

# Continuous Profiling in Production: What, Why and How



Richard Warburton (@richardwarburto)

Sadiq Jaffer (@sadiqj)

<https://www.opsian.com>

# Why Performance Tools Matter

Development isn't Production

Profiling vs Monitoring

Continuous Profiling

Conclusion

# Known Knowns

# Known Unknowns

# Unknown Unknowns

# Why Performance Tools Matter

## Development isn't Production

### Profiling vs Monitoring

### Continuous Profiling

### Conclusion

# Development isn't Production

Performance testing in development can be easier

May not have access to production

Tooling often desktop-based

Not representative of production

# Unrepresentative Hardware




VS






# Unrepresentative Software



# Unrepresentative Workloads



VS



# The JVM may have very different behaviour in production

Hotspot does adaptive optimisation

Production may optimise differently

**I have no idea**



**what I'm doing**

Why Performance Tools Matter

Development isn't Production

Profiling vs Monitoring

Continuous Profiling

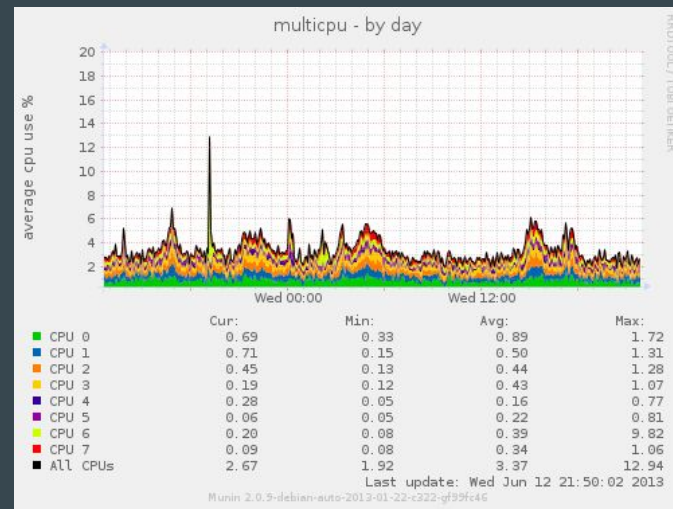
Conclusion

# Ambient/Passive/System Metrics

Preconfigured Numerical Measure

CPU Time Usage / Page-load Times

Cheap and sometimes effective





# Logging

Records arbitrary events emitted by the system being monitored

log4j/slf4j/logback

Logs of GC events

Often manual, aids system understanding, expensive



# Coarse Grained Instrumentation

Measures time within some instrumented section of the code

Time spent inside the controller layer of your web-app or performing SQL queries

More detailed and actionable though expensive

# Production Profiling

What methods use up CPU time?

What lines of code allocate the most objects?

Where are your CPU Cache misses coming from?

Automatic, can be cheap but often isn't

# Where Instrumentation can be blind in the Real World

Problem: every 5 seconds an HTTP endpoint would be really slow.

Instrumentation: on the servlet request, didn't even show the pause!

Cause: Tomcat expired its resources cache every 5 seconds, on load one resource scanned the entire classpath



# Surely a better way?

Not just Metrics - Actionable Insights

Diagnostics aren't Diagnosis

What about Profiling?

# Why Performance Tools Matter

## Development isn't Production

## Profiling vs Monitoring

## Continuous Profiling

## Conclusion

# How to use Continuous Profilers

- 1) Extract relevant time period and apps/machines
- 2) Choose a type of profile: CPU Time/Wallclock Time/Memory
- 3) View results to tell you what the dominant consumer of a resource is
- 4) Fix biggest bottleneck
- 5) Deploy / Iterate



# CPU Time vs Wallclock Time





# You need both CPU Time and Wallclock Time

CPU - Diagnose expensive computational hotspots and inefficient algorithms

Spot code that should not be executing but is ...

Wallclock - Diagnose blocking that stops CPU usage

e.g blocking on external IO and lock contention issues

# Profiling Hotspots

Packages to exclude: com.sun.org.foo.bar

Method	Calls	Time	
<a href="#">java.util.zip.ZipFile.getEntry</a>	589 (8%)	84,227ms	
<a href="#">java.io.UnixFileSystem.getBooleanAttributes0</a>	419 (5%)	59,917ms	
<a href="#">java.security.AccessController.doPrivileged</a>	135 (1%)	19,305ms	
<a href="#">sun.misc.Unsafe.park</a>	119 (1%)	17,017ms	
<a href="#">sun.nio.ch.EPollArrayWrapper.interrupt</a>	108 (1%)	15,444ms	
<a href="#">java.lang.String.charAt</a>	103 (1%)	14,729ms	

## Hot Spot line information

String.java:657

String.java:660

## Samples

64 (62%)

39 (38%)

# Profiling Treeviews

Flame Graph

**Tree View**


Hot Spots

**i** 11% of your time was spent in the JVM [Learn more.](#)

Select thread group..

