An Embedded Service Platform for Uninterruptible Processing

Lessons Learned

Tim Biernat
Paul Schmirler

Rockwell Automation

Agenda

- Industrial Automation
- Decisions, Decisions
  - Which Java? What Container? Datastore?
- OSGi
- Embedded Database
- Demo
- Some Challenges
  - Performance
  - Flash Memory
  - Hardening
  - Troubleshooting
- Q & A

Introductions

Rockwell Automation

- 20,000 employees, $4.8 billion sales
-
- Automotive, Food & Beverage, Pharma, Material Handling, Mining, Oil & Gas, Electronics, and more
- Components, motor drives, industrial control and information systems

- Tim Biernat
  - Worked with General Dynamics, Motorola, IBM and SoftwareMentor
  - Interests: java, distributed computing, real-time fault-tolerant systems
- Paul Schmirler
  - Worked with eFunds, Eagle Technology
  - Interests: mobile computing, cloud computing

Industrial Automation Primer

Many kinds of production processes

- Discrete (auto assembly)
- Batch (beer brewing)
- Continuous (metal production)

- Challenging Environments
  - Hot, cold, dusty, wet, EM, G shock
- Safety Concerns

Manufacturing is extremely competitive

- Downtime unacceptable

- Long-lived systems
  - 15 to 20+ years in service is not uncommon
  - Maintenance, spares, support can be a real challenge

Where will my Linux kernel be 10 years from now?

Software in Industrial Automation

Growing role

- Visualizing, communicating, integrating, controlling, monitoring
- Historically dominated by MS tech: Windows OS, OPC (OLE for Process Control), D/COM, VBA
- Desire to connect factory with enterprise

- Many different platforms
  - Cloud, Virtual, PC, embedded (ARM, x86)
  - Windows, Linux, RTOS
- Java is compelling
  - 3Ps: Portable, Productive, Performant
  - Large open source palette
- Java challenges
  - Largish footprint for many embedded applications
  - Need for fast response, deterministic execution
  - High level of abstraction ➔ isolated from hardware

Decisions: Which Java?

Requirements

- embedded, headless
- target multiple hardware architectures
- full SE (Standard Edition) APIs
- performant as native code

-
  - OpenJDK, Oracle SE, Oracle Embedded, proprietary
- Discoveries
  - OpenJDK JIT still immature on ARM
  - benchmarks indicated decent Oracle Embedded performance
  - Oracle Embedded Java is JRE only, not JDK; client JVM only

Benchmarking

Desire to evaluate various hardware/JVM combinations

- X86, ARM
- OpenJDK, Oracle SE, Embedded JREs

- Representative applications
  - service provider
  - embedded datastore
  - IDE
- DaCapo (http://dacapobench.org/

