A decorative graphic on the left side of the slide consists of numerous circles of varying sizes and colors, including green, cyan, blue, and purple. The circles are arranged in a pattern that suggests movement or a trail, with some circles having soft shadows behind them.

The Tracing and Monitoring Framework in 2012

or Why Research Associates are Awesome


Matthew Khouzam & Alexandre Montplaisir

What will we see today

- › Blatant self-plugs
- › Part I:
 - Review
 - View Recap
 - New Developments
 - Creating your own view
- › Part II:
 - Statistics in state
 - Multiple state systems
 - Histogram from state
- › Part III: How to contribute
 - Making a small patch
 - Pushing it to Gerrit for review

It looks like you're trying to view a trace. Would you like me to open it for you?

- I'll use eclipse, thank you
- 1997 called, they want you back



imageGenerator.net

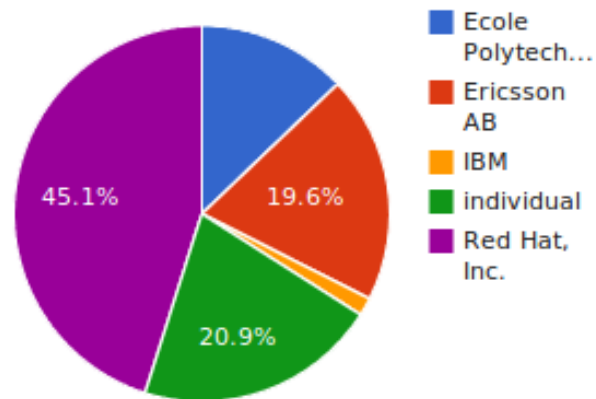
Who are we?

- › Alexandre:
 - King of charts
 - Java enthusiast
- › Research associates working for Polytechnique
- › Collaborating with Ericsson
- › Eclipse Linux Tools committers
- › Matthew:
 - Spice pusher
 - Likes profilers a bit too much
- › Working with the TOSIDE team to improve the trace viewer

Poly gets some visibility at Eclipse!

Organization Commit Activity


Commits on this project by supporting organization over the last three months.



*Chart is for the whole Linux Tools project
TMF = blue + red*

Part I: What is Eclipse? TMF?

- › Eclipse is an IDE... please tell me you know Eclipse :)
- › TMF is the Tracing and Monitoring Framework
- › Eclipse-LTTng is the trace viewer for LTTng 2.0 traces
- › Download at <http://www.eclipse.org>
- › Git: <http://git.eclipse.org/c/www.eclipse.org/linuxtools.git/>



When we met last time

- › LTTng viewer had many tools and views:
 - Detailed events
 - Filters
 - Highlighting
 - Bookmarking
 - Histogram (Events density)
 - Statistics
 - Tracer Control
 - Control flow
 - Resource
 - Trace organization by experiments
 - Trace project navigation
 - Colors
 - Environment Variables
 - Time-graph
 - Text dump in examples
 - Any Questions?

New features

- › Event “editors”, a.k.a. multi-tabbed view
 - Time synchronization between traces for comparison purposes
- › New human-readable timestamps
- › CTF 1.9 support
- › Multiple state system support
- › State system now drives statistics
- › State system now drives histogram
- › Bug fixes

Event Editor

- › Changes the paradigm of the LTTng viewer.
- › All views reflect the currently selected event editor
- › Experiments are now traces with children, simpler.

Human-Readable Timestamps

- › Timestamp format in one place (preferences)
- › Histogram input no longer in ns

CTF 1.9 support

- › Support for callsites added
- › Improved test coverage and code clarity

TMF General Demo

Let's make a view!

- › CPU usage per thread
- › Requires a state system input and a view

State system input

- › Only handle “sched_switch” event in the kernel
 - If a pid is switched in, write it as “active”
 - If a pid is switched out, write it as “idle”
 - If the pid is 0, ignore it (swapper)
 - Store how many cpus we have

Viewer

- › Display using a stacked bar chart
- › Cannot lose details so it will be slower than CFV
- › Show total time
- › Show Top 10 threads at the time unless threads consume less than 1% of CPU
- › Try not to freeze UI