Thread Group	Method	Time share	Self samples	Total samples
http-nio-?-exec-?	- java.lang.Thread.run:748		0 (0%)	6384 (90%)
	- org.apache.tomcat.util.threads.TaskThread\$WrappingRunnable.run:61		0 (0%)	6384 (90%)
	- java.util.concurrent.ThreadPoolExecutor\$Worker.run:624		0 (0%)	6384 (90%)
	- java.util.concurrent.ThreadPoolExecutor.runWorker:1149		1 (0%)	6339 (89%)
	- org.apache.tomcat.util.net.SocketProcessorBase.run:49		0 (0%)	6336 (89%)

# Profiling Flamegraphs



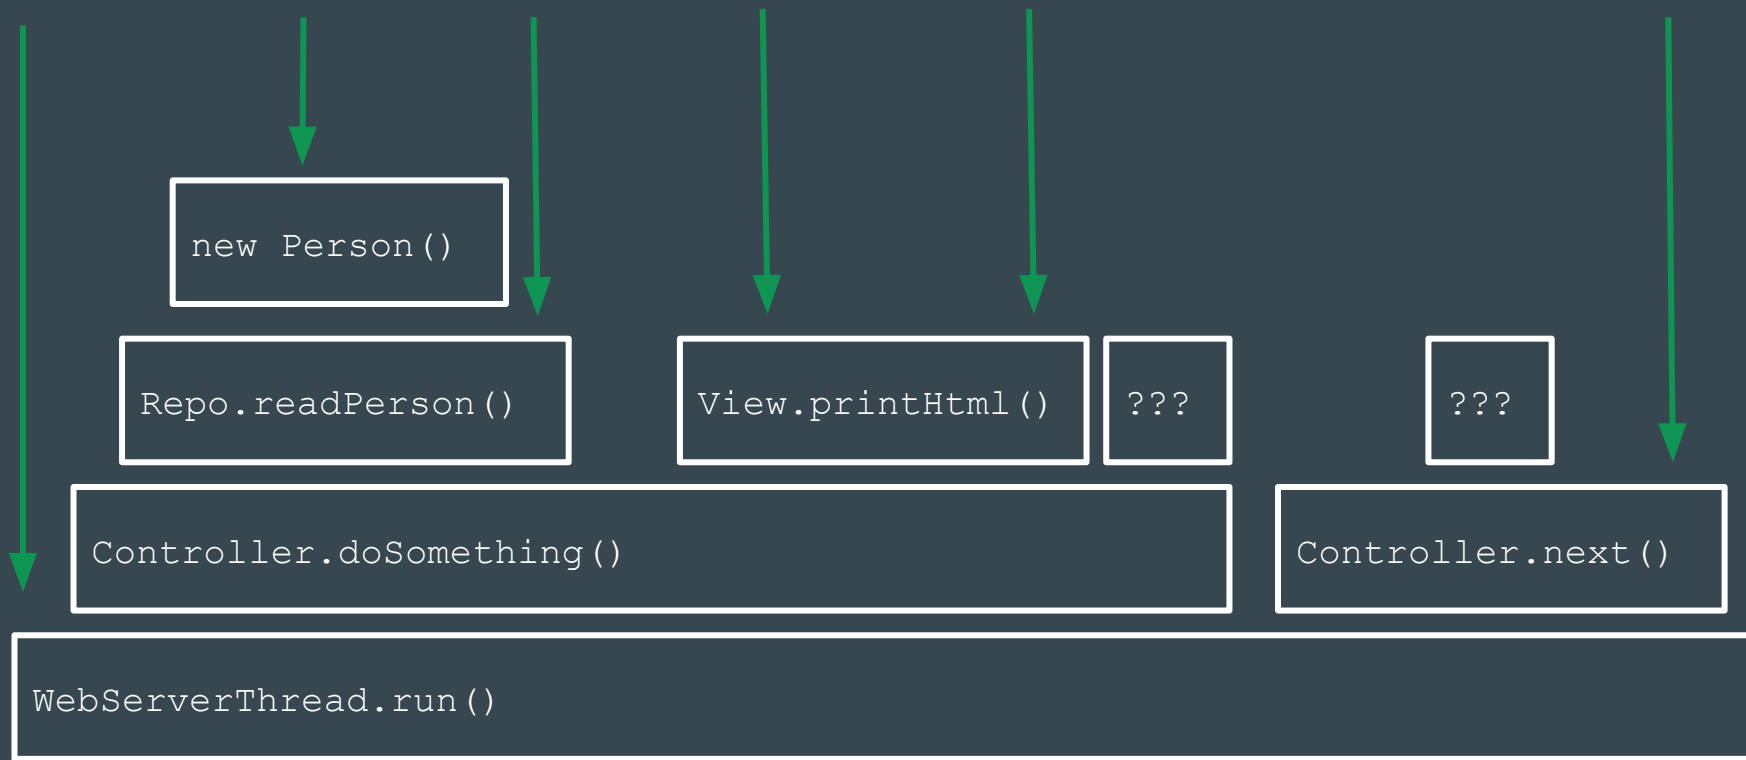
# Instrumenting Profilers

Add instructions to collect timings (Eg: JVisualVM Profiler)