external concurrency, multiple iterations, JVM warm up

x86 JVMs – IDE Bench

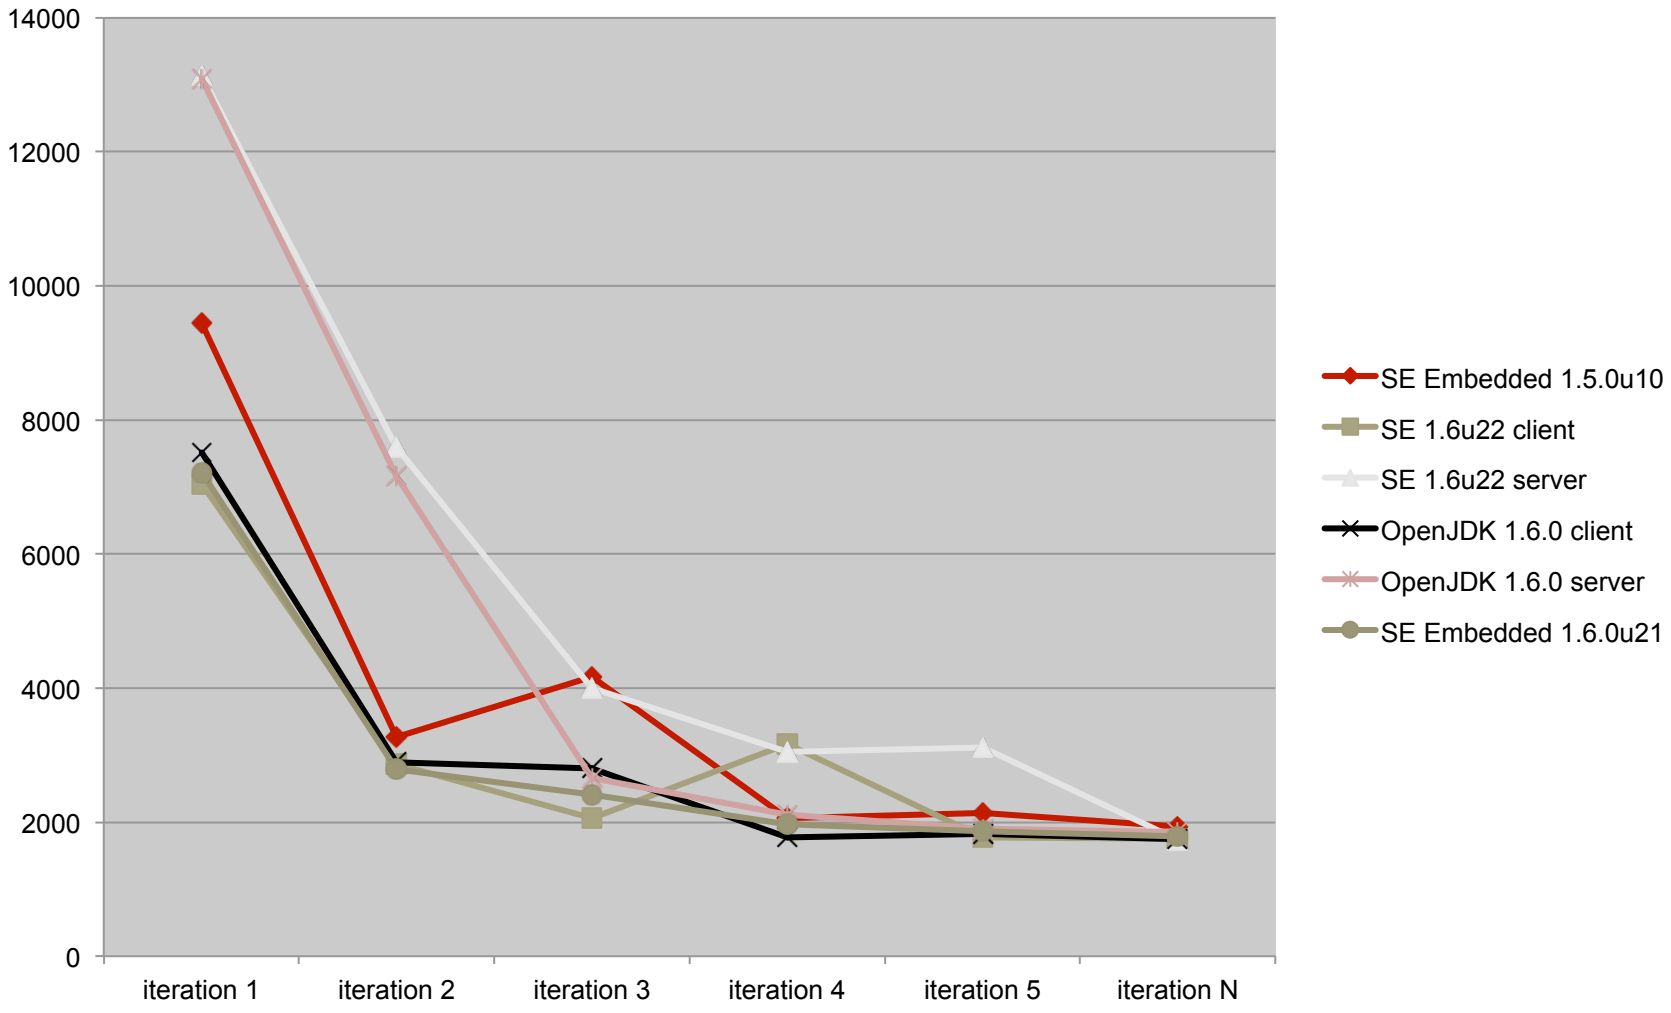

ARM & x86 – IDE Bench

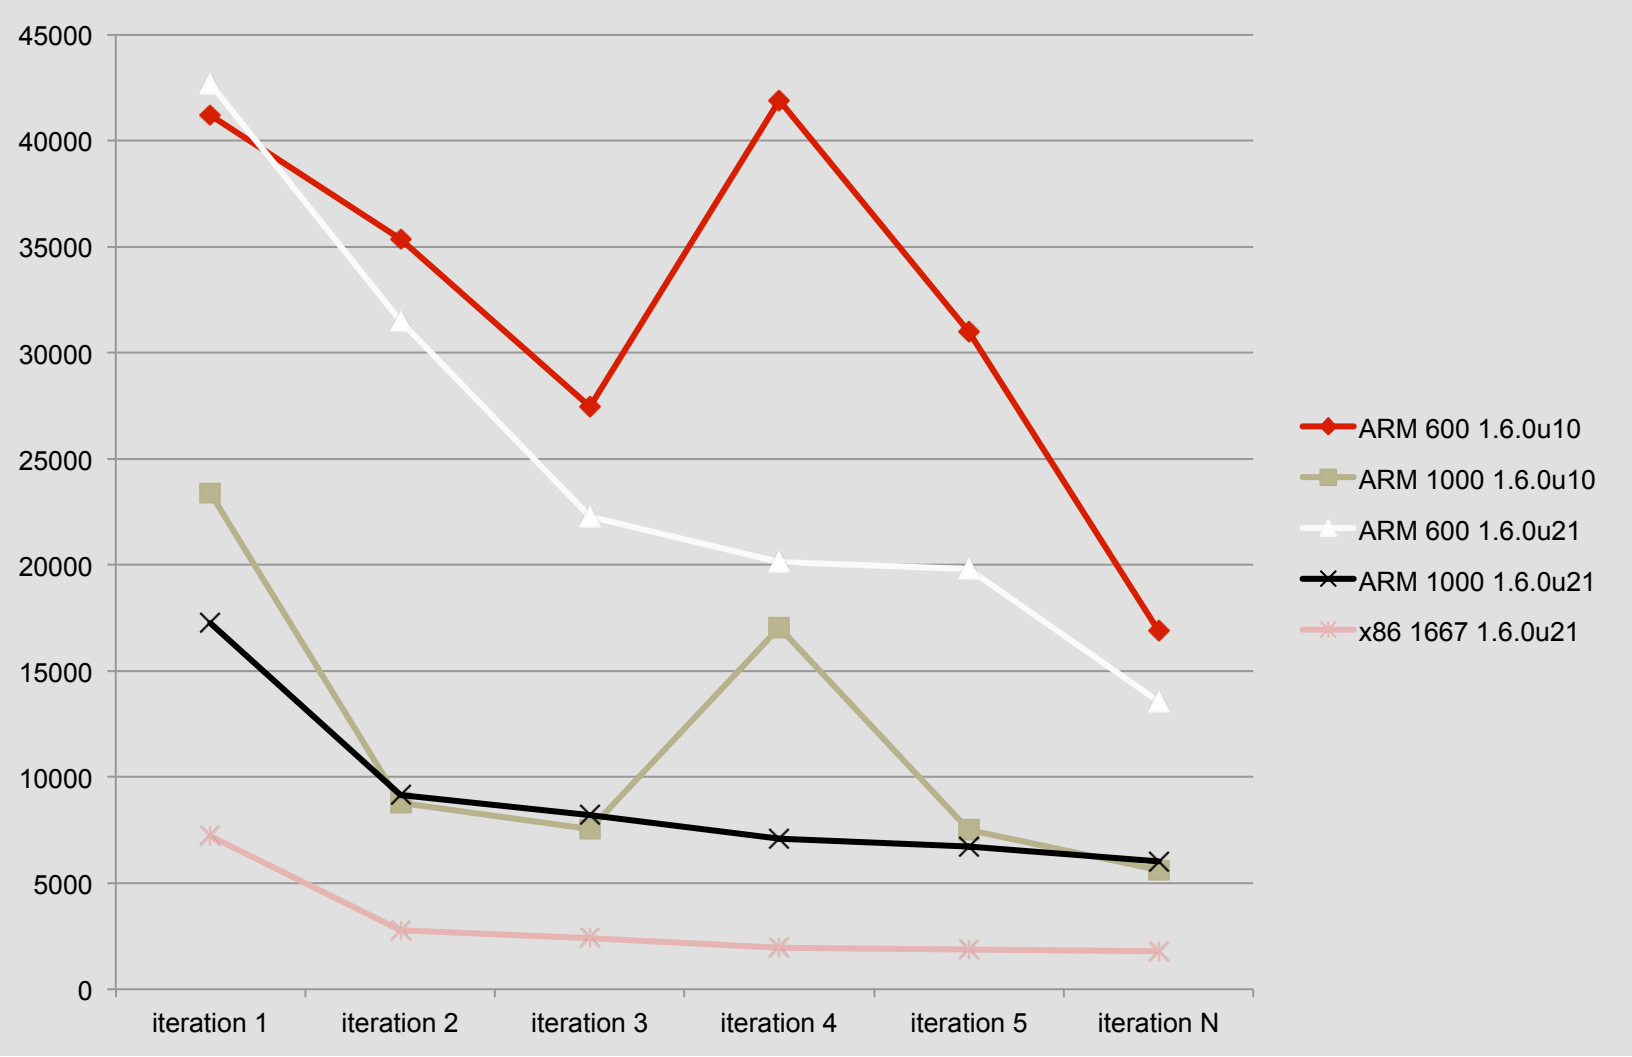

Decisions: Which Container?

- Requirements
  - embeddable (lightweight, proven, manageable)
  - modular deployment
  - dynamic “hot” deploy, update; concurrent software versions
  - Dependency Injection (DI) to minimize hardwiring, reduce coupling
  - support for native code
- Options
  - iPOJO, Guice, SpringDM, straight OSGi
- Discoveries
  - SpringDM is feature full, well integrated DI support
    - but fairly steep learning curve
  - OSGi a good fit
    - lightweight, mature and manageable
    - mature implementations (using Equinox, Felix DI was incomplete)
    - course-grained DI support, introduce full DI framework later

OSGi

- Modularity
  - Bundle – Physical and logical unit of modularity
  - Classloader model – Classloader per bundle
  - Imports/Exports – Restricts visibility to public API
  - Identity – Bundle-SymbolicName + Bundle-Version
  - Native code – Embedded in bundle / multi-platform
- Lifecycle
  - Dynamic – Independent of JVM
  - States – Installed, Resolved, Starting, Active, Stopping, Uninstalled
  - Activators – Hook for lifecycle events / access to OSGi framework
- Services
  - Decoupled, Dynamic, Pluggable
- Management Console

OSGi LIfecycle

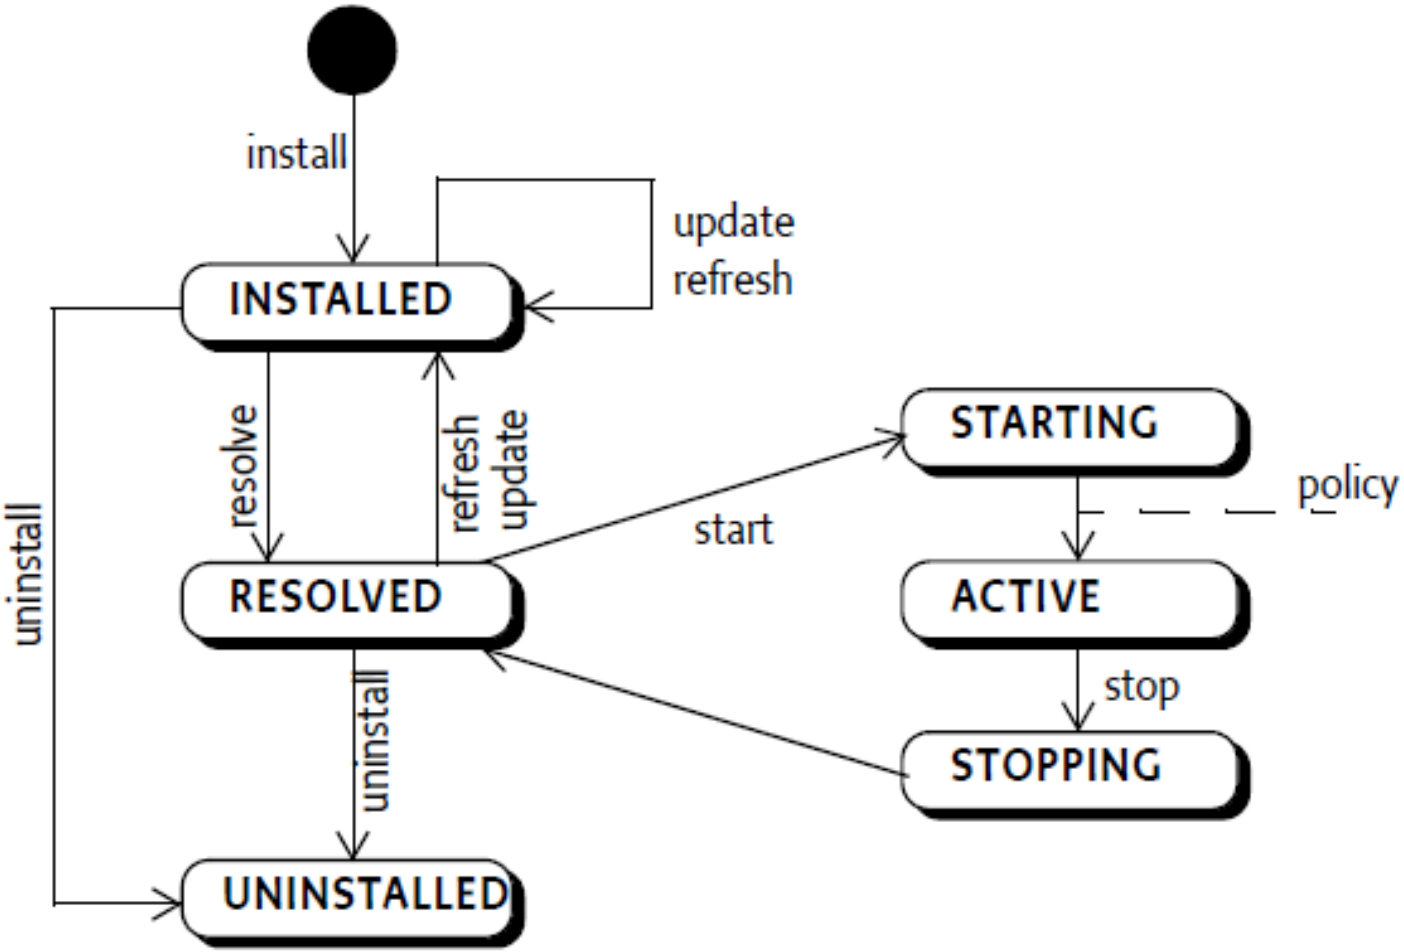

ref. OSGI Service Platform Release 4, The OSGi Alliance

OSGi Footprint

- Single OSGi JVM vs. multiple service provider JVMs
  - shared Java code loaded once
  - shared 3rd party libraries loaded once
  - shared native code loaded once
  - additional one time OSGi runtime overhead: 2 - 3 MB