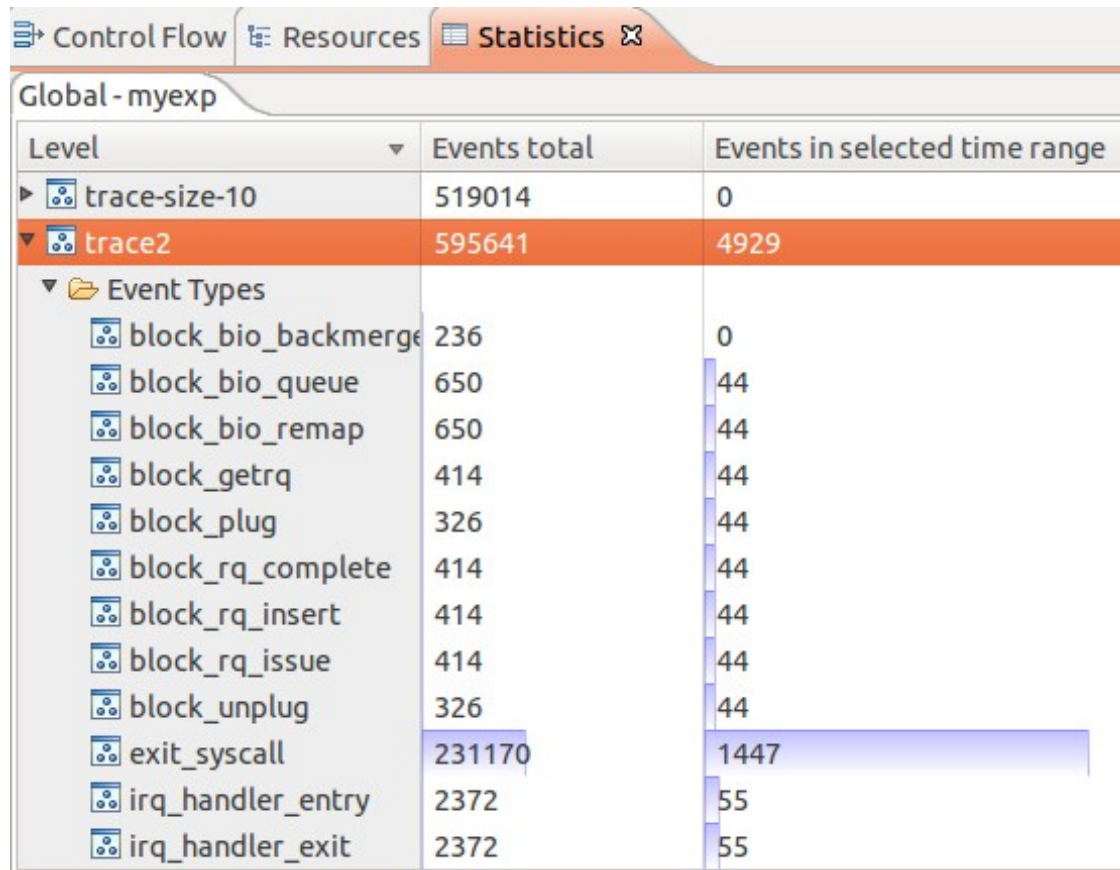
Demo of the new view

Part II: From research project to production

- › Finished my master's one year ago
- › Back then, “Future work” items were:
 - Statistics in the state
 - Hybrid storage
 - Partial history

Statistics in the state

- › “Time range” statistics were now available, but slow.



Level	Events total	Events in selected time range
▶ trace-size-10	519014	0
▼ trace2	595641	4929
▼ Event Types		
block_bio_backmerge	236	0
block_bio_queue	650	44
block_bio_remap	650	44
block_getrq	414	44
block_plug	326	44
block_rq_complete	414	44
block_rq_insert	414	44
block_rq_issue	414	44
block_unplug	326	44
exit_syscall	231170	1447
irq_handler_entry	2372	55
irq_handler_exit	2372	55

Statistics in the state

- › Initial work by Mathieu Denis
- › We use one attribute per event type, incremented every event.
- › When doing a range query:
 - Query the value at the start and end
 - Difference between the two values = number of events in that range
- › $O(1)$ with regards to the length of the interval!
- › $O(\log n)$ with regards to the size of the trace.

Multiple State System Support

- › Support for several state systems in parallel
- › Each trace type can define its own state system
 - TmfTrace (base class) defines a state system for statistics
 - CtfKernelTrace (specific) defines a state system for kernel traces
- › This is different than the “hybrid storage” that was first envisioned, but fills the same needs.

Next contender : Histogram

- › Like the “current range” statistics, the “current range” histogram has to re-read the trace all the time.
- › Couldn't it use the statistics state system too?

- › Demo...

Part III: Contributing code to TMF

- › Code sitting on your desktop is not solving world hunger
- › Contribution is easy!
 - Legal issues notwithstanding...
- › Linux Tools uses:
 - Git for version control. It's 2012, you should know Git!
 - Gerrit for code review

Gerrit

- › Web-based Git front-end
- › Code review
- › Push approved patches straight from the web

Gerrit flow of action

- › *git push* to a special branch
- › Anybody can review code
- › Reviewers can assign a score to the patchset
 - 2 = NAK NAK!
 - 1 = Yes but some things need fixing
 - 0 = Not yet reviewed
 - +1 = Looks good to me
 - +2 = Looks good to me, and I endorse this
- › Once the patch has approval of at least one committer (+2), it can get merged.

Gerrit Demo

So long, and thanks for all the fish!

- › We hope to see contributions, and are here to help your work get immortalized in Eclipse.

Reminders

- › Download at <http://www.eclipse.org>
 - Git: <http://git.eclipse.org/c/www.eclipse.org/linuxtools.git/>
 - Instructions at <http://ltnng.org/eclipse>
- › LTTng: <http://ltnng.org>
- › Alex: alexandre.montplaisir-gon.alves <at> ericsson.com
- › Matthew: matthew.khouzam <at> ericsson.com
- › Slides: <https://ahls.dorsal.polymtl.ca/2012decembermeeting>