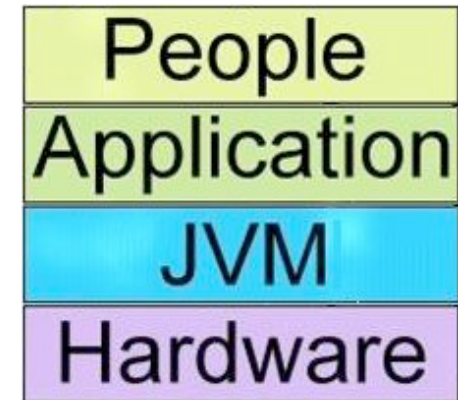


Holistic View



Dynamic Nature of Systems

- Systems by their nature are dynamic
 - Mix of static and dynamic elements
- Static aspects of a Java based system
 - Not bottlenecks onto themselves
 - Hardware/OS
 - Defines the physical constraints of the system
 - Java Virtual Machine
 - Primarily a translation layer
 - Application
 - Expression of what is needed to be done




Dynamic Nature of Systems

🌀 Dynamic aspects of a system


➔ People

- Abstraction for system drivers
 - Batch processing
 - External systems
- create flows through the system
 - maybe beyond the capacity of the system
 - Can put pressure on pinch points (or bottlenecks) in the system

🌀 How does this work?



Resource Contribution



Forward Propagation of Actions

- 👤 People drive the application
- 👤 Application drives the JVM
 - ➔ Direct consequence of what the people are asking
 - ➔ And how application was coded
- 👤 JVM Drives the hardware
 - ➔ Direct consequence of what the application is asking
 - ➔ And how JVM was coded and configured
- 👤 Hardware executes instructions
 - ➔ Limited by speed and capacity



Backward Propagation of Problems

- 🔄 Problem: hardware lacks capacity or is slow
 - ➔ people experience poor response times
- 🔄 Problem: JVM is poorly configured
 - ➔ People experience poor response times
- 🔄 Problem: Application suffers from contention
 - ➔ People experience poor response time
- 🔄 Our starting point; people are experiencing poor response times
- 🔄 How do we start our investigation?
 - ➔ It is at this point JoGoSlo ran into trouble




Performance Anti-pattern: Shot in the Dark

- 🔗 Developers dove into the code
 - ➔ Found many ugly bits
 - Interactions with database
 - ➔ Wasted valuable time fixing them
 - None of the ugly code bits had any consequence on performance
 - ➔ Ignored key pieces of information
 - DBA reported millisecond response times
 - System sometimes recovered
- 🔗 Developers started guessing at the cause of the problem



Solution to Shot in the Dark

Measure
Don't Guess



Measure Don't Guess

- 🔄 Solid Measurements
 - ➔ Show you what needs to be done
 - ➔ Focus efforts
 - ➔ Facilitate planning
 - ➔ Instill confidence
 - ➔ Deflect finger pointing



Measure Don't Guess

- 🕒 Review all performance requirements
- 🕒 Construct a realistic test environment
- 🕒 Use “The Box” as a roadmap
- 🕒 Tackle one layer at a time
- 🕒 Start with the people
- 🕒 Start the investigation with the hardware
 - ➔ Work up the stack
- 🕒 Let the user experience guide all decisions



Investigative W5

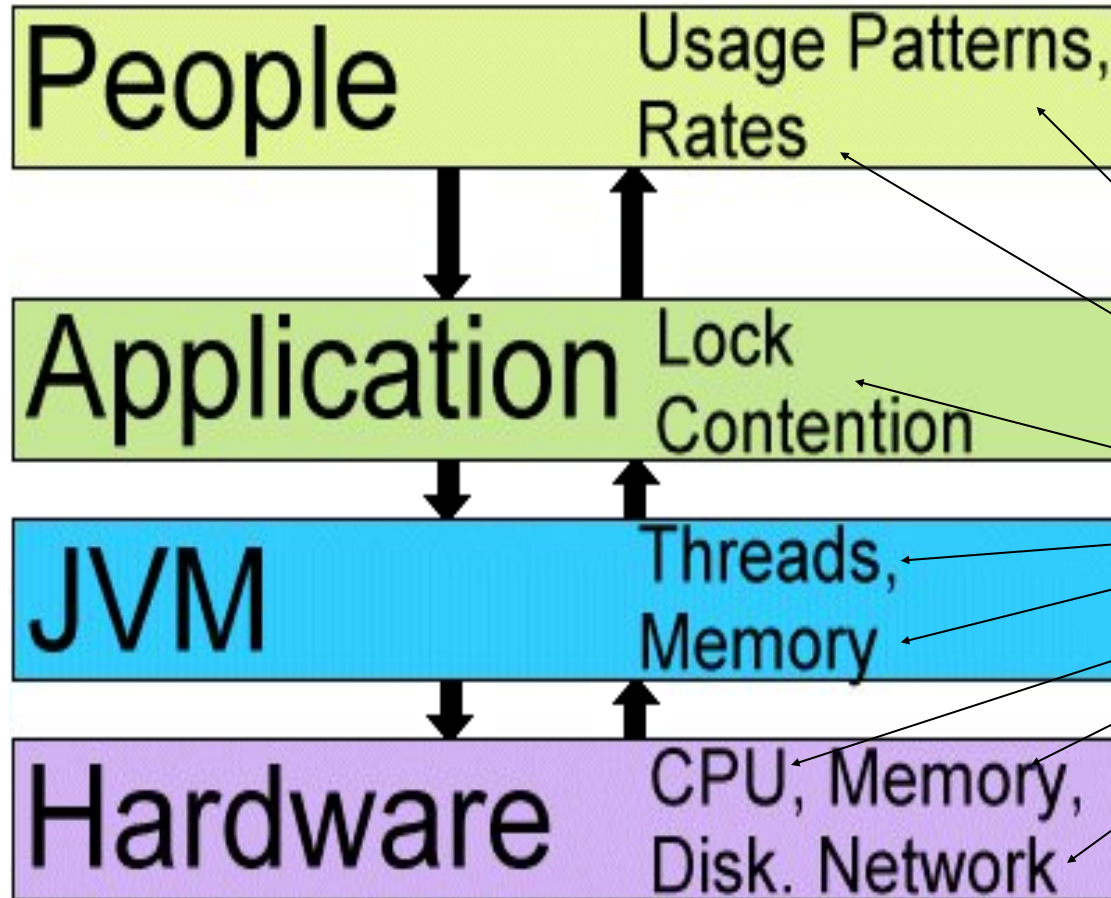
- Five questions asked by investigators:
 - Who ?
 - Who (which resource) is exhibiting the problem?
 - What ?
 - Observation: what do the users see?
 - Where ?
 - Which layer is exhibiting the problem?
 - When ?
 - Are there any peculiarities about when the problems occur?
 - Why ?
 - An explanation (hypothesis) of the observation from system perspective



Actors in the Performance Profile

What

Where



Simple Process

- ➊ Form a hypothesis from observed behavior
- ➋ Devise a test to validate the hypothesis
- ➌ Measure for effect
- ➍ Make changes
- ➎ Test for desired effect
- ➏ Repeat until performance profile is in tolerance



People

Usage Patterns, Rates

- Provide the dynamics for the system
 - Use system in their own way
 - Use the system at their own leisure
- Need to capture the dynamics
- Usage pattern
 - Sequence of user actions
 - Timing information
 - Pauses between actions
 - Time of day for activity



People

- 🔗 System utilization is an aggregate of all usage patterns
 - ➔ How system copes with the aggregation defines its performance profile
- 🔗 Stress testing
 - ➔ Use mix of usage patterns to load the system
 - Ideally driven by a load testing tool
 - ➔ Measure system activity
 - Careful use of a selected tools
 - ➔ Must be run against a production like environment
- 🔗 Goal: understand the user experience



Stress Testing Environment

- 🌐 Production environment?
 - ➔ Not desirable and usually not an option
- 🌐 Test environment should
 - ➔ Perfectly resemble your production environment
 - Data sizes, memory sizes, cache sizes, disk speeds, network speeds, should be the same
 - ➔ Be isolated
 - Introduce other systems/processes in a controlled fashion



Stress Testing Environment

☞ Caching

- ☞ Protects your application from an underlying slower technology
- ☞ Reduces response times
- ☞ May reduce the effects of I/O waits

☞ May need to simulate external systems

- ☞ Do this with care

☞ Don't extrapolate!