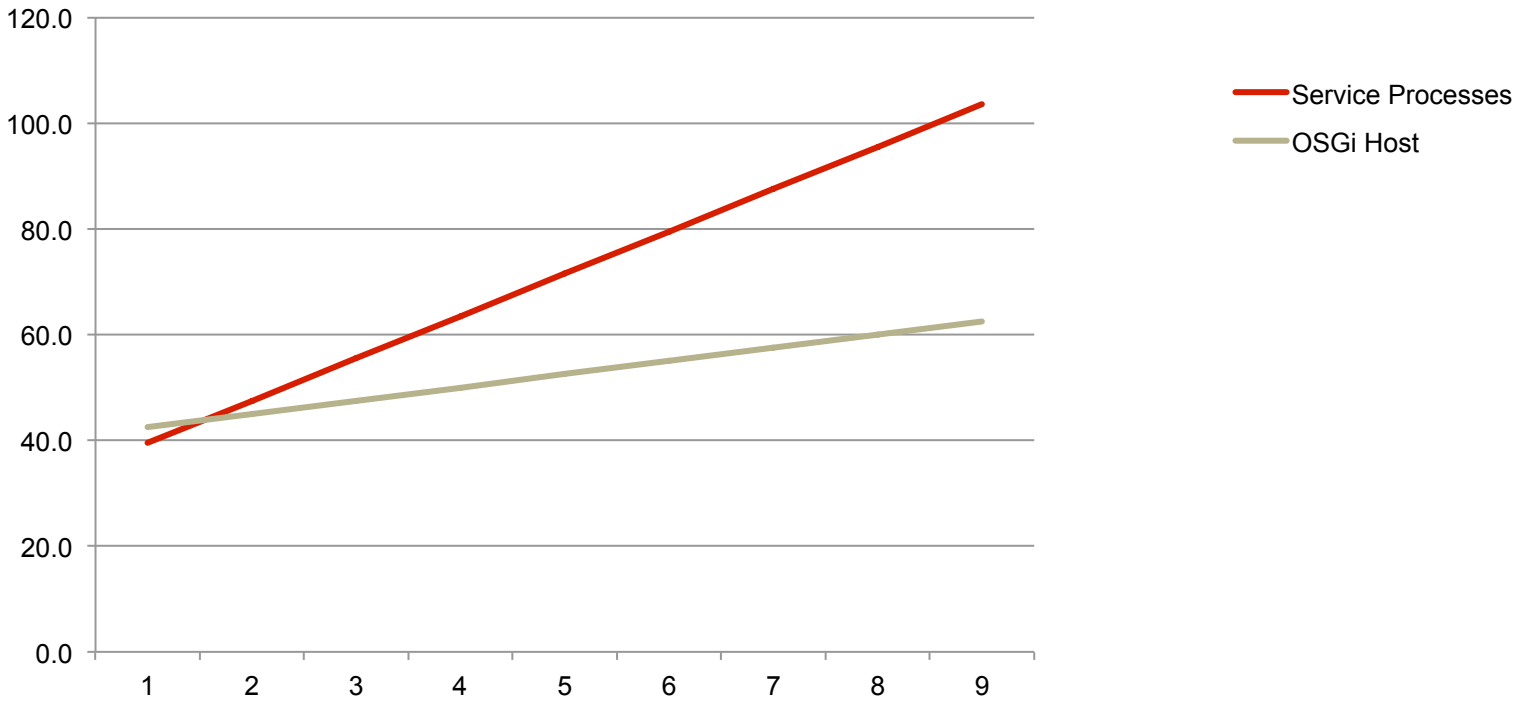

H2 Database

- Small footprint (1MB jar), pure Java solution
- Capable with 16MB of heap
- Best performance in class (vs. Derby, HSQLDB, PostgreSQL)
- Well-tested, good support ecosystem
- Other features
  - standard SQL support
  - dual open source license
  - fully transactional
  - highly tunable (buffers, cache, sync)
  - embedded and client-server modes
  - user-defined functions and stored procedures (in Java!)
  - built in full-text search or Lucerne support
  - built in profiling and performance statistics
  - engine-level encryption (2-3X slower)

Stack

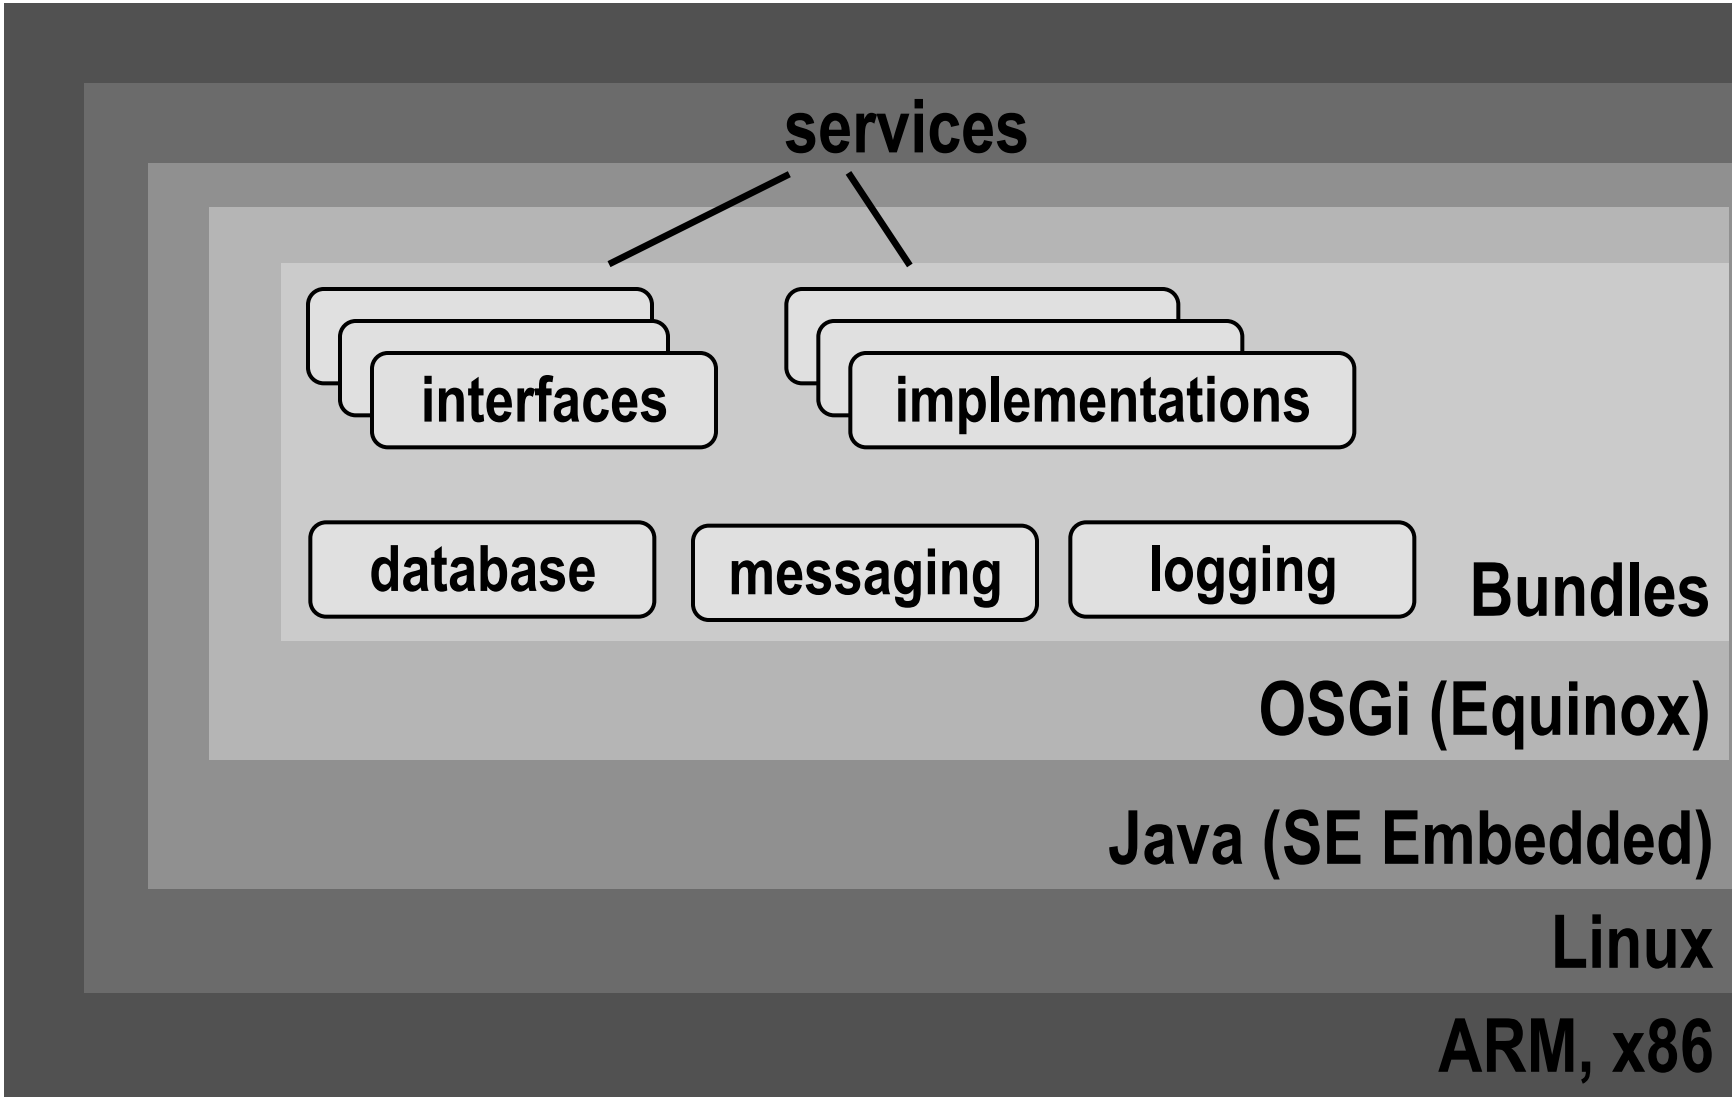
services
interfaces
implementations
database
messaging
logging
Bundles
OSGi (Equinox)
Java (SE Embedded)
Linux
ARM, x86