Inaccurate - modifies the behaviour of the program

High Overhead - > 2x slower

# Sampling/Statistical Profilers



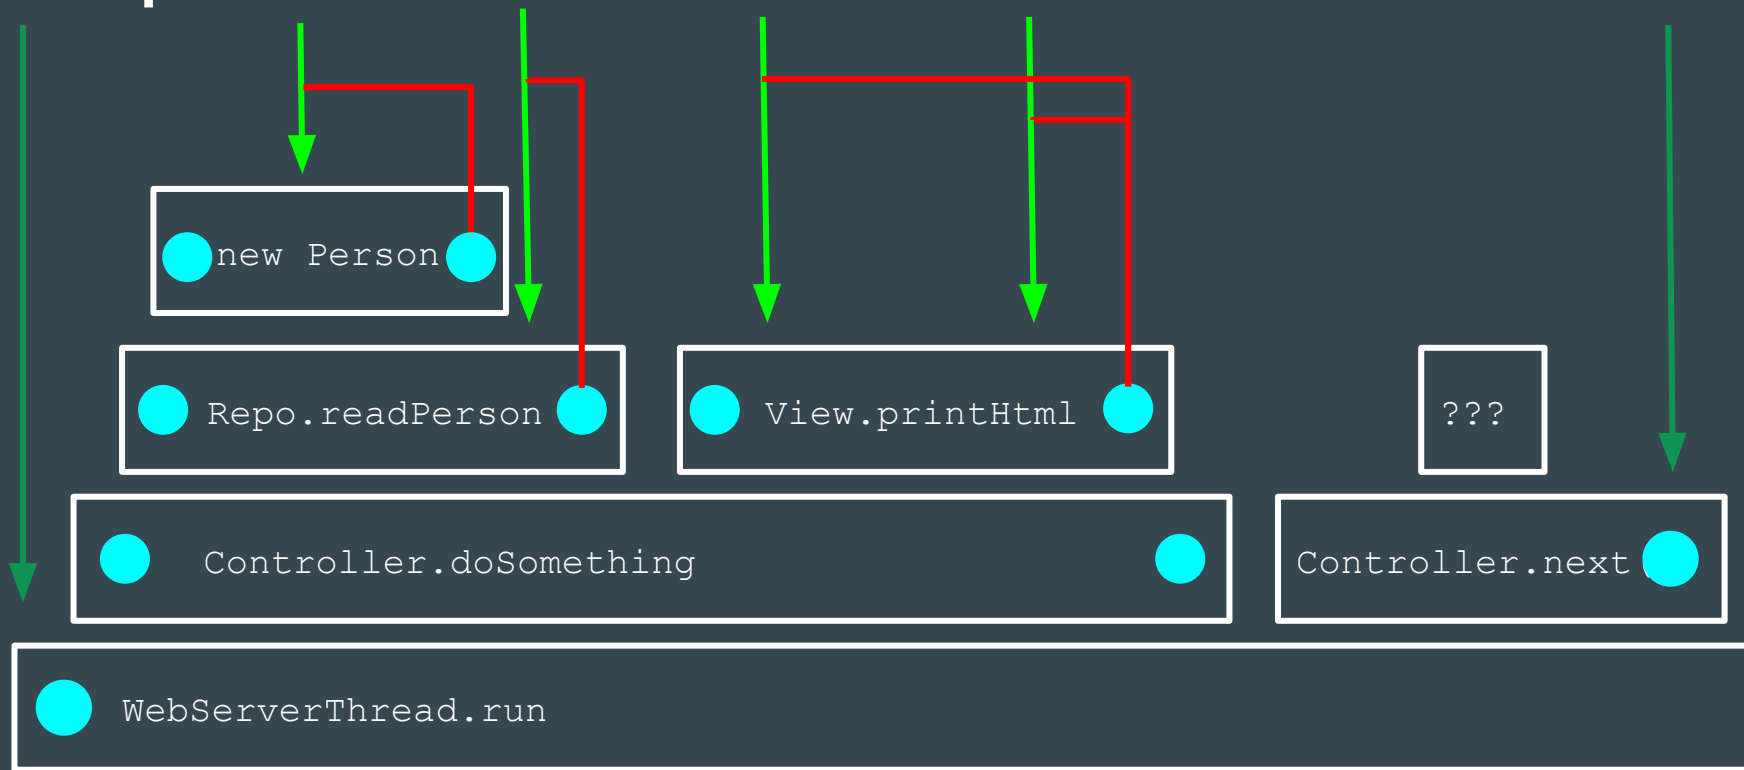
# Safepoints

Mechanism for bringing Java application threads to a halt