- ☞ Difficult to know when you will hit the wall
- ☞ E.g. Application using 15Mbits is moved from a gigabit to 10Megabit network
 - Shifts the bottleneck

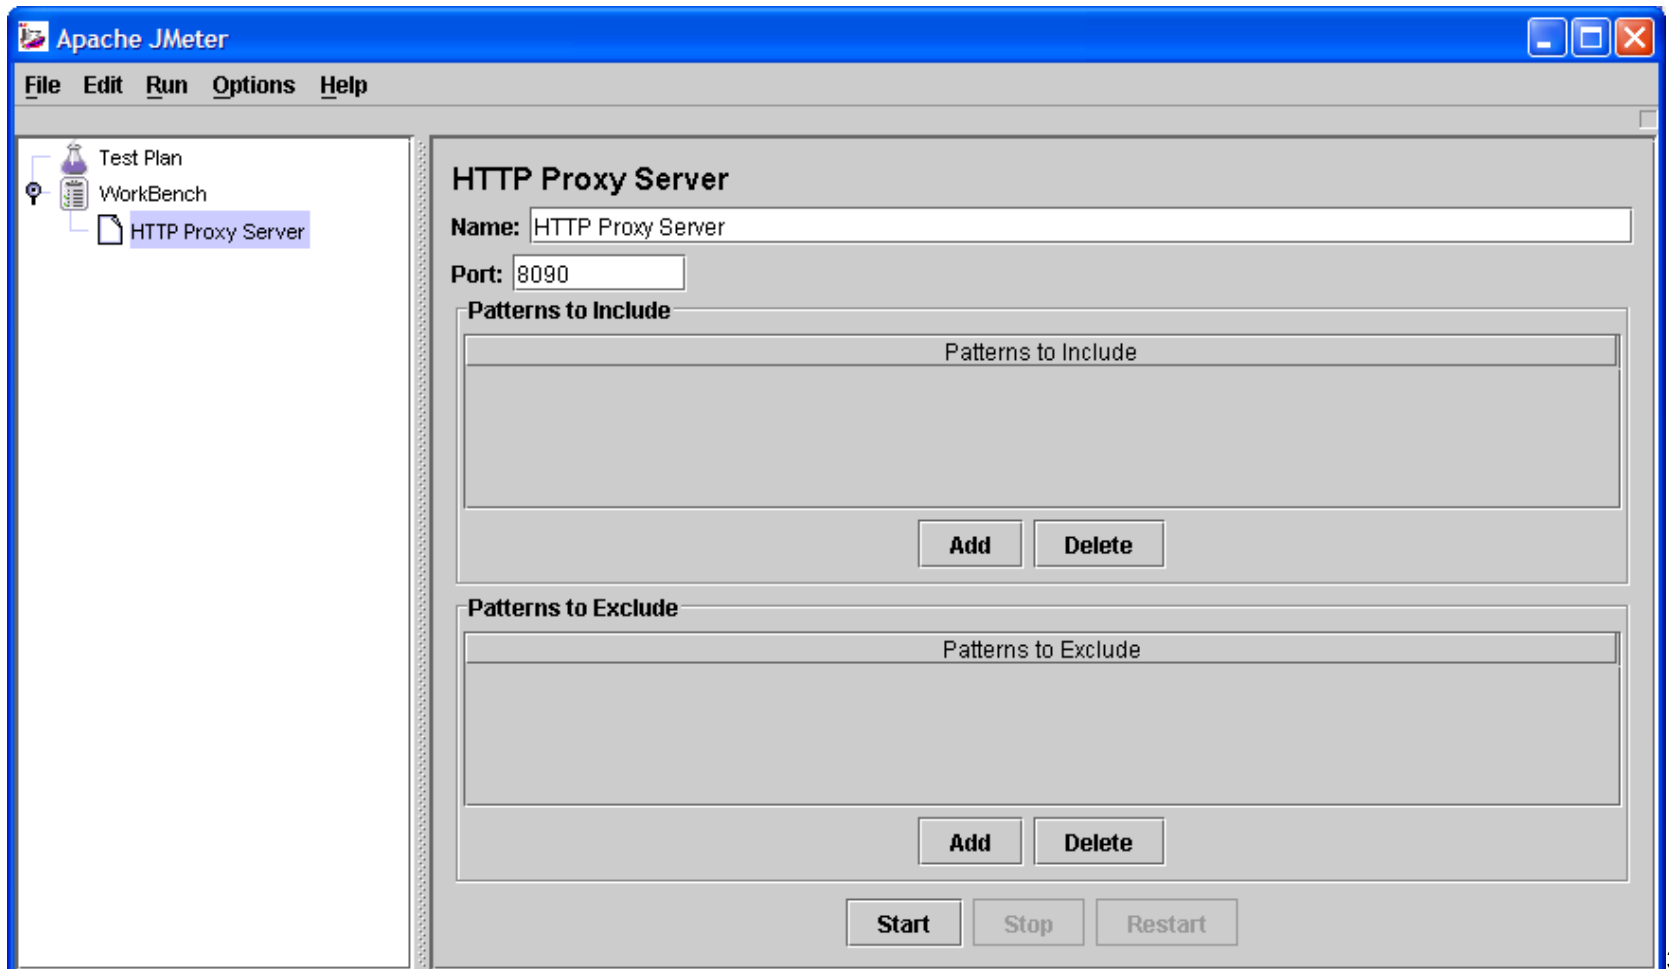


Stress Testing

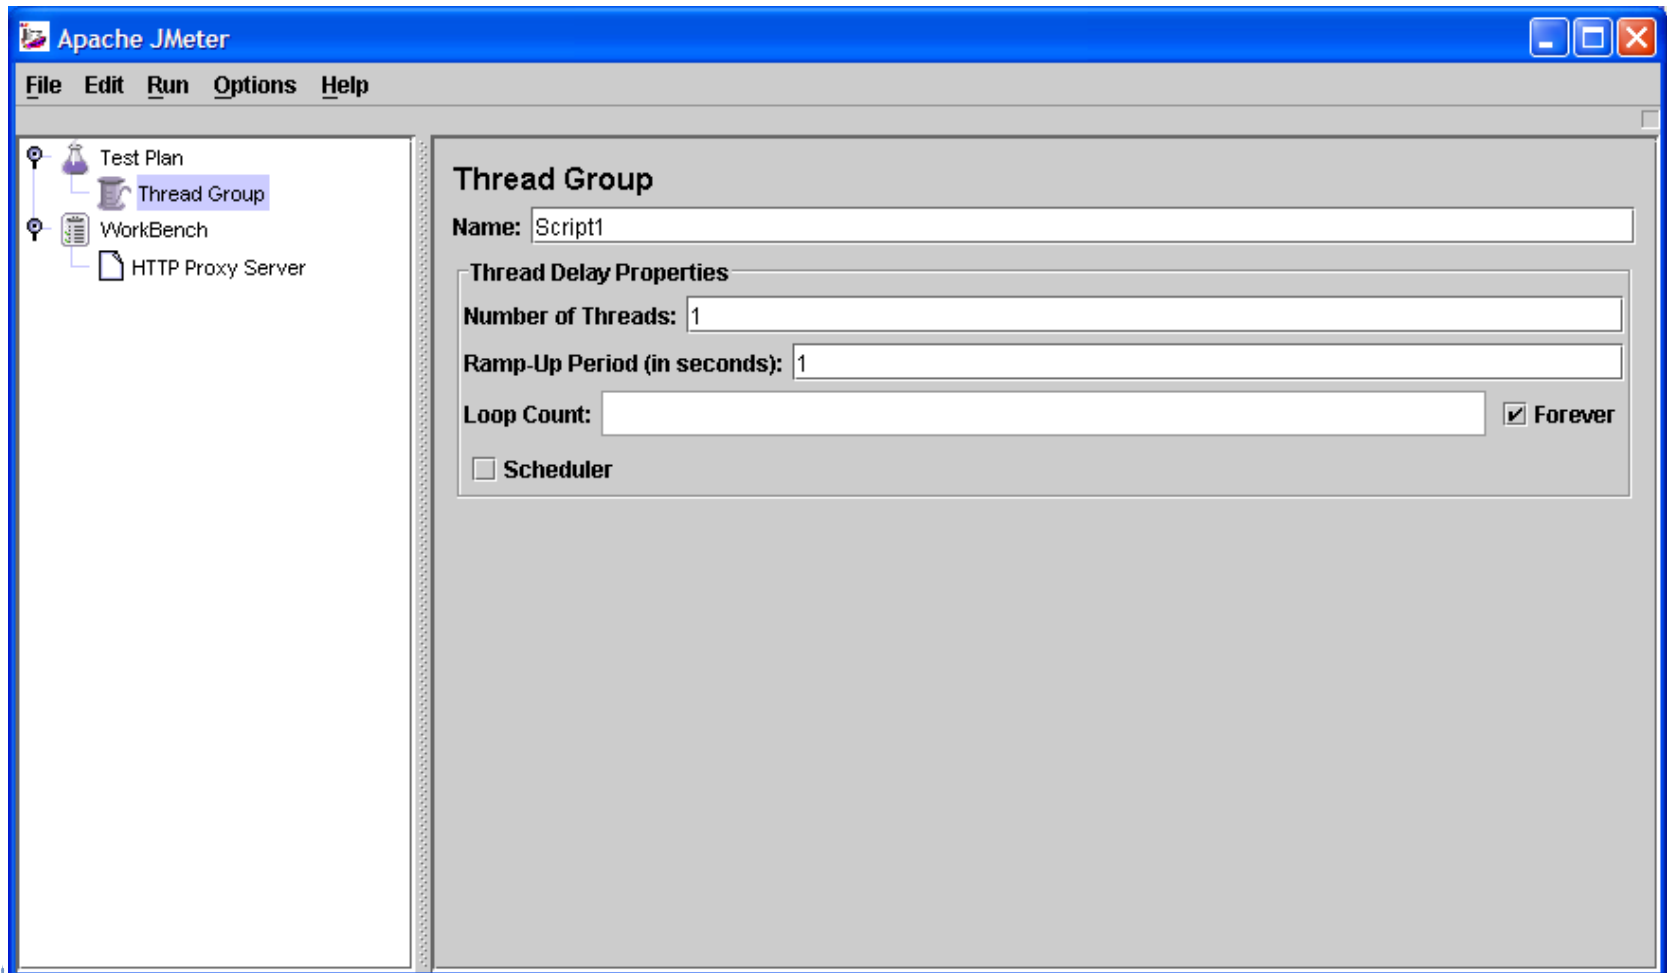
- 🌀 Stress testing tool feature list
 - ➔ Easily scripted to support many users doing many different things
 - ➔ Supports randomization of inputs
 - ➔ Throttles request rates
 - ➔ Randomized request rates
 - ➔ Reports on response times (from clients perspective)
 - ➔ Vary loads
 - ➔ Generate high loads
- 🌀 Introduced Apache JMeter to JoGoSlo