Cloud Architecture

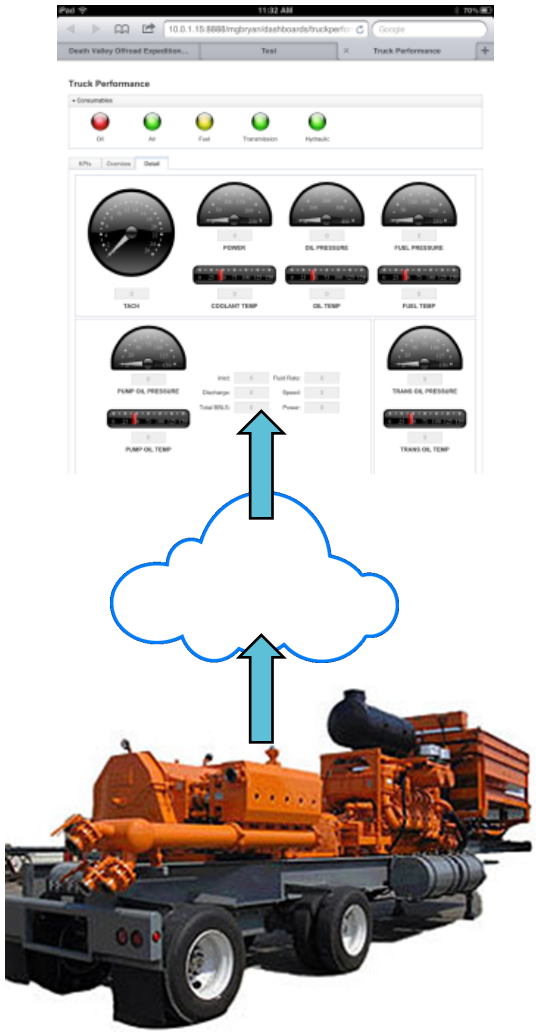

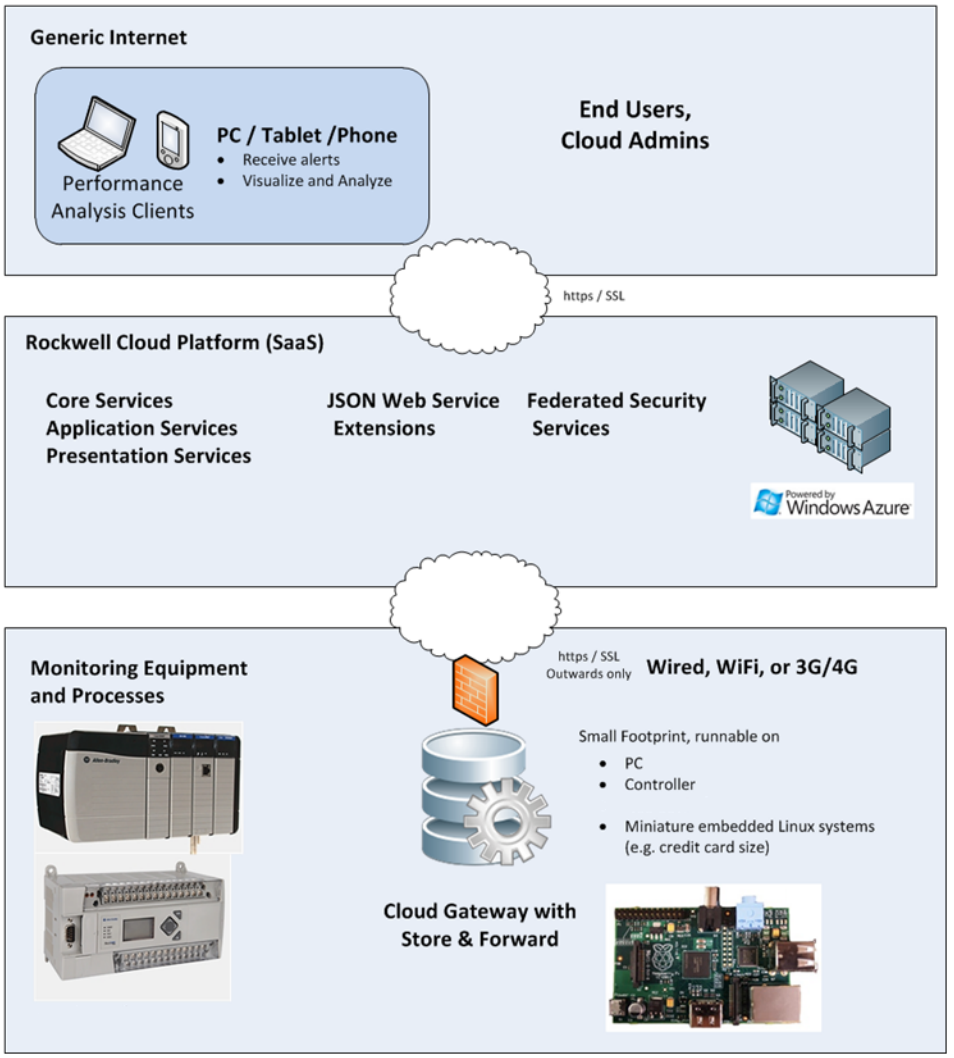

Demo - Cloud Gateway

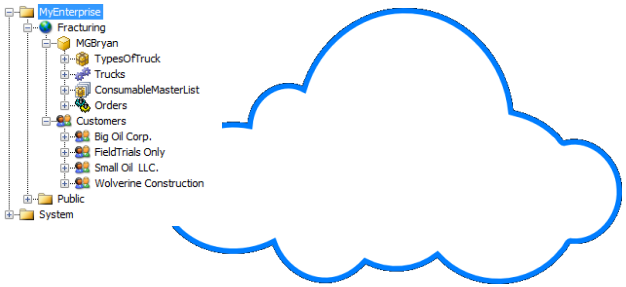

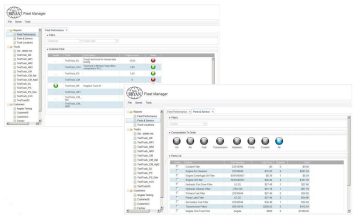

Cloud Gateway

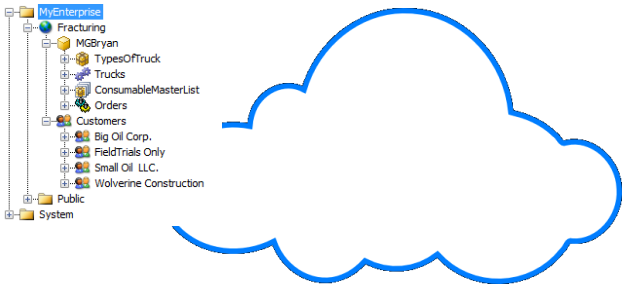

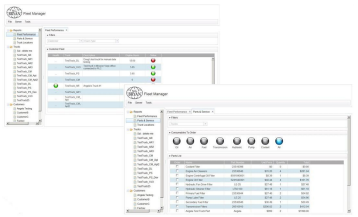

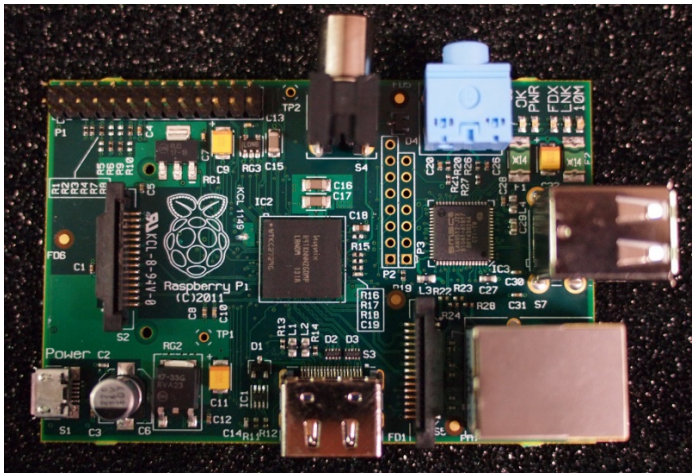

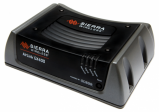

Cloud Gateway

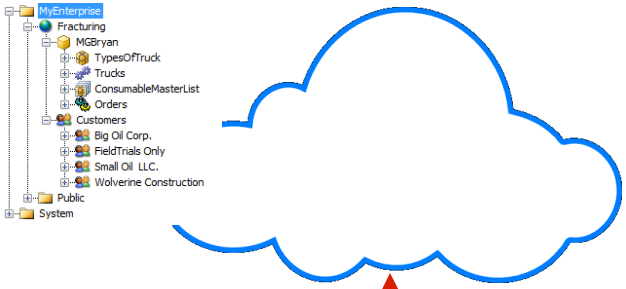

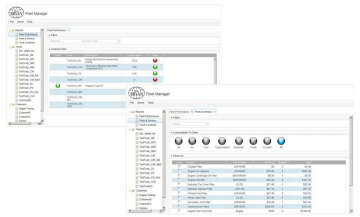

HTTPS / Basic Auth ↑ RESTful / JSON

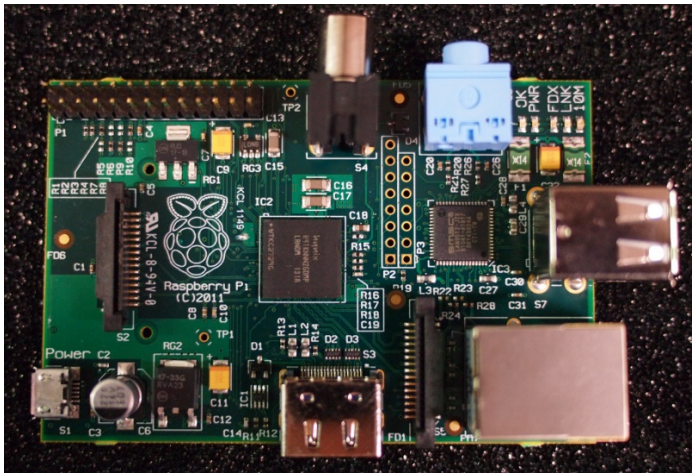

CIP

PCCC

TAIP

(GPS)

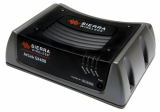

Cloud Gateway

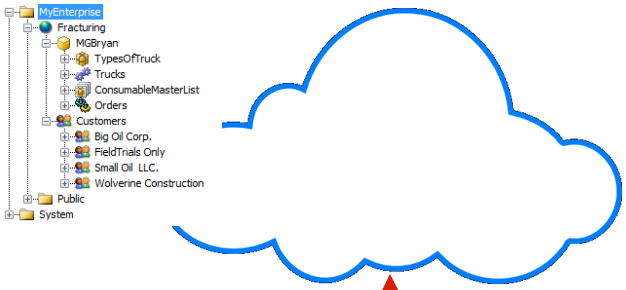

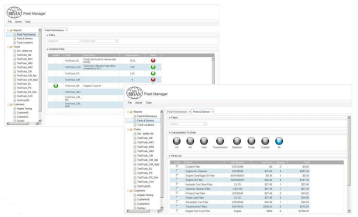

HTTPS / Basic Auth ↑ RESTful / JSON

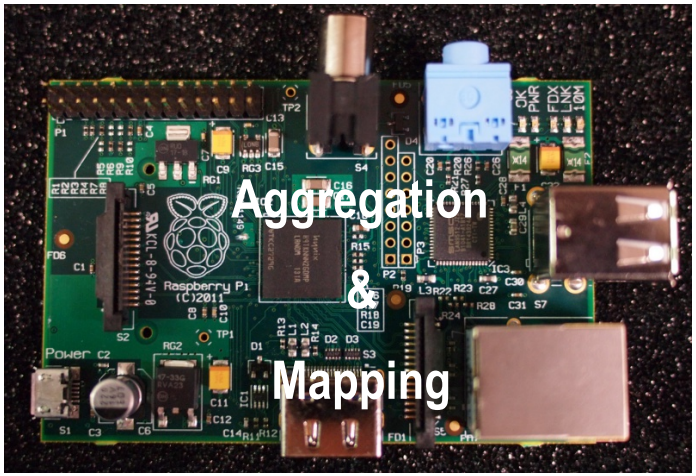

CIP

PCCC

TAIP

(GPS)

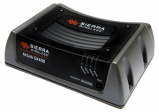

Cloud Gateway

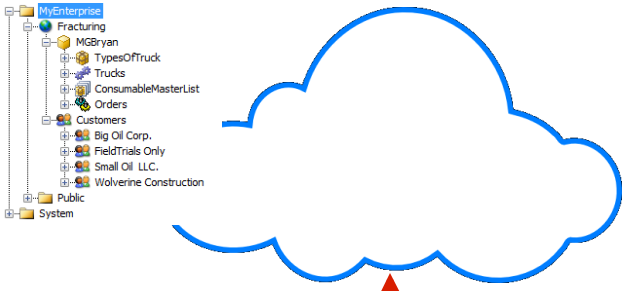

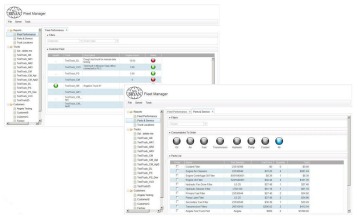

HTTPS / Basic Auth

RESTful / JSON

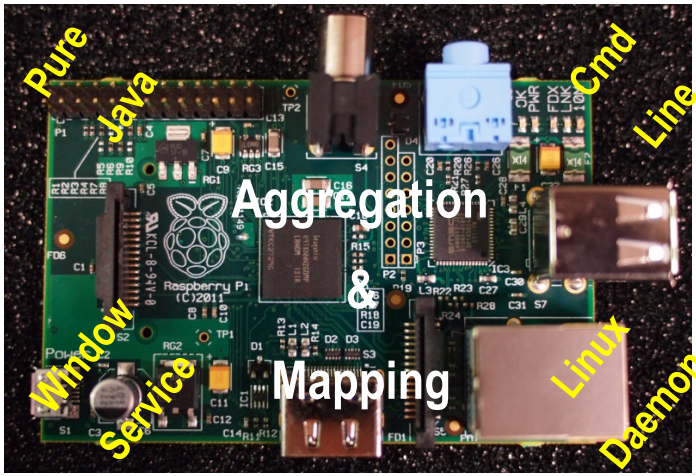

CIP

PCCC

TAIP

(GPS)