Safepoint *polls* added to compiled code read known memory location

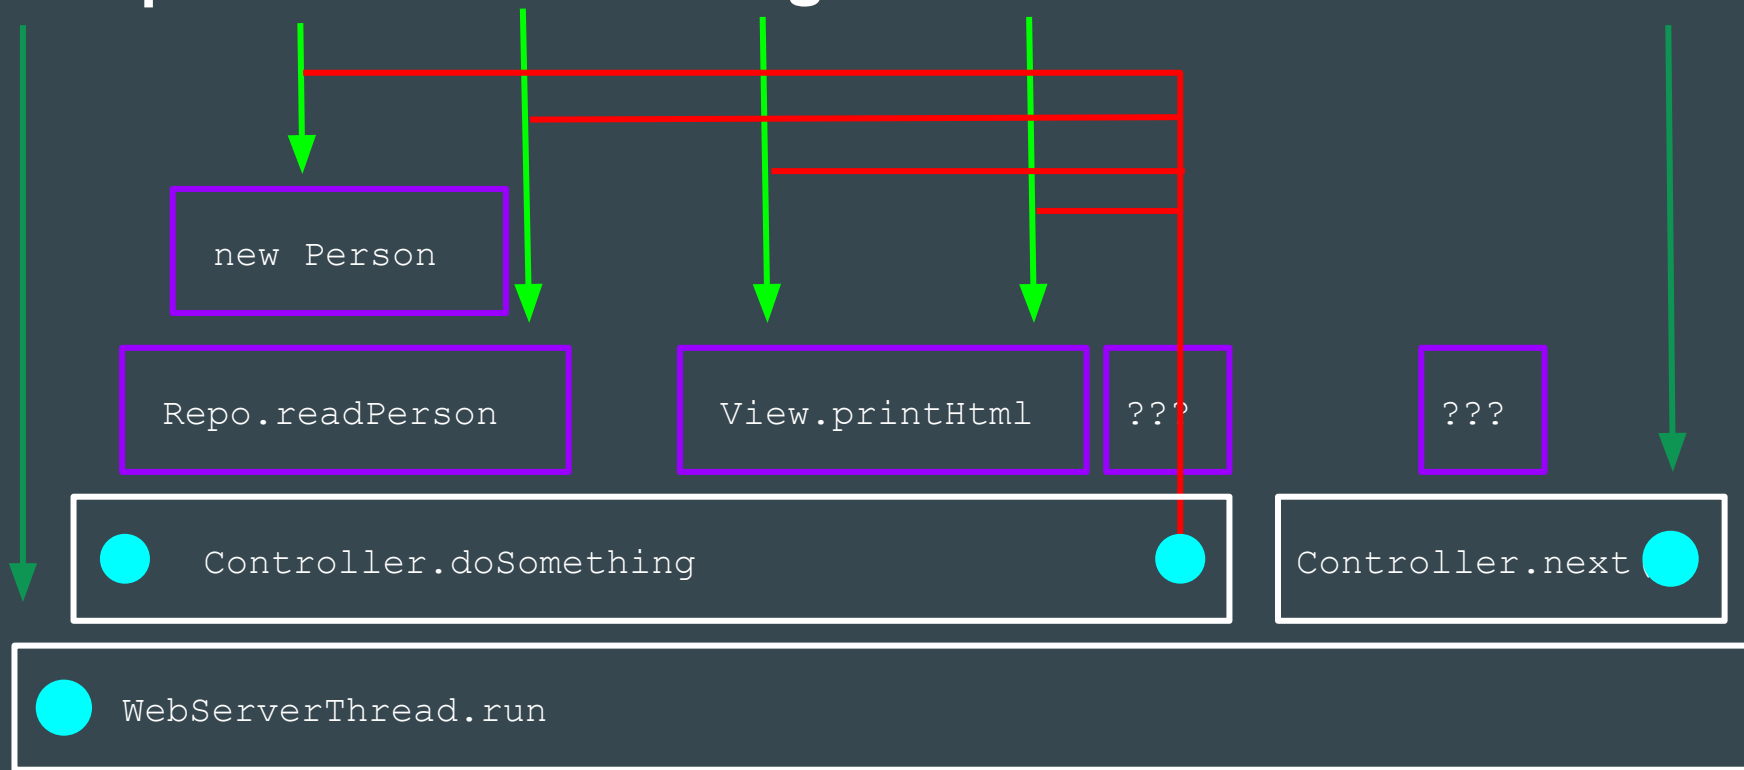
Protecting memory page triggers a segfault and suspends threads