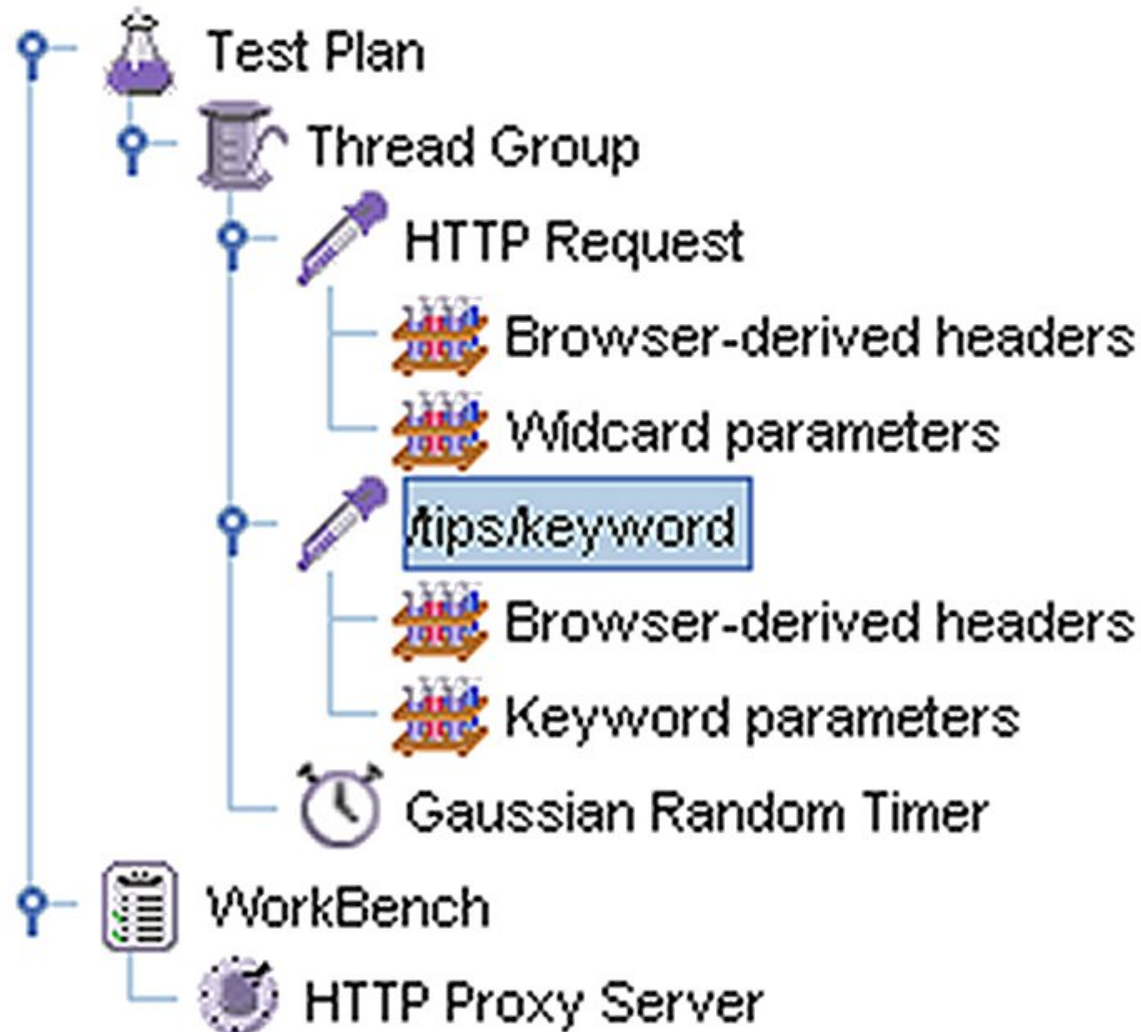
Apache JMeter



Apache JMeter



Apache JMeter



Apache JMeter

The screenshot displays the Apache JMeter GUI. On the left, a tree view shows the test plan structure: Test Plan, Thread Group, HTTP Request, Browser-derived headers, Wildcard parameters, /tips/keyword, Browser-derived headers, Keyword parameters, Gaussian Random Timer, WorkBench, and HTTP Proxy Server. The main panel is titled 'HTTP Request' and contains the following configuration fields:

- Name:** /tips/keyword
- Web Server:** Server Name or IP: localhost, Port Number: 8080
- HTTP Request:** Protocol: http, Method: GET (selected), POST
- Path:** /tips/keyword
- Options:** Redirect Automatically, Follow Redirects, Use KeepAlive
- Send Parameters With the Request:** A table with columns: Name, Value, Encode?, Include Equals?.

Name:	Value	Encode?	Include Equals?
keyword	\${keyword}	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Buttons: Add, Delete

Send a File With the Request: Filename: [text box] Browse..., Parameter Name: [text box], MIME Type: [text box]

Optional Tasks: Retrieve All Embedded Resources from HTML Files, Use as Monitor

Apache JMeter

The screenshot shows the Apache JMeter application window. The title bar reads "Apache JMeter". The menu bar includes "File", "Edit", "Run", "Options", and "Help". On the left, a tree view shows a "Test Plan" containing a "Script1" with four test elements: two "/app1/keyword" and two "/app1/wildcard". Below the tree are options for "View Results Tree", "View Results in Table", "Graph Full Results", and "Aggregate Report" (which is selected). The main area displays the "Aggregate Report" dialog. The "Name" field is set to "Aggregate Report". There is a section for "Write All Data to a File" with a "Filename" input field, a "Browse..." button, and a "Log Errors Only" checkbox. Below this is a table with the following data:

URL	Count	Average	Min	Max	Error%	Rate
/app1/keyword	6	647	0	1302	0.00%	1.5/sec
/app1/wildcard	6	40	10	50	0.00%	2.1/sec
TOTAL	12	343	0	1302	0.00%	2.9/sec




Apache JMeter Simple Setup

- Setup proxy
- Use browser to generate desired traffic
- Add in timers
- Randomize input
- Add in listeners
- Configure ThreadGroup properties
- Run load test





Practical



Instructions

- 🔄 Start up `apachejmeter.bat`
- 🔄 We'll skip the proxy setup. Load `mixed.jmx` JMeter plan
- 🔄 Add random delay that ranges between 1 and 4 seconds between calls for both keyword and wildcard
- 🔄 Add in a listener of your choice
- 🔄 Use 2 threads (concurrent users)
 - ➡ Don't forget to set the repeat count
- 🔄 Run and watch



Hardware

CPU, Memory,
Disk, Network

- Hardware is our physical constraint
- If we don't have enough
 - ➔ Get more
 - ➔ Reduce utilization of what we have
 - Strength reduction (algorithms)
 - Trade one resource for another
 - Caching trades memory for I/O
- Judge utilization in relation to the task at hand
 - ➔ Reading 1 megabyte from disk should not stress a modern I/O channel
 - Are you really reading 1 meg?