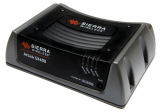

Provisioning the Cloud Gateway

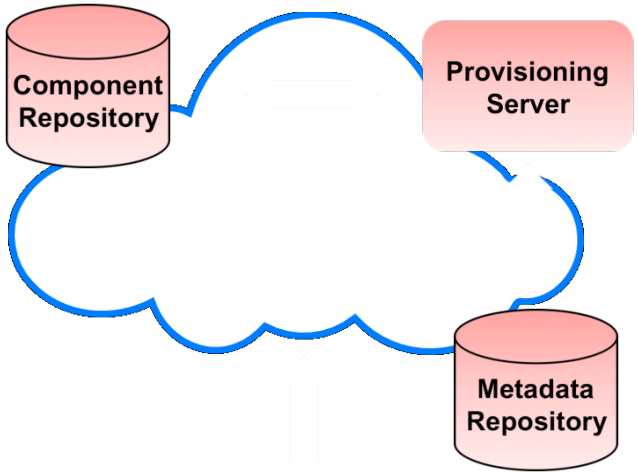

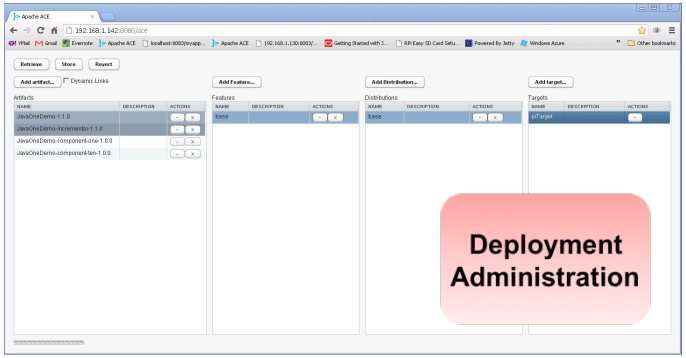

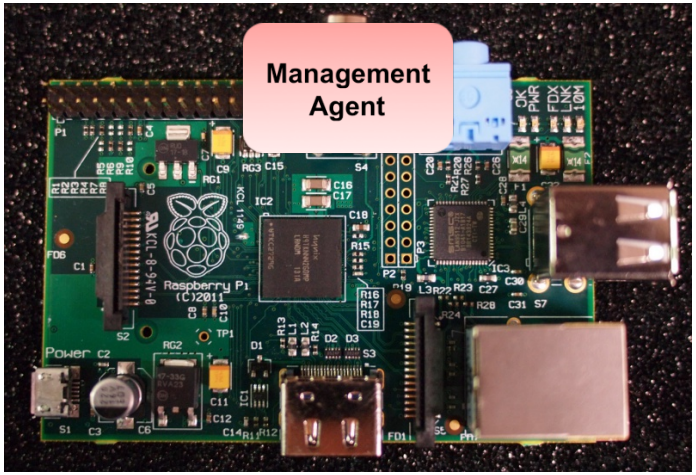

Apache Ace

- Management Agent
  - Identification, Discovery, Scheduler, Deployment, Audit Logs
- Deployment Administration
  - Versioned, Transactional, Deltas, Digital Signatures, Extensible
- Provisioning Server
  - Maps components to targets
  - Only stores metadata
- Component Repository
  - OBR, Maven, anything reachable by URL
  - Can be deployed separately from the server
  - May be replicated

Challenges - Performance

Startup Time

- How to speed up OSGi initialization?
  - Use bundle cache
  - Deferred startup
  - Concurrent startup: bundle activator threads
  - Strict bundle loading faster than dynamic
- Jar consolidation

- Runtime
  - Deterministic Response Time
    - Impeded by periodic processing or DB transaction log
    - Impacted by various DB housekeeping chores
  - Limit use of JNI
- Managing Memory
  - Consider more aggressively releasing heap to OS