# Safepoint Bias

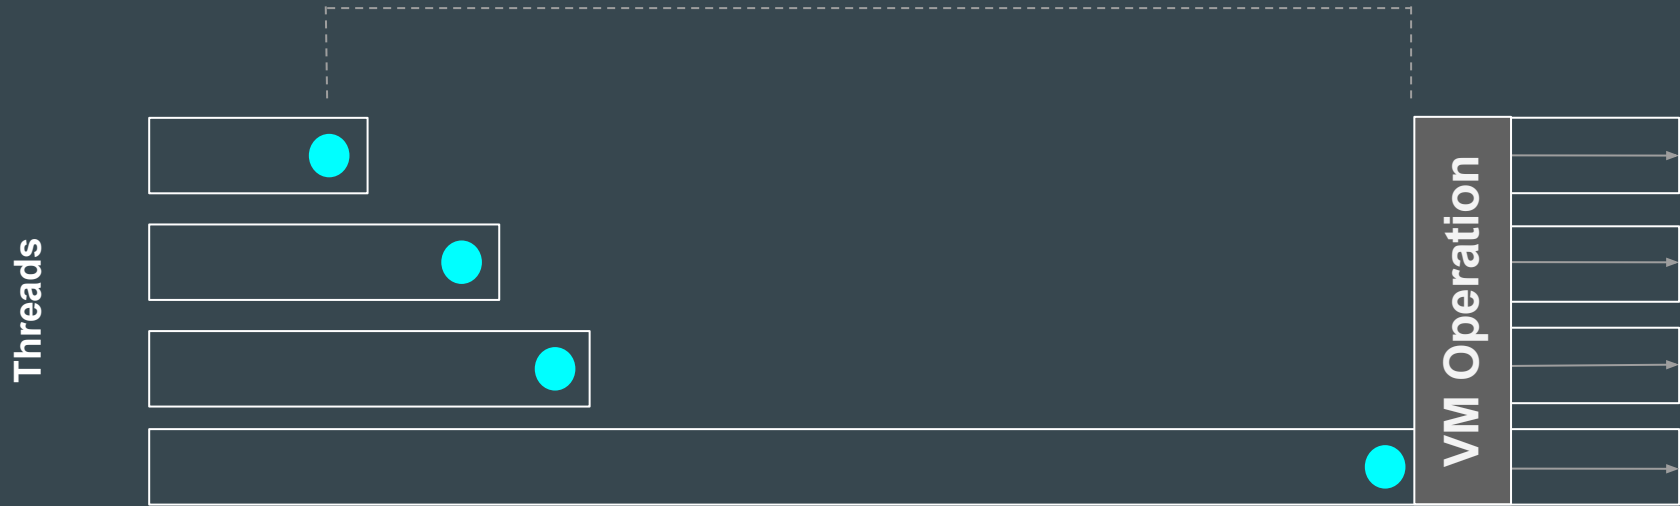




# Safepoint Bias after Inlining



# Time to Safepoint



-XX:+PrintSafepointStatistics

● Safepoint poll

# Statistical Profiling in Java

Problem: *getAllStackTraces* is expensive to do frequently and inaccurate, also only gives us Wallclock time

Need ways to:

1. Interrupt application
2. Sample resource of interest

# Advanced Statistical Profiling in Java

- Interrupt with OS signals
  - Delivered to handler on only one thread
  - Lightweight
- Sample resource of interest
  - Use *AsyncGetCallTrace* to sample stack
  - Examine JVM internals for other resources

# Advanced Statistical Profiling in Java

Approach not used by existing profilers (VisualVM and desktop commercial alternatives)

Can give very low overheads (<1%) for reasonable sampling rates

**People are put off by practical as  
much as technical issues**

# Barriers to Ad-Hoc Production Profiling

Generally requires access to production

Process involves manual work - hard to automate

Low-overhead open source profilers without commercial support



What if we profiled all the time?