Measuring Hardware Unix

- System activity
 - ➔ Maintained in kstat structures by the kernel
 - ➔ Collection of counters
- Reported on by command line tools
 - ➔ Includes vmstat, mpstat, iostat
 - ➔ Values reported as activity since last call
 - ➔ Provides instantaneous view on how hardware is coping with load



Measuring Hardware Windows

- System activity
 - ➔ Maintained in registry
 - ➔ Collection of counters
- Reported on by taskmgr and perfmon
 - ➔ Graphical windows on system performance
 - ➔ perfmon is configurable
 - ➔ Taskmgr has few configurations
 - You can (and should) turn on reporting of system time (CPU)



CPU

- High utilization is easily measurable
 - vmstat (Unix) or taskmgr (Windows)
- Different types of utilization
 - ➔ Application
 - ➔ JVM
 - ➔ System/OS



Application

Source

- ⇒ Heavy workload
 - Add CPU
- ⇒ Remove processes from machine
- ⇒ Inefficient algorithms
 - Use method profiler to identify bottlenecks.
 - prof
 - hprof
 - NetBeans (JFluid)

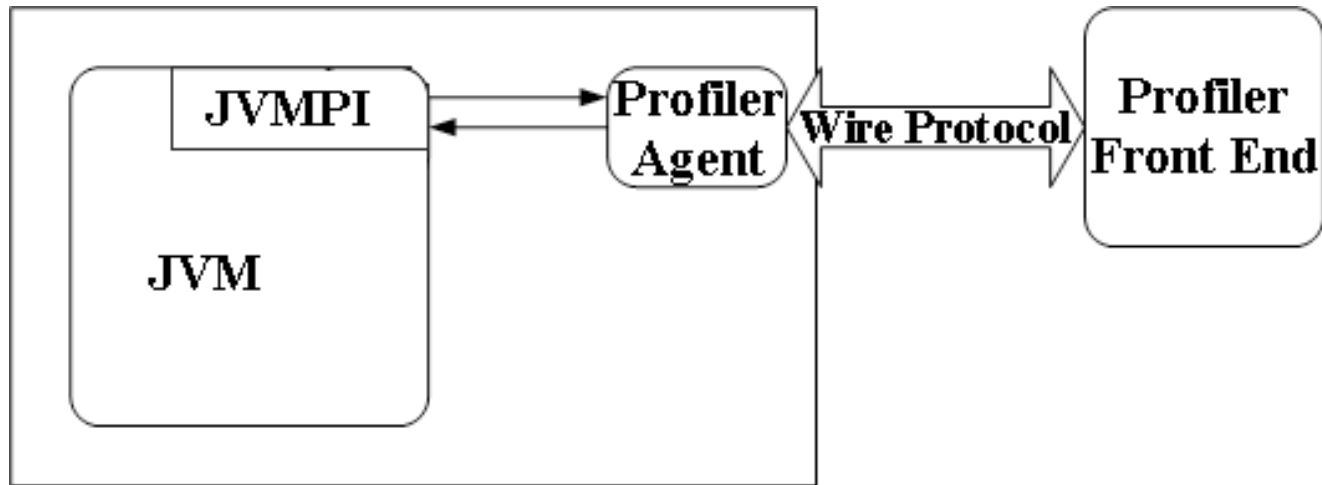


Application Profiling

- JVMTI interface
 - New to 1.5
 - Combination of old JVMPi and JVMDb interfaces
 - Supported by `-Xrunyourlib:parameters`
 - Loads yourlib (dll or so)
 - Initializes with parameters



Application Profiling



Application Profiling

-Xprof

➔ Original execution profiler

➔ Sampling profiler

- Adds 1 to a counter for each method when it is found at the top of stack
- Timings are inclusive

➔ Reports on a thread bases

➔ Dumps report to screen when thread dies



Application Profiling

- 🔄 -Xrunhprof
 - ➔ Original heap profiler
 - ➔ Extended for thread and execution profiling
 - ➔ Built off of JVMTI interface but no wire protocol
 - ➔ Much more data than prof
 - Best viewed with a profiling tool (HPJMeter)



Java Virtual Machine

- 🔄 Heavily threaded (measure with vmstat)
 - ➔ Runnable (r) queue consistently 2x number of CPUs
 - ➔ Stresses scheduler
 - ➔ Introduce thread pooling to limit activity
 - ➔ Reduce number of threads in current pools
- 🔄 Java heap management
 - ➔ Monitor gc with `-verbose:gc` flag
 - ➔ View output with HPJ Tune



Operating System

Context switching

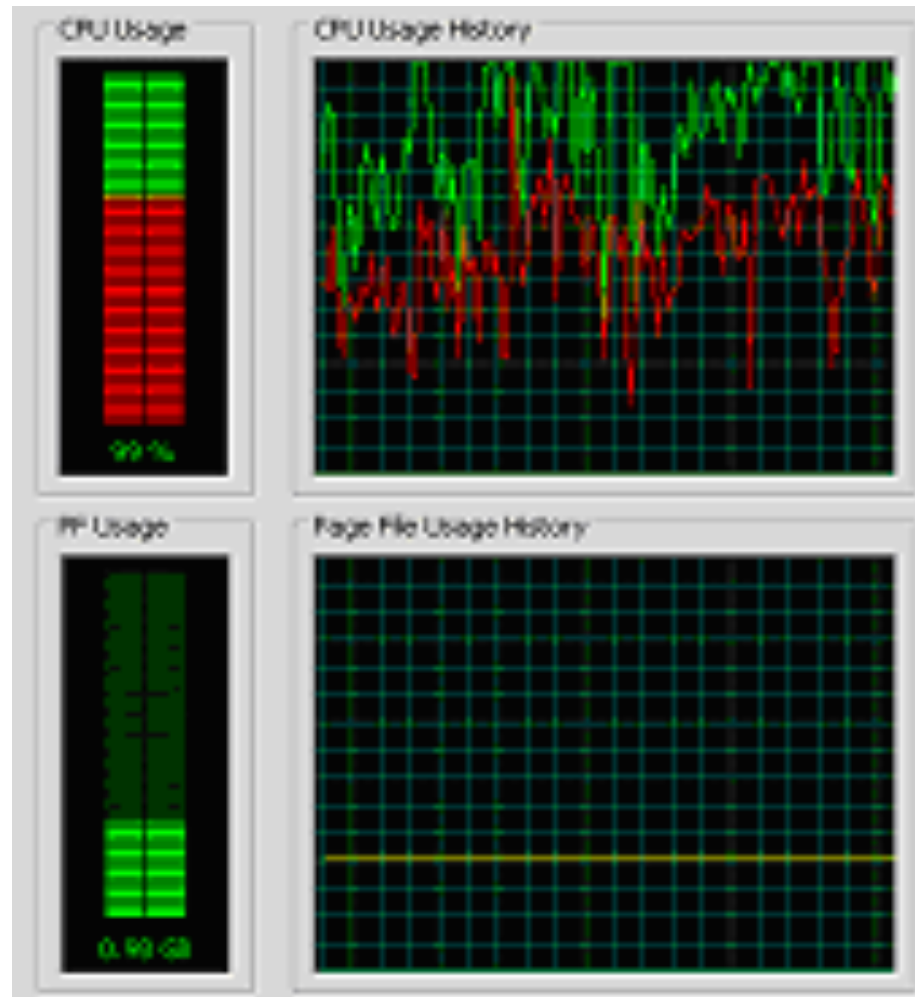
- ➔ Threads not completely using quantum
- ➔ I/O
- ➔ Lock acquisition
- ➔ Interrupt handling

Memory management

- ➔ non-zero scan rates (sr) for more than a few seconds at a time
 - OS is thrashing



Operating System



Memory

- High utilization is easily measurable
 - ➔ memstat (Unix) or taskmgr (Windows)
 - ➔ Can look like high CPU utilization
- Real memory
- Virtual memory
 - ➔ An outdated optimization
- Ideally we want to pin JVM into real memory
 - ➔ Eliminate paging
 - ➔ Reduce memory utilization
 - ➔ Add memory



Disk and Network I/O

- Heavy utilization will most likely prevent application from fully utilizing CPU
- Source (iostat)
 - Reading/writing large data sets or many network calls
 - Use counters to calculate rates
 - Use I/O channel specs to understand capacity
 - For disk, introduce buffering in hardware or application
 - E.g. Databases use paging
 - For network introduce caching
 - Bulk up operations
- Wrap I/O calls with timer



JDBC Monitoring

- 🔗 Common problem is interactions with database
 - ➔ Can measure activity using JDBC interceptor
- 🔗 P6Spy looks like a JDBC driver
 - ➔ Logs all JDBC calls
 - ➔ Logs can be viewed with IronEye



IronEye

IronEyeSQL

File View Server Help

Connect Disconnect Config Purge Import Export About

Legend: Preparation (blue), Execution (green), Retrieval (orange). Slowest (red), Most Run (yellow), Both (red).