Java Cipher Performance on ARM

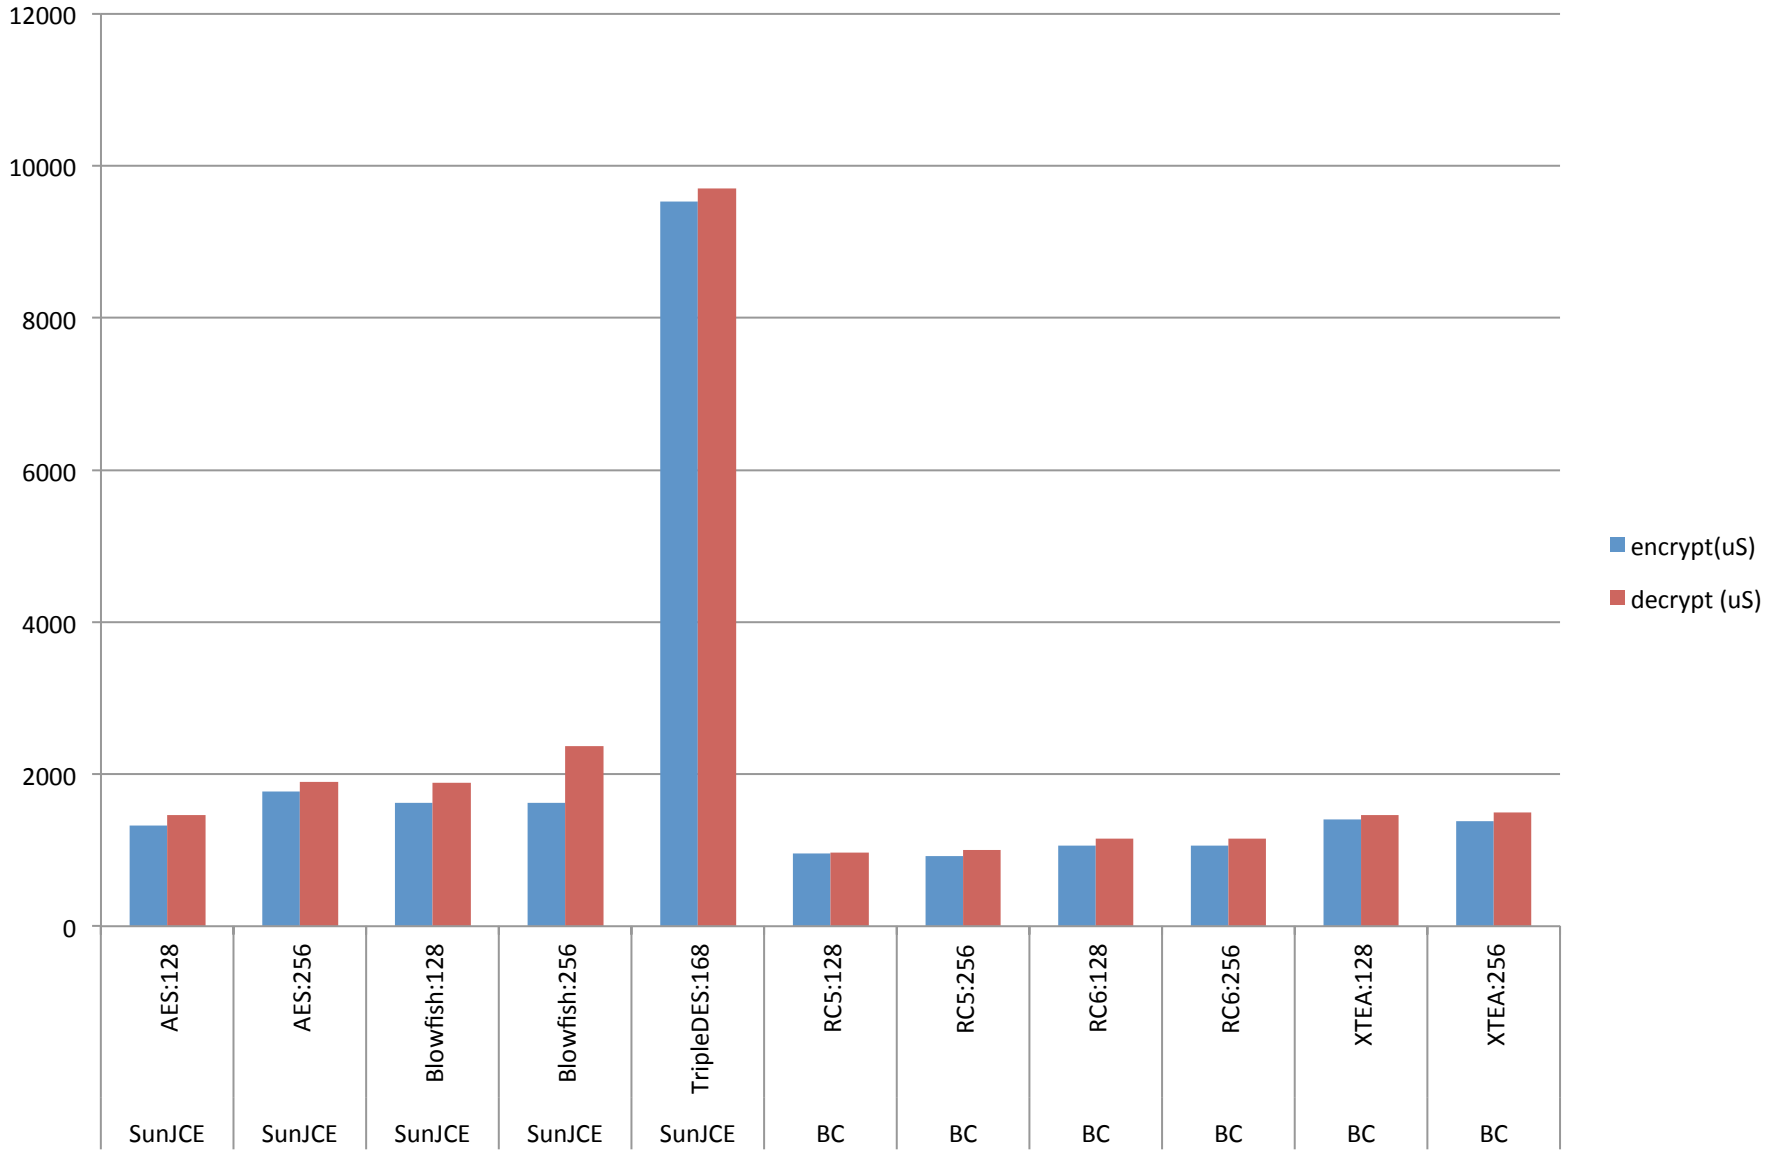

H2 Encryption Performance

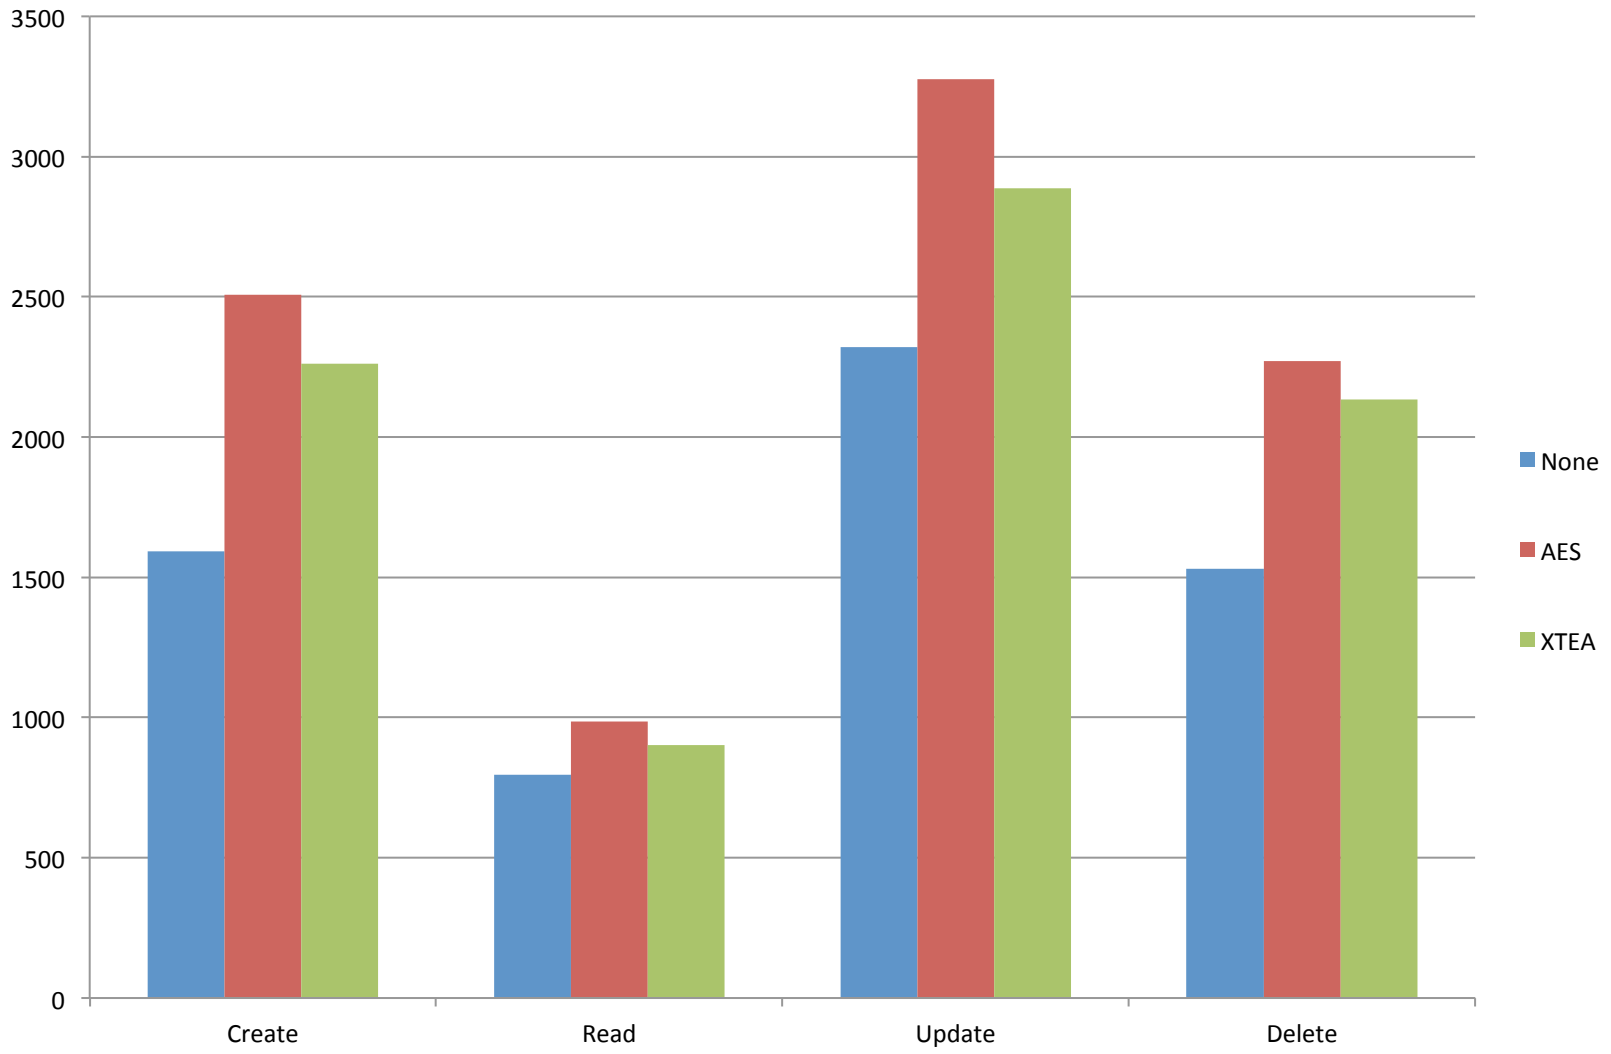

Challenges – Flash Memory

- Automation systems are long lived
- Flash memory wears out: limited # of write/erase cycles
- Write cycles must be considered
- Database
  - updates
  - index generation
  - transaction logs
  - Housekeeping
- Swap
- Temp file systems
- Trace / log data
- Diagnostics

Challenges - Hardening

- How to protect Intellectual Property
- How to make it tamperproof
- How to handle secrets?
  - Device must have access to keys when disconnected
  - Hide in filesystem
  - Hide in code
  - Encrypt on dongle
- Build time techniques
  - Identify critical components
  - Use “Clean room” development
  - Code signing, encryption
  - Obfuscation
    - Challenge for debugging

Code Obfuscation

- Desire to protect code and intellectual property
- Approaches
  - Layout, Data, Control, Encryption
- Consequences
  - Difficult to debug
  - May execute more slowly, particularly with encryption
  - Bytecode manipulation may introduce unintended behavior, as when a dynamically loaded a class
  - Engineering work must be performed to determine which classes must be excluded from obfuscation.
  - Additional build steps may be required to process code.
- ProGuard
  - Flexible, open source obfuscator
  - Also optimizes execution, shrinks jar footprint ~ 1/3
  - Retracing file allows stack traces to be reconstructed

ProGuard jar sizing

reduction (%)	shrink	optimize	obfuscate
8.8	X		
0*		X	
29.6			X
11.7	X	X	
34.9	X		X
36.2	X	X	X

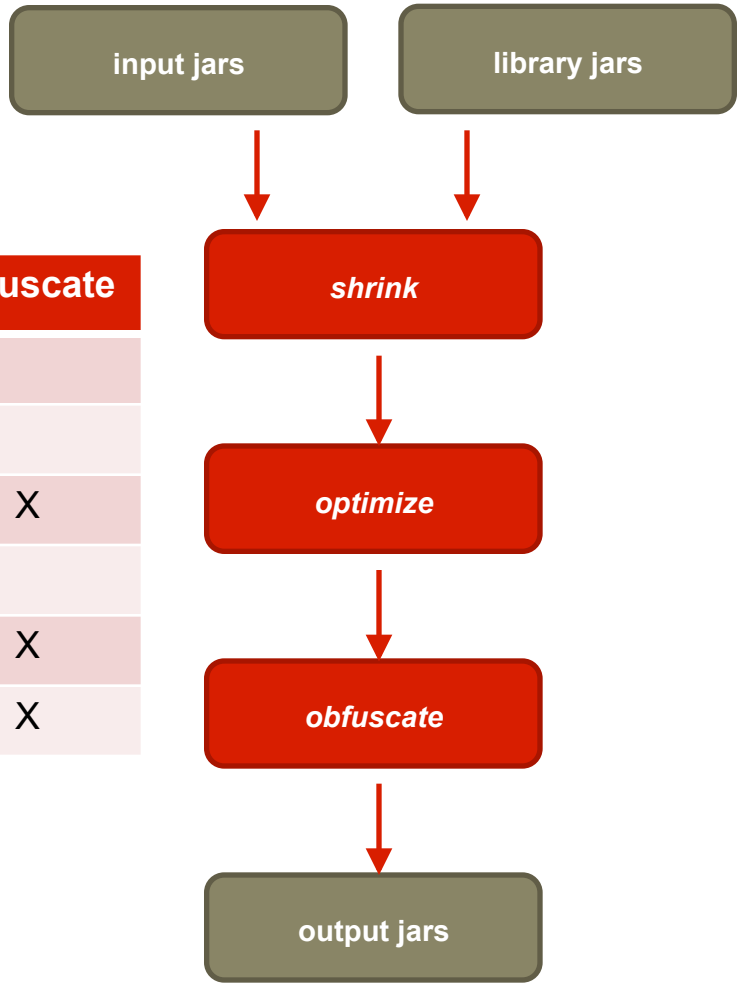

Challenges - Troubleshooting

- Embedded Java provides a JRE not a JDK
- Typical JDK diagnostic tools are unavailable on the target
  - jmap, jps, jstat, etc.
- The JRE can support remote connection via JMX protocol
- Attach VisualVM from remote workstation
  - Note: run with same major Java version and arch
  - Monitor memory & CPU
  - Generate and analyze heapdumps
  - Profiling (sampling)