# Historical Data

Allows for post-hoc incident analysis

Enables correlation with other data/metrics

Performance regression analysis

# Putting Samples in Context

Application version

Environment parameters (machine type, CPU, location, etc.)

Ad-hoc profiling we can't do this

# How to implement Continuous Profiling

# Google-wide profiling

Article: *Google-Wide Profiling: A Continuous Profiling Infrastructure for Data Centers*

Profiling data and binaries collected, processed and made available for browser-based reporting

“The system has been actively profiling nearly all machines at Google for several years”

<https://ai.google/research/pubs/pub36575>

# Self-build

- Open source Java profilers suitable for production
  - Async-profiler
  - Honest profiler
  - Flight Recorder
- Need to collect and store profiles in a database
- Tools for retrieving and visualising stored profiling data
  - Browser-based
  - Command line

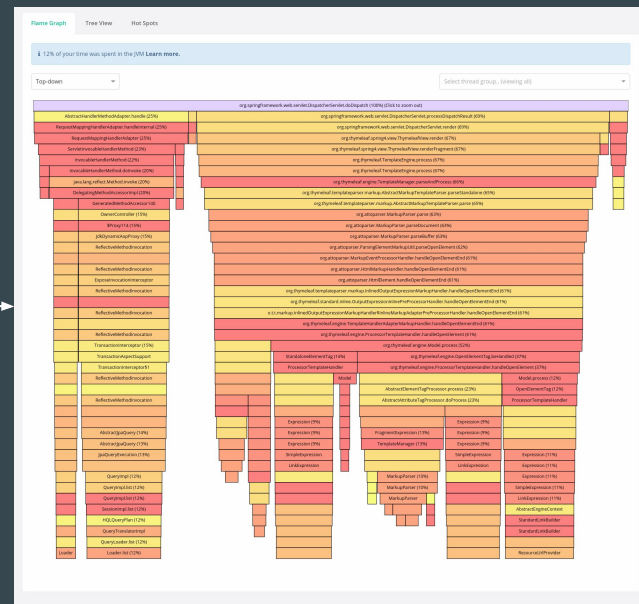
# Opsian - Continuous Profiling

## JVM Agents



**Opsian  
Aggregation  
service**

## Web Reports



# Summary

It's possible to profile in production with low overhead

To overcome practical issues we can profile production all the time

We gain new capabilities by profiling all the time

Why Performance Tools Matter

Development isn't Production

Profiling vs Monitoring

Continuous Profiling

Conclusion



Performance Matters