Filtering (click to open)

SQL	Count	Avg Time	Max Time
SELECT LASTUSERID, OIDFIRMA, OIDYEARVALIDFROM, OIDBASEC...	30735	2	47
SELECT LASTUSERID, OIDFIRMA, OIDYEARVALIDFROM, OIDBASEC...	30735	2	47
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	30726	2	78
SELECT LASTUSERID, OIDCOSTCENTRE, OIDBUDGETCT, CALCU...	1031	10	16
SELECT LASTUSERID, OIDCOMPANY, OIDROOT, GROUPTYPE, CO...	228	2	16
SELECT LASTUSERID, OIDGROUP, OIDCOSTTYPE, ITEMPOSITION ...	226	2	16
SELECT LASTUSERID, OIDPARENT, OIDLANGUAGE, TEXT, OBJEC...	73	1	16
SELECT LASTUSERID, OIDGROUP, OIDCOSTCENTRE, ITEMPOSITI...	73	3	16
SELECT LASTUSERID, OIDGROUP, OIDCOSTCENTRE, ITEMPOSITI...	73	3	16
SELECT LASTUSERID, OIDCOSTTYPE, OIDCOSTCENTREBASE, OB...	73	2	16
SELECT LASTUSERID, OIDCOMPANY, MANAGER, VALIDFROM, VALL...	73	2	16
SELECT LASTUSERID, OIDCOSTCENTRE, CODE, SHORTNAME, VA...	73	1	16
SELECT LASTUSERID, SETTINGUSER, SETTINGFILE, SETTINGSEC...	61	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	1	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	1	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	3	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	3	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16
SELECT LASTUSERID, OIDBUDGETROW, OIDCURRENCY, BALANC...	54	2	16

SQL Statement Syntax: OIDACUMULATIVEECC, ISEXTENDED AUDIT, ISCHANGEABLE, ISAUTOREVALUATION, TARGETCOSTSCIS, CONCURRENTCUCALC, OBJECTTS, OBJECTID FROM KORESETUP_V WHERE (OIDFIRMA = ?)

Rows Returned: max avg min

Time Performance: max avg min

Count: no. of executions

Data loaded from V:\barcelona\logs\varial\server\spy4.log

Not Connected


JAMon 2.2

- Specify JAMon JDBC driver
- Can be viewed using supplied WAR file
- To bind it in without code or config changes:
 - ➔ <http://www.cretesoft.com/archive/newsletter.do?issue=136>





Practical



Instructions

- Make sure Tips is running
- Use 30 threads (concurrent users)
 - ➔ Don't forget to set the repeat count
- Run and watch the hardware
- What do we see?
- What do we do next?



- If hardware is able to cope with the load, move to investigate JVM
- Threading
 - Maybe hints of problem when investigating hardware
 - Examine threading with kill -3 or ctrl-break
 - Dumps activity to console
 - Look for many busy threads
 - Control level of threading using thread pooling
 - Traffic calming



Java Heap Memory

- Java Virtual Machine C/C++ process
 - Structure depends upon OS
 - Shared text
 - Stack
 - Heap
 - Java Heap allocated from process heap
- Java object allocated from Java heap space
- Java heap space managed by garbage collection
 - Object that are no longer reachable will be collected
 - Memory that is no longer referenced will be returned to the free list



Java Heap Space

- C struct defines Java object
 - OOP
 - Contains references to other object
 - Depends on the class declaration

```
public class A {  
    public Object x;  
    public Object y;  
}  
  
struct OOP {  
    int refCount;  
    byte *refs;  
} OOP, *OOP;  
  
...  
refs[0] = x;  
refs[1] = y;
```




Java Heap Space


- 🌐 Java heap maintains a references to OOP
 - ➔ Reference to all object maintained in OOP table
 - ➔ Root objects are at the top of object graphs
 - Define live objects
- 🌐 Object not reachable from GC roots will be collected
 - ➔ Three step process known as Mark and Sweep:
 - Traverse OOP table and clear mark bit
 - Traverse object graphs starting at GC roots and set mark bit
 - Sweep across OOP table de-allocating OOP structures



Mark & Sweep GC



Mark & Sweep GC



Mark & Sweep GC

- Triggered on allocation failure
 - ➔ `new Object();` fails
- Needs exclusive access to all of heap
 - ➔ Cannot share heap with application threads
 - ➔ Concurrency issue known as “stop-the-world” GC
- Single threaded
- Must manage entire heap space
 - ➔ Large heaps == long pauses



Mark & Sweep GC Optimizations

- When GC runs only 1 CPU is hot
 - Develop multi-threaded GC algorithms
 - Still have pause times but hopefully shorter
- Application pauses
 - Develop concurrent GC algorithms
 - Application and GC can run together
 - Reduced contention == reduced pause time
 - Higher overhead (ie trading CPU for shorter pause)




Mark & Sweep GC Optimizations


- Most Objects live for less than 100 μ s or for a long time
 - IBM defines pinned clusters, wilderness (not so generational)
 - Sun/HP/JRocket added Generation Spaces
- Generational spaces
 - Choose a different collector for young and old
 - Collect young first
 - Collect old only when there will not be enough room for old objects



Object Lifespan



Sun Generational Spaces



Generational Spaces

- Heap sizing
 - ➔ Can size generational spaces using ratios or absolute sizes
- -Xmx defines maximum size of entire heap
- -XX:MaxNewSize=<N>
- -XX:NewRatio
 - ➔ Ratios: 8 for -client and 2 for -server
- -XX:SurvivorRatio
- -XX:PermSize=<size>
- -XX:MaxPermSize=<size>
- Old space is what is left over



Survivor Spaces

$$\text{Eden} = \text{New Size} - \frac{\text{New Size}}{\text{Survivor Ratio} + 2} * 2$$

e.g. new size = 2M, Survivor ratio = 8

$$\text{Eden} = 2\text{M} - 2\text{M} / (8 + 2) * 2$$

$$= 2048\text{K} - 204.8\text{K}$$

$$= 1843.2\text{K}$$

Each Survivor Space = 102.4K



Monitoring GC

- `-verbose:gc` prints one log record for every GC event
 - ➔ `-Xloggc:file`
- Log entries provides a picture on how
 - ➔ your application is behaving
 - ➔ GC is coping
- Want to calculate GC throughput
- Want to find long GC pause times