VisualVM Monitoring

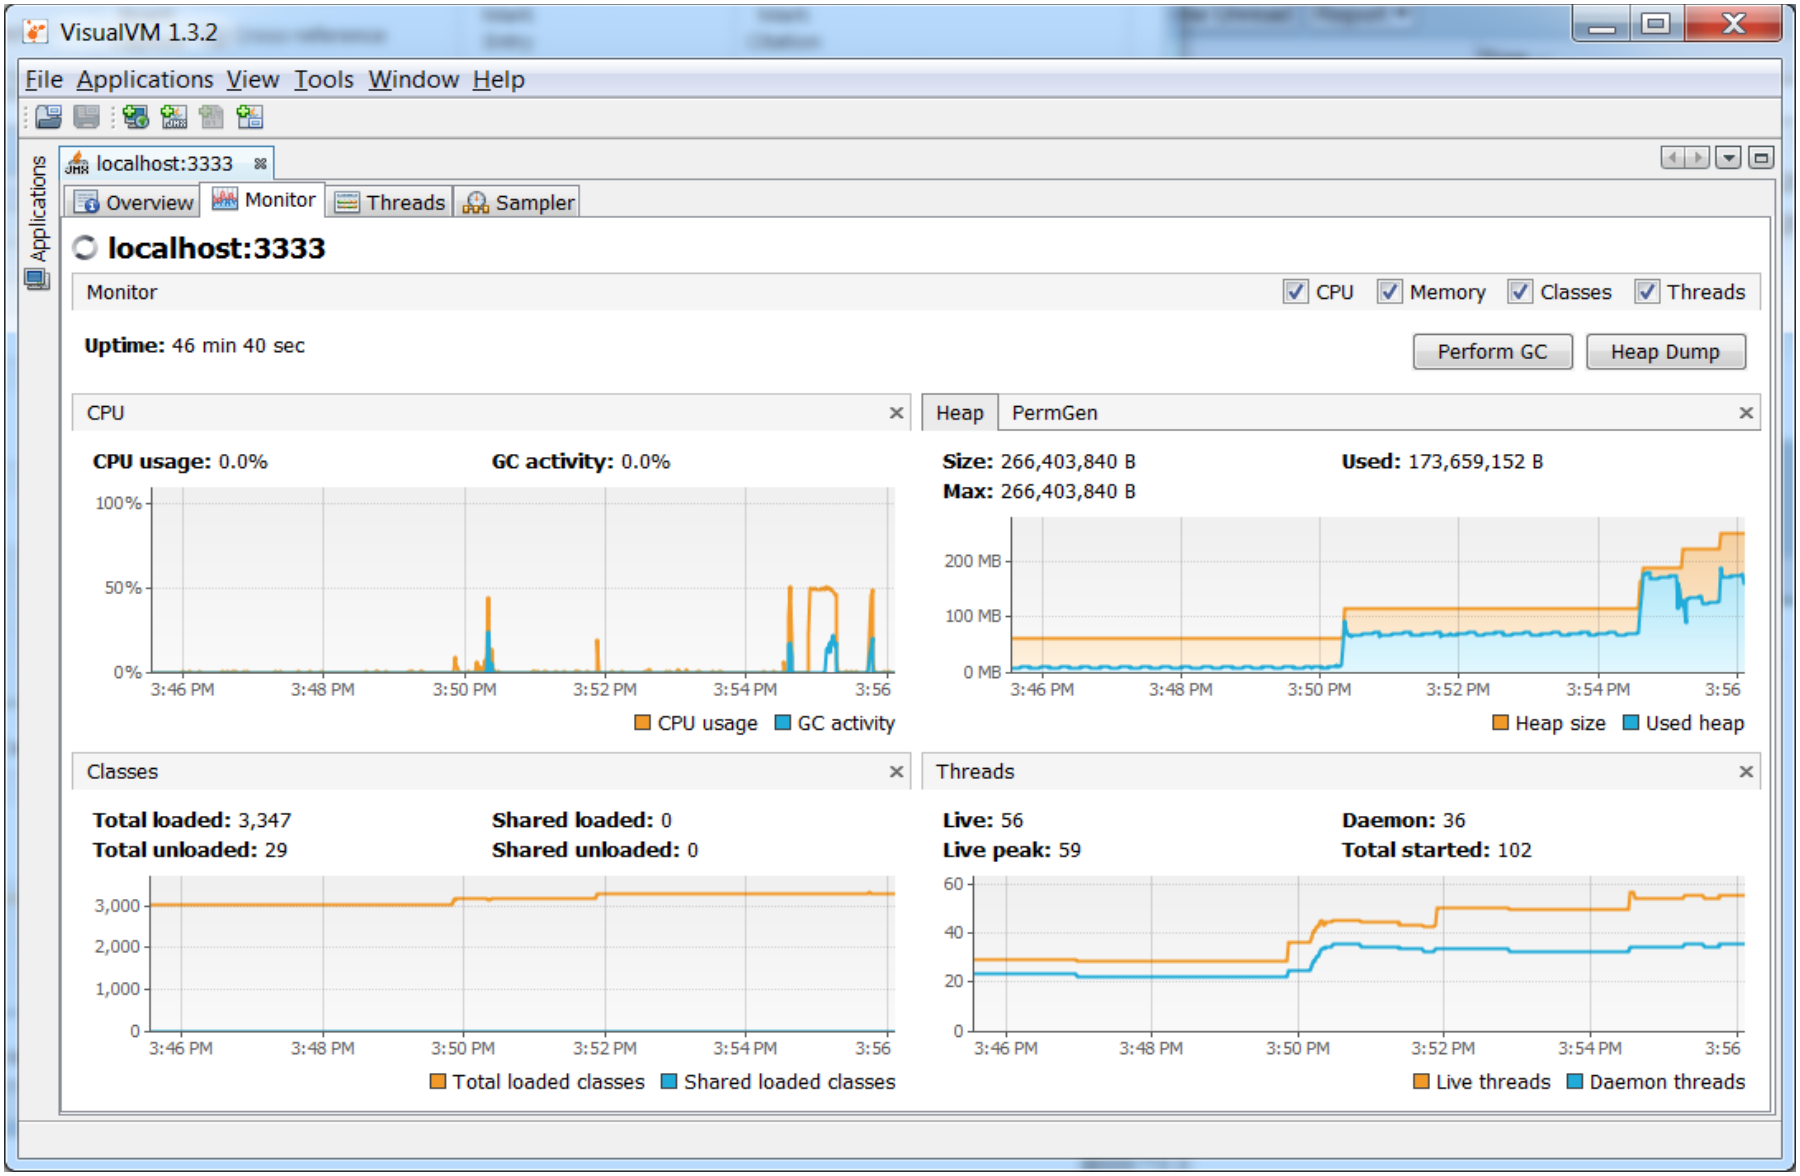

VisualVM Heap Analysis

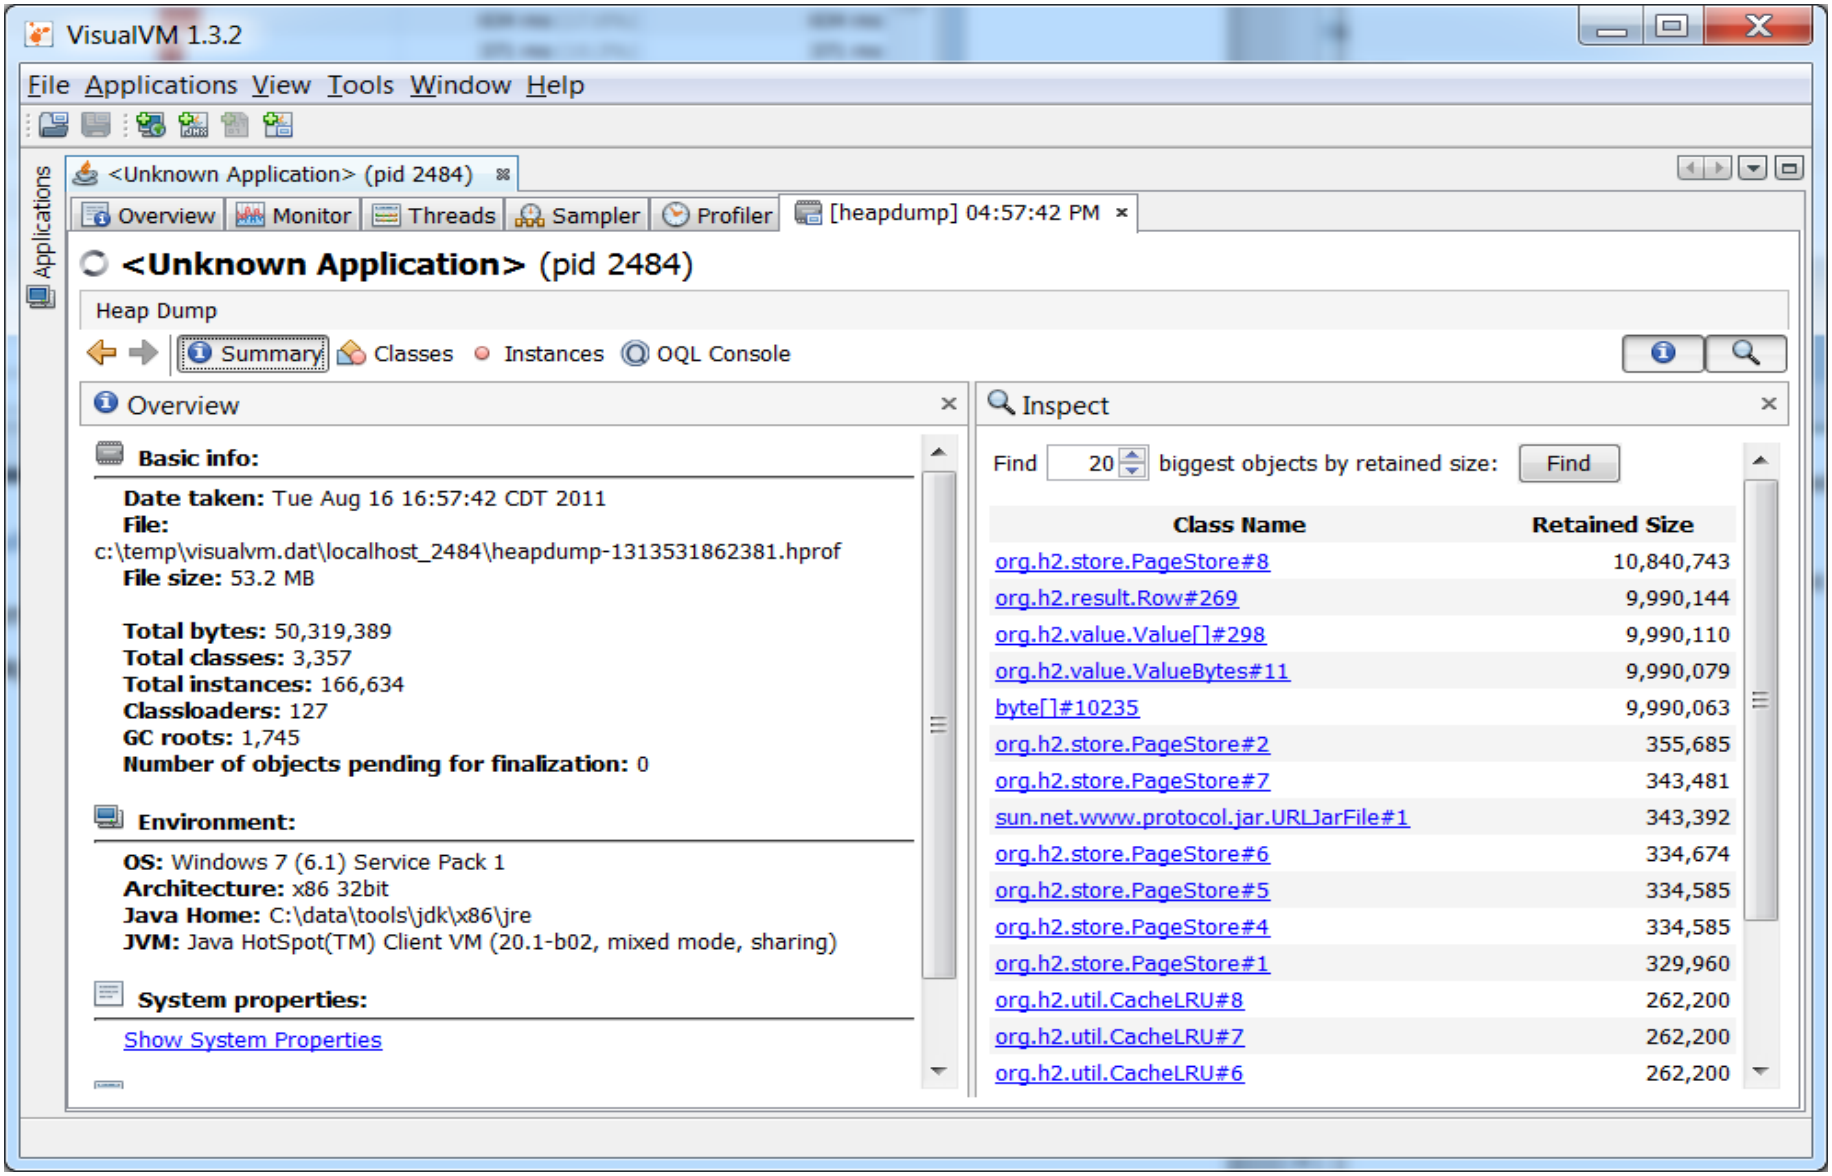

VisualVM Profiling

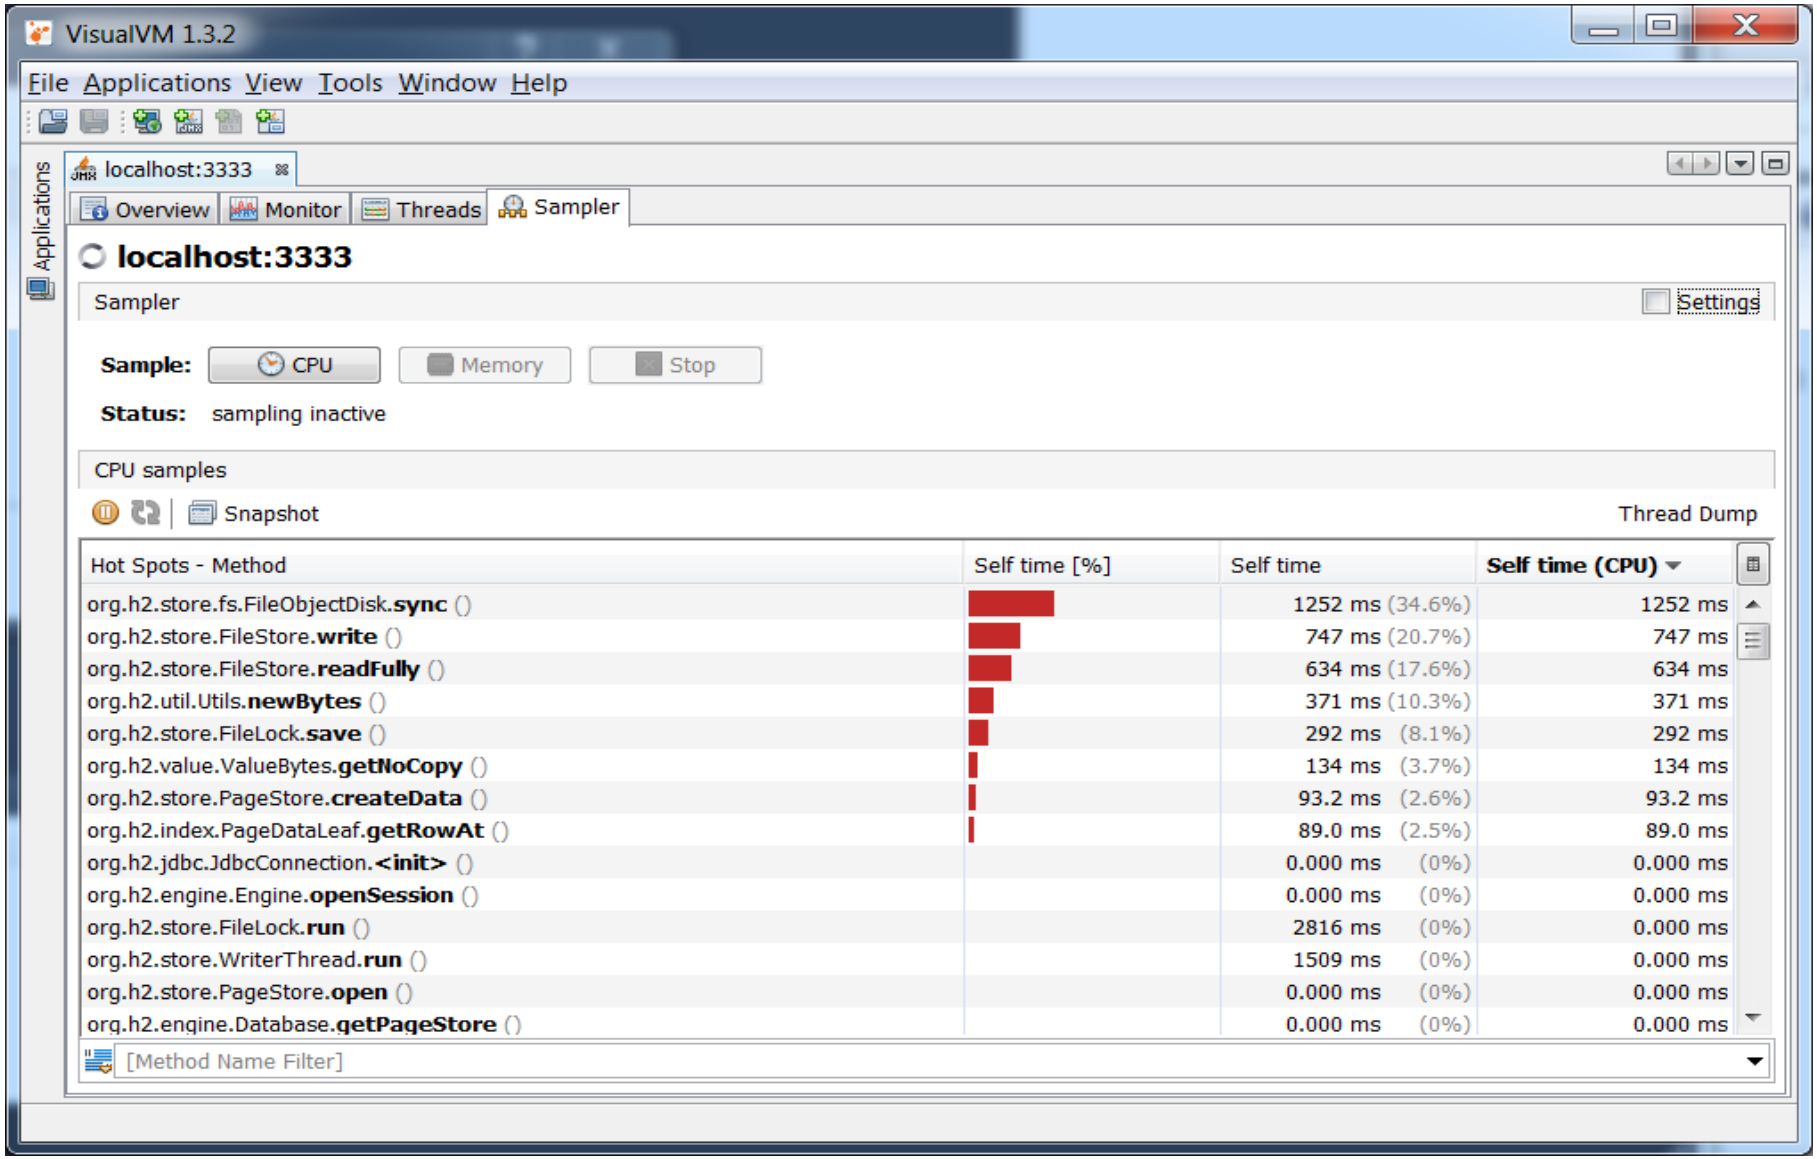

More Troubleshooting

- Development
  - Monitoring: memory leaks, GC
  - Sampling / Profiling: hotspots, latencies
- Production
  - Postmortem: system crashes, hangs
- Lots of troubleshooting tools available
  - JVM cmd line options like -XX:+HeapDumpOnOutOfMemoryError
  - JDK tools like *hprof*, *jmap, jstack* and *jhat* (not on Embedded JRE!)
  - JConsole and VisualVM
  - Native tools like *dtrace* (Solaris) and *strace* (Linux)
- Troubleshooting Guide for Java SE 6
  - http://www.oracle.com/technetwork/java/javase/tsg-vm-149989.pdf

Troubleshooting Scenarios

Memory Utilization	• *-verbose:class* displays classes loaded • *-XX:+HeapDumpOnOutOfMemoryError* generates a heap dump on OOME • stack dump summarizes heap memory usage • heap analysis shows detailed memory usage, potential leaks • monitoring exhibits runtime memory allocation and GC behaviors
Garbage Collection	• *-verbose:gc* displays GC statistics • monitoring exhibits runtime memory allocation and GC behaviors • heap analysis shows large instances counts and GC roots
Poor Performance	• stack dump summarizes memory usage, indicating potential memory utilization issues • monitoring exhibits runtime CPU usage, thread count & activity • monitoring exhibits runtime memory allocation and GC behaviors • stack dump indicates deadlocks
Hanging / Looping	• stack dump indicates busy threads, deadlocks (may need to force dump) • monitoring exhibits runtime CPU usage, thread count & activity
System Crash	• examine application and Fatal Error logs for clues • look in filesystem working dir for stack/heap dumps • restart JVM with monitoring enabled • restart JVM with *-XX:+HeapDumpOnOutOfMemoryError* option • restart JVM with *-XX:OnError* executing script to assist w/ debugging

Q & A