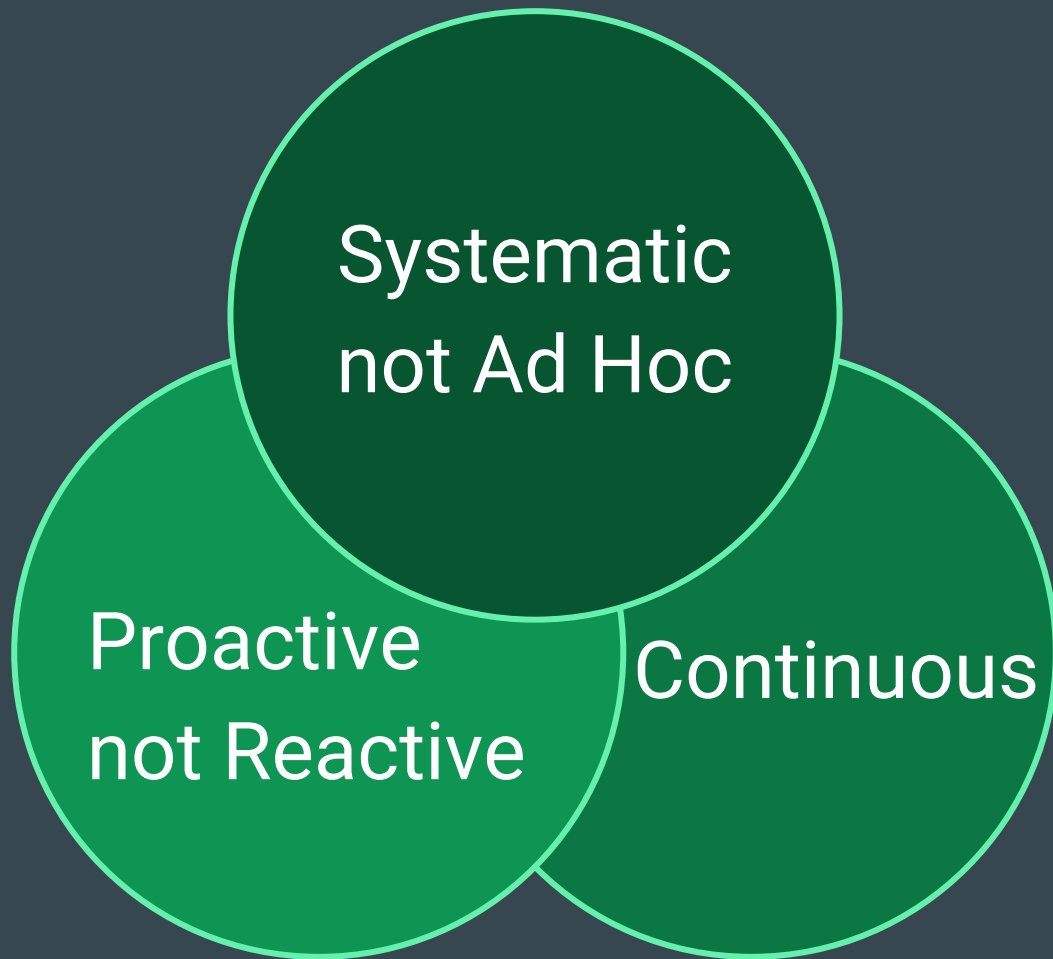
Development isn't Production

Metrics can be unactionable

Instrumentation has high overhead

Continuous Profiling provides insight

**We need an attitude shift on profiling  
+ monitoring**



**Please do Production Profiling.  
All the time.**

**Any Questions?**  
<https://www.opsian.com/>

# Live Demo?

# Links

[Collector - Flame Graph](#)

[Collector - Hot Spots](#)

# The End



# Existing tools are blind

Traditional profilers don't work in production

Metrics aren't code level visibility

Instrumentation must be done ahead of time

# How do we help?

Reduce the risk of change

Help you scale with happy customers

Cut the cost of infrastructure

# Production Visibility

Actionable reports for causes of latency and CPU usage

From high-level reports to line-level granularity


Very low overhead (<1%) and always-on

# Reduce the risk of change

On-demand performance comparison between releases

Accelerate root-cause analysis for performance regressions

# Improve Developer Productivity



Source: ZT RebelLabs Developer Productivity Report 2017

# Understand don't Overwhelm

## Too Little

You can't understand production problems

## Too Much

Needle in a Haystack

You are the problem (overhead)

# Normalisation of Deviance

“Some of the tests always fail, so we just ignore them.”

“Some of the alerts get triggered regularly, so we just ignore them.”