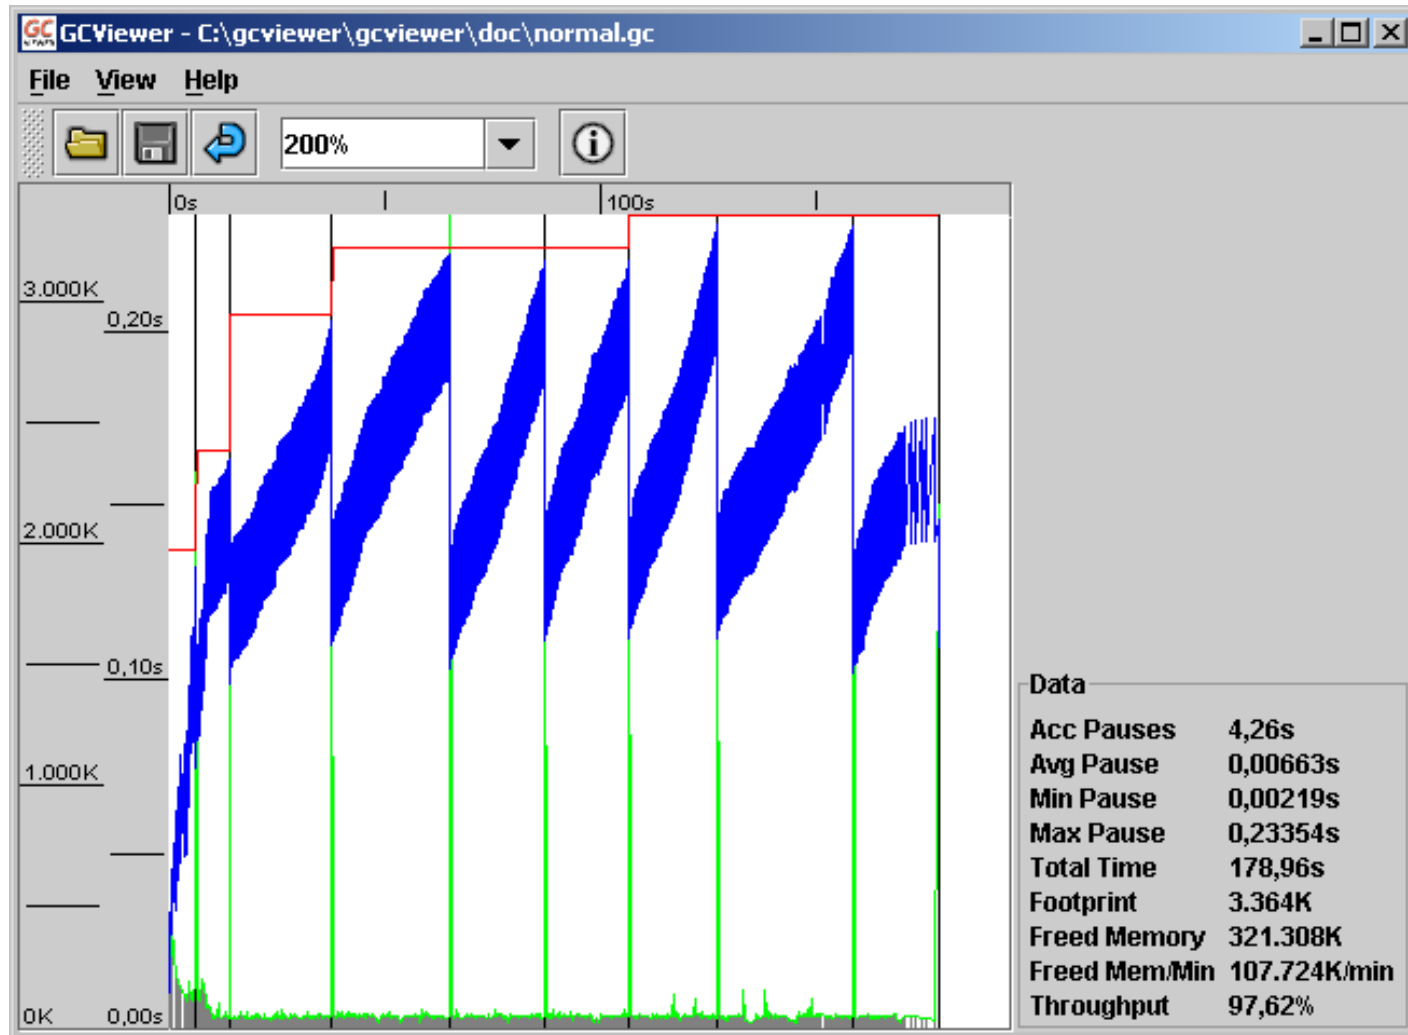


GC Throughput

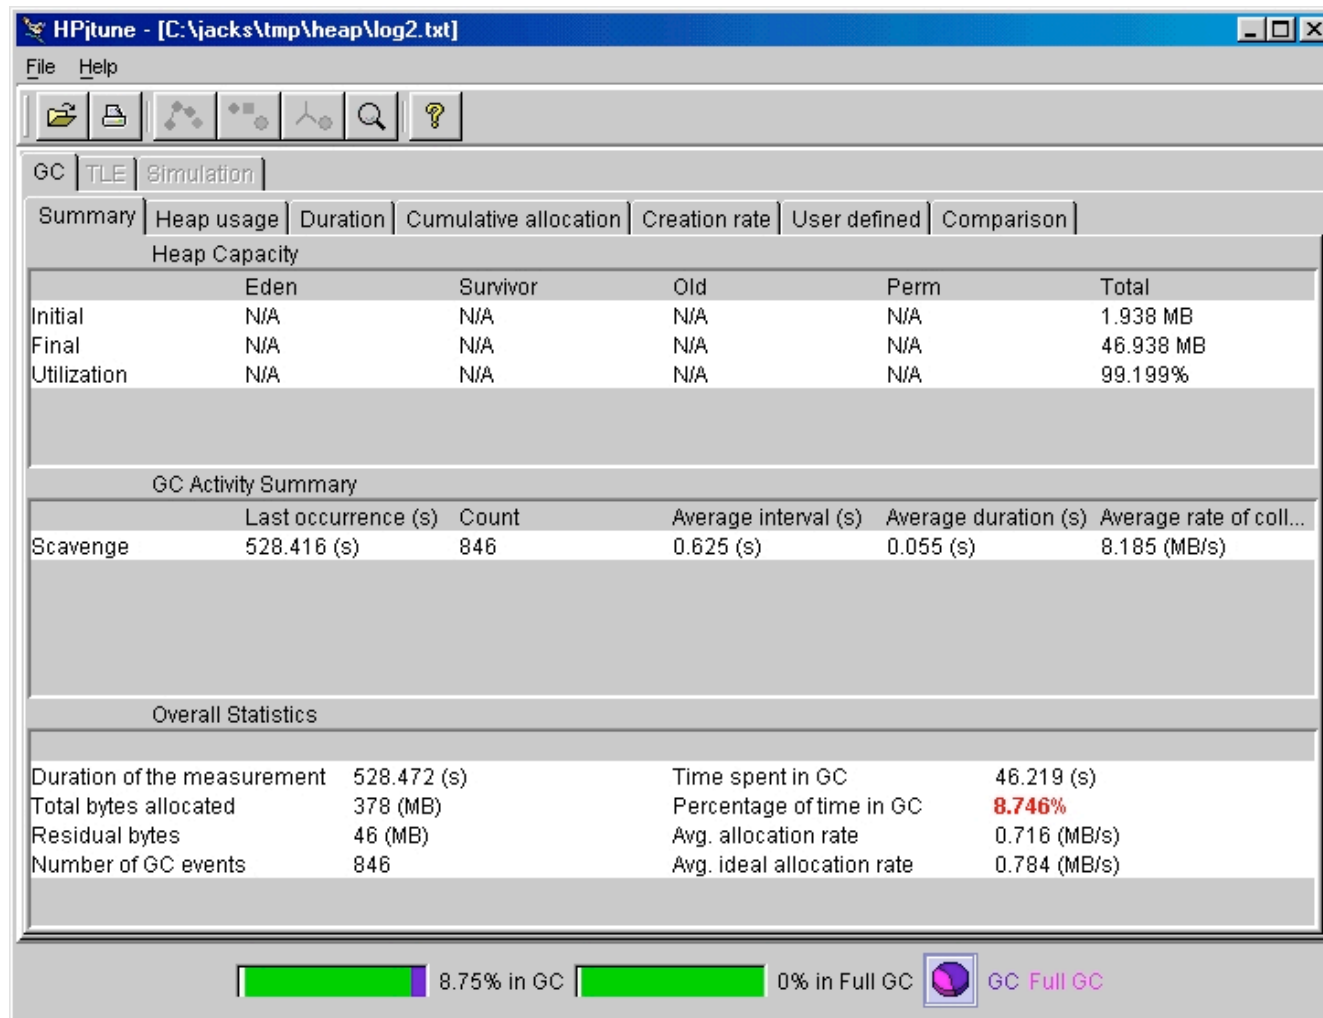
- ⌚ “Time application is suspended by GC” divided by “total run time”
- ⌚ E.g. 5 minutes of a 20 minute runtime is spent performing GC
- ⌚ 25% efficiency
- ⌚ GC bottleneck
- ⌚ Requires many records to calculate
 - ➔ Better tooling
- ⌚ GCViewer (TagTram)
- ⌚ HPJTune (HP)