Alert false positives have a cost

# Open Source Java Profilers

## High Overhead

VisualVM

hprof

Twitter's CPUProfile

Anything *GetAllStackTraces* based

## Low Overhead

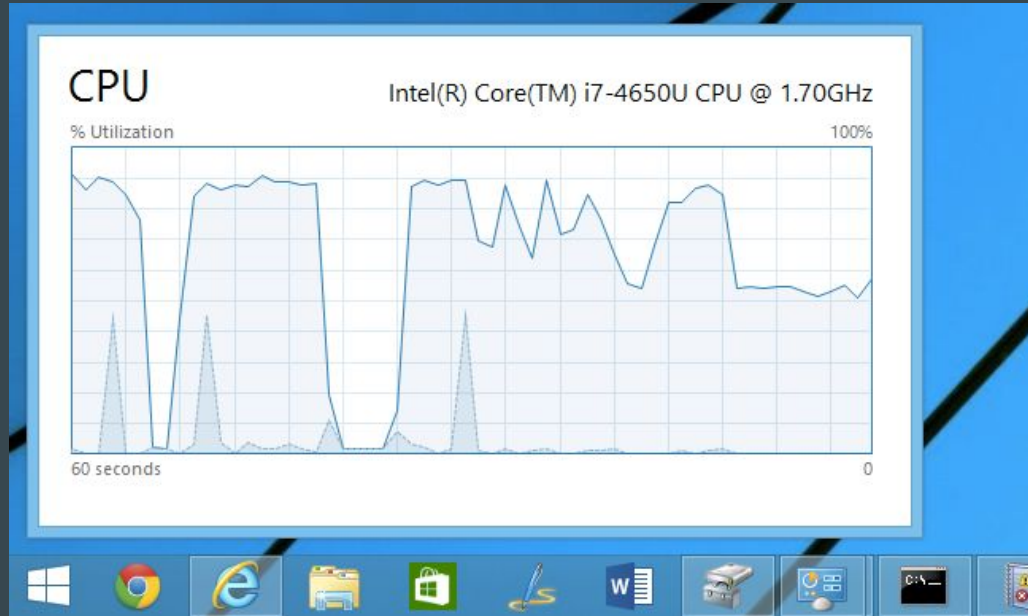
Async Profiler

Honest Profiler

Java Mission Control



# Unactionable Metrics



Many metrics provide pretty graphs but don't progress treatment

# Profiling Support in the Linux Kernel

perf and eBPF

perf-map-agent for the JVM

Hardware events (L1/L2/L3 cache misses, branch mispredictions, etc.)

Take heed: potential security issues

# Customer Experience



# Responsive Applications make more Money



Amazon: 100ms of latency costs 1% of sales

Google: 500ms seconds in search page generation time drops traffic by 20%

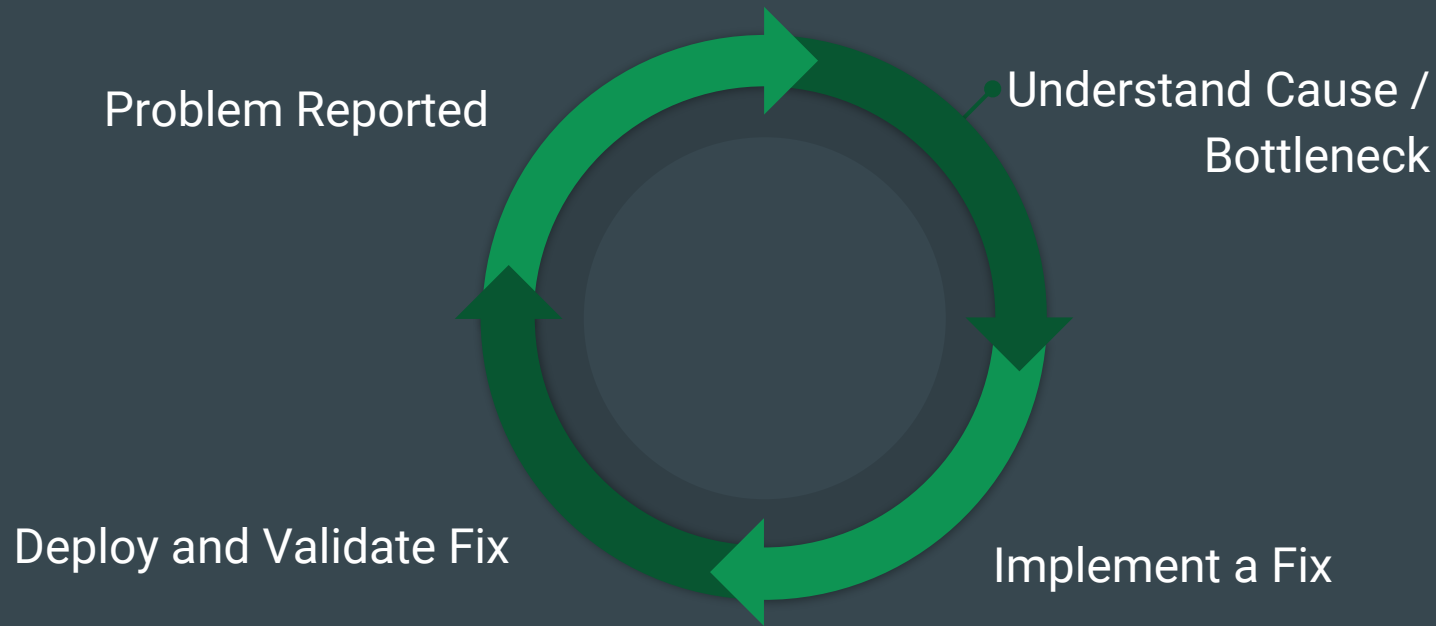
# Stop Costly Downtime



# Reduce Costs



# Performance Optimisation Cycle



# What's Hard?





**How do you find performance  
bottlenecks?**