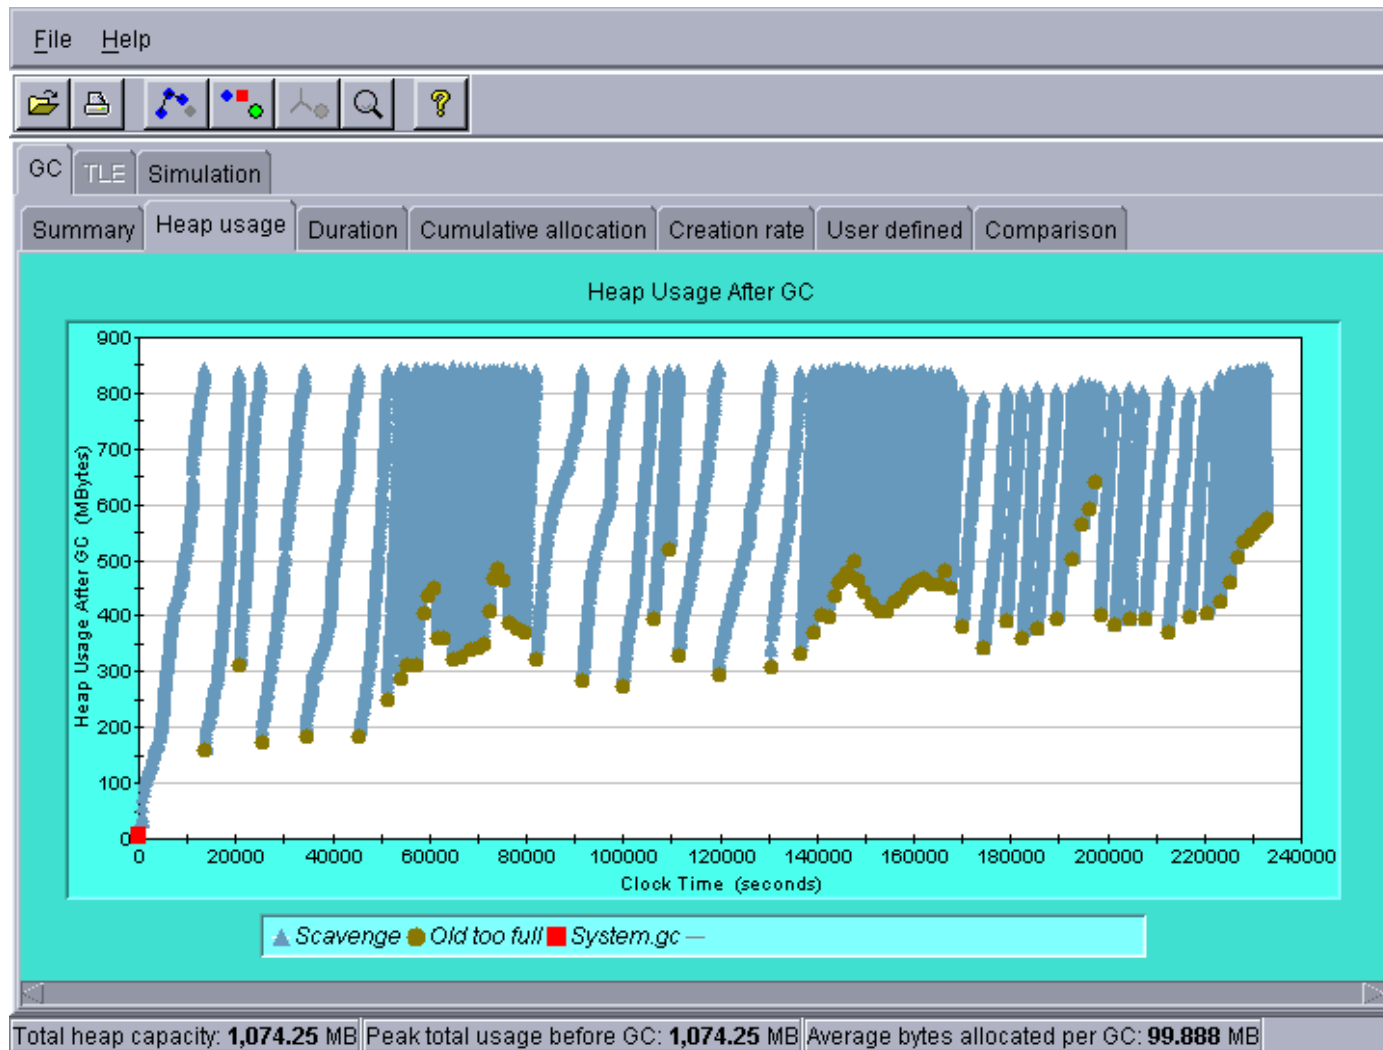
Tagtram GCViewer



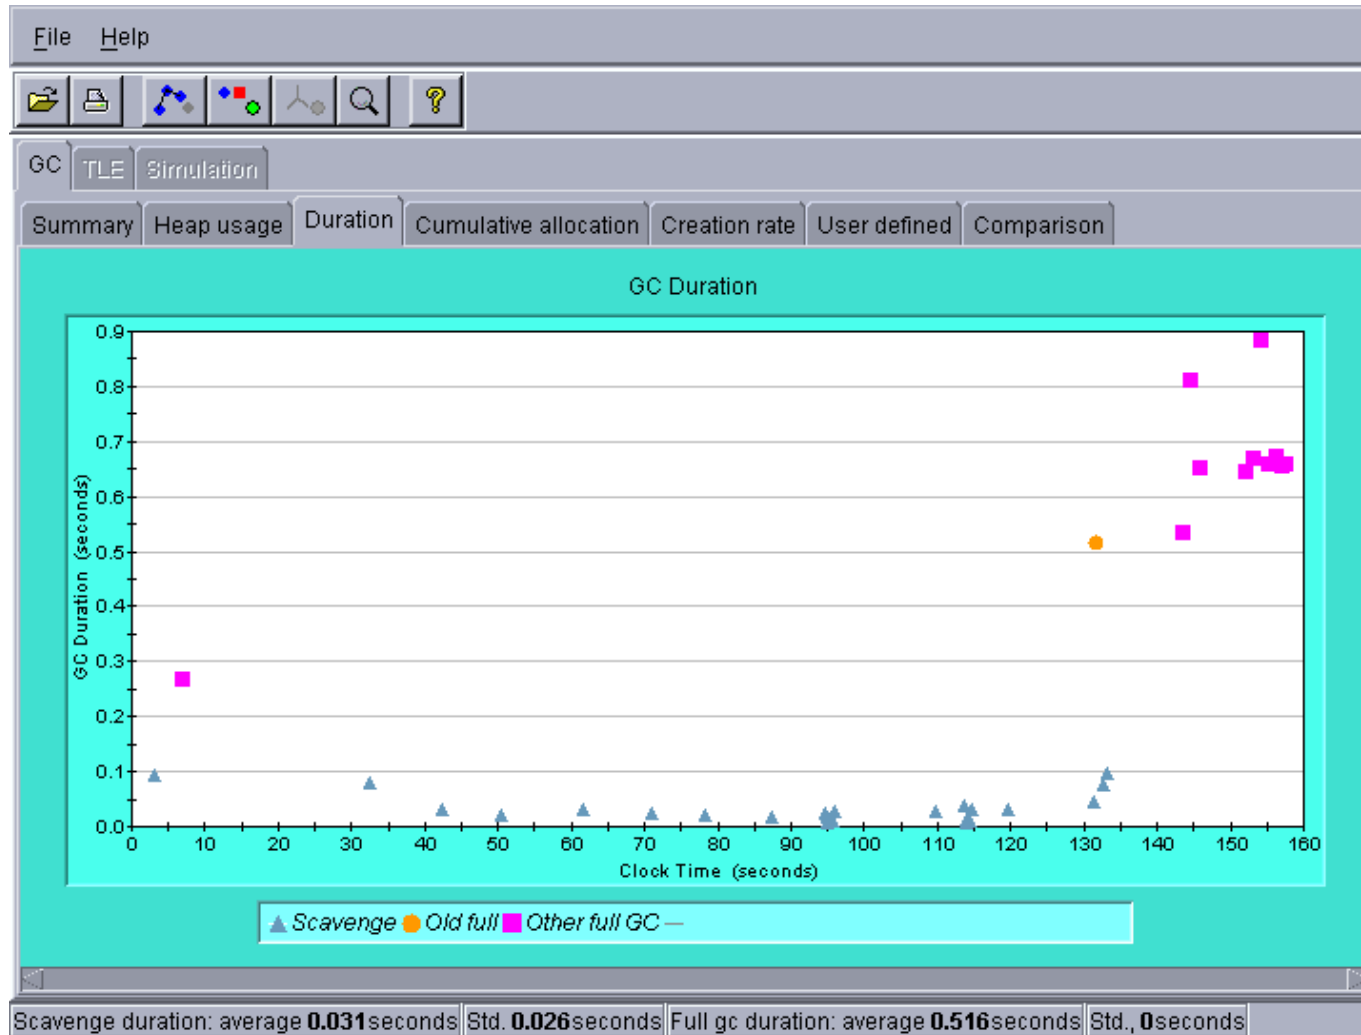
HP Jtune



HP Jtune Heap Usage After GC



HP Jtune Pause Time



Heap/GC Tuning

- 🔄 Use graphics to decide how to tune memory
 - ➔ Let the user experience to temper your choices
- 🔄 Strategy: eliminate full GC
 - ➔ Adjust size of total heap and survivor spaces
 - ➔ Tune other parameters as needed
- 🔄 Strategy: eliminate long pauses
 - ➔ Use Parallel (if multi-cored)
 - ➔ Use concurrent if you can tolerate overhead



Heap/GC Tuning

- ☞ Tuning GC cannot eliminate
 - ☞ Extremely high rates of churn
 - ☞ Temporal or permanent memory leaks
- ☞ Need to fix the problem in the code
 - ☞ Use a memory profiler to direct your search
- ☞ `-Xrunhprof:heap=all`
 - ☞ Dumps heap when JVM exits
 - ☞ Dumps with `kill -3` or `ctrl-break`
- ☞ `-XX:+HeapDumpOnOutOfMemoryError`
 - ☞ New for latest version of 1.6, 1.5, and 1.4



Heap Dump

- Contains enough information to reconstruct a picture of memory
- Picture contains references to all objects
 - ➔ Dead objects held by OOP table
 - ➔ Live objects
- Call GC twice before dumping heap
- Data volume and complexity calls for tooling
 - ➔ HPJMeter

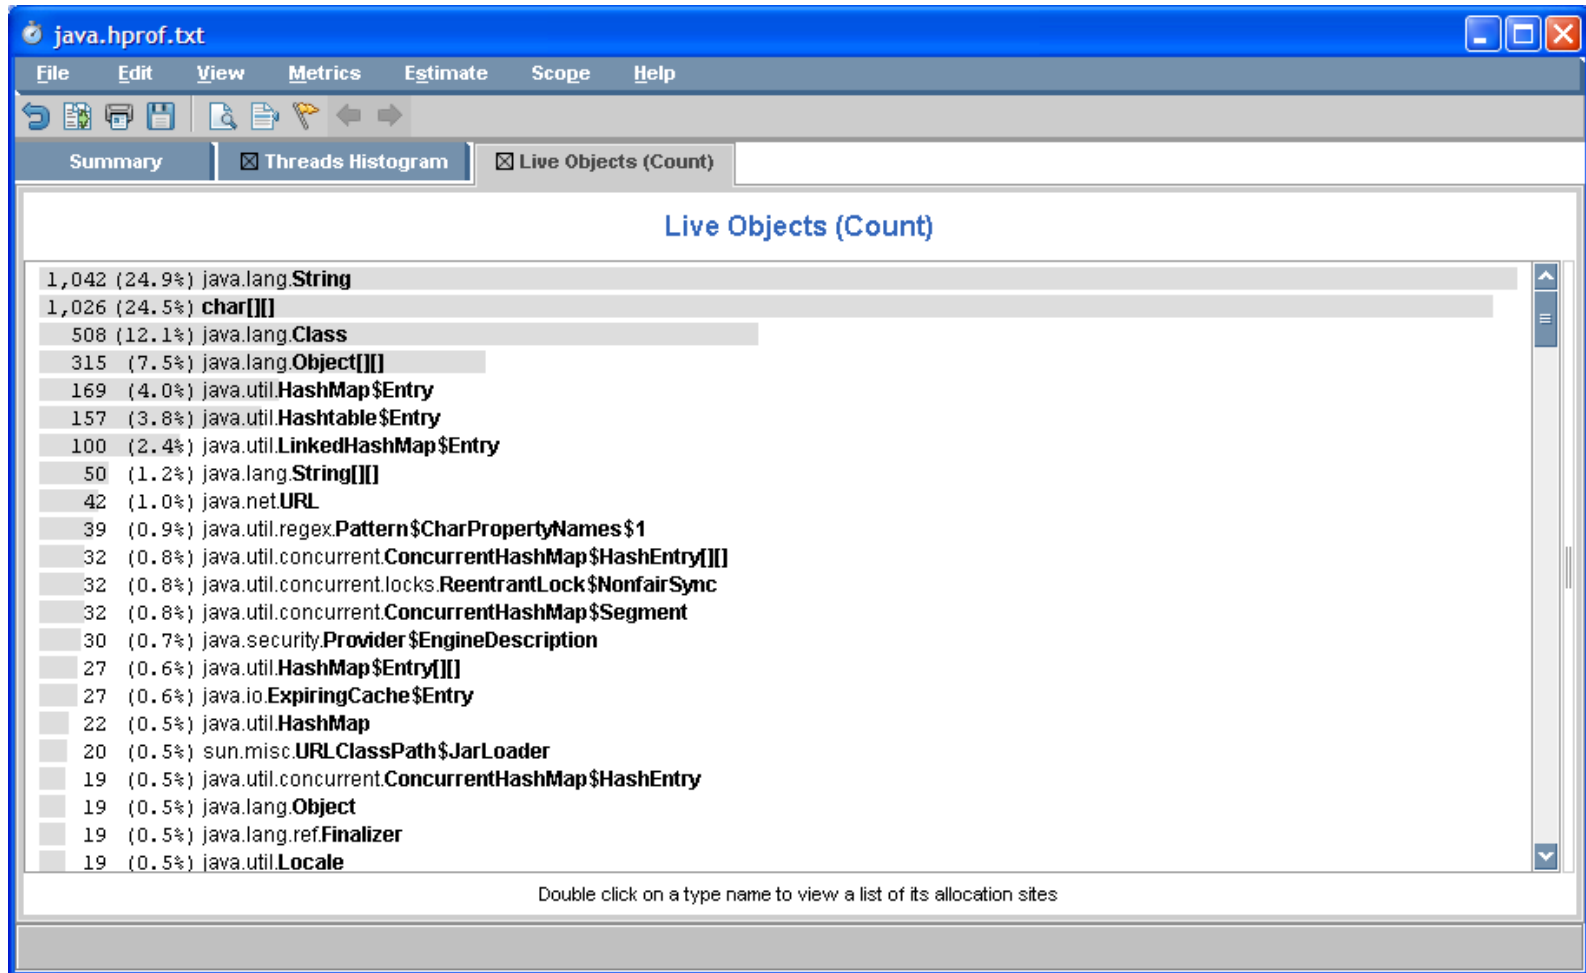


HPJMeter

- 🔄 Read hprof dump
 - ➔ Limited to single snapshot
- 🔄 Provides rudimentary views of heap
 - ➔ Live object numbers and sizes
 - ➔ Dead objects numbers and sizes
- 🔄 Can guess at memory leaks
 - ➔ Single snapshot analysis is limited
 - ➔ Can be good enough if you are methodical
- 🔄 Memory leaks usually are found in collections
 - ➔ Strategy: focus on collections



HPJMeter Live Object View




HPJMeter Leak Detection

The screenshot displays the HPJMeter interface with the 'Reference Graph Tree' view selected. The tree shows a hierarchy of objects, with several highlighted in blue to indicate memory leaks. The most significant leak is for the class `java.lang.System` at address `5000002d`, which holds 66,498 bytes. Other notable leaks include `java.io.OutputStreamWriter` (8,240 bytes), `java.io.PrintStream` (25,057 bytes), `java.io.BufferedWriter` (16,424 bytes), `char[]` (16,384 bytes), `java.util.Properties` (8,004 bytes), and `java.util.Hashtable$Entry[]` (7,956 bytes). The interface includes a menu bar (File, Edit, View, Metrics, Estimate, Scope, Tree, Help), a toolbar with navigation icons, and a tabbed interface with 'Live Objects (Count)' and 'Live Objects (Bytes)' selected. A legend at the bottom identifies object types: Leaf (black), Expanded Here (grey), Expanded Elsewhere (red), Visited (purple), and Other (blue). A note at the bottom states: 'Double click on object name to view all references to it'.



Practical



Instructions

- 🔗 Let's profile heap with JVM switch `-Xrunhprof:heap=all`
 - ➔ For fun, add switch `-Xloggc:gc.log`
- 🔗 Restart application server and run JMeter plan
- 🔗 Confirm that there is a memory leak with HPJTune
 - ➔ Open `gc.log`
 - ➔ Look at "Heap Usage After GC"
 - ➔ Look at "GC Duration"
- 🔗 Open HPJMeter and find the leak
 - ➔ You may need to shut down everything first



Application Lock Contention

- The only problem left is lock contention
- Characterized by inability to utilize CPU
 - ➔ Similar to I/O bound (call to external system)
- High system time (% of total)
 - ➔ Locks are a kernel resource
- Find by performing a thread dump (kill -3)
 - ➔ For live lock you may need many thread dumps
- Techniques to reduce lock contention is an emerging topic



- If you haven't found anything
 - Re-investigate the people
 - Are they really doing what you think they are doing?
 - Read logs
 - Visit the floor and watch
 - Re-do usage patterns
 - Compare JMeter scripts with real life
 - Re-test
 - Validate that QA == Production
 - Even the smallest difference can hide the problem



JoGoSlo Reload

- Introduced Apache JMeter
- Introduced HPJTune to monitor memory
- Confirmed memory leak hypothesis
 - Resting the application allowed application to recover
 - Recovery was tied to HttpSessionState timeout
 - Developers were working on persistence framework
- Isolated memory leak to single usage pattern
 - Filtered off a vast majority of the application
 - Identified, fixed and re-tested with-in budget




Summary

- 🔄 Systems are dynamic, code is static
- 🔄 Be methodical
- 🔄 Review performance requirements
- 🔄 Prepare stress testing environment
- 🔄 Define Usage patterns
- 🔄 Investigate hardware, JVM, and Application
- 🔄 Use measurements from tooling to direct your efforts
- 🔄 Let the user experience guide your decisions




Measure Don't Guess





Q&A



Thank you for your attention!

