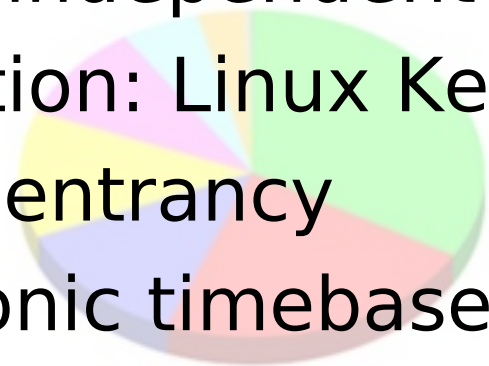


An Overview of LTTng and LTTV

Michel Dagenais
Mathieu Desnoyers
Pierre-Marc Fournier
Gabriel Matni

Department of Computer and Software Engineering
Ecole Polytechnique, Montreal

LTTng key aspects

- Low overhead
 - Architecture independent core
 - Instrumentation: Linux Kernel Markers
 - Very good reentrancy
 - Solid monotonic timebase
- 

LTTng key aspects (2)

- Atomic, buffered, flight recorder or zero-copy write to disk
- Integration into the mainline Linux kernel in progress
 - Low level primitives: ***merged***
 - Kernel Markers: ***merged (2.6.24)***
 - Instrumentation: ***pending***
 - Tracer: ***pending***

LTTV (Viewer)

- Modular, plugin-based
- Handles huge traces
- Views:
 - ◆ Raw events
 - ◆ Control flow
 - ◆ Histogram
 - ◆ Statistics
 - ◆ ...



LTTV

Typical workspace

The screenshot displays the Linux Trace Toolkit Viewer (LTTV) interface. The window title is "Linux Trace Toolkit Viewer <@latence>". The interface includes a menu bar (File, View, Tools, Plugins, Help) and a toolbar with various icons for file operations, navigation, and analysis.

The main area is divided into three sections:

- Traceset:** A graph showing the CPU usage of the traceset over time. The y-axis represents CPU usage (0 to 3402.0) and the x-axis represents time (1793s to 1809s). The graph shows a significant spike in CPU usage around 1797s.
- Process List:** A table listing the processes being traced, including their names, brands, and PIDs.
- Trace:** A detailed view of the events captured during the trace, showing the tracefile, CPUID, event name, time, and event description.

The process list is as follows:

Process	Brand	PID
su		4336
bash		4337
lftctl		4401
lftd		4418
lftd		4419
/usr/bin/find		4420
/usr/local/bin/lftctl		4421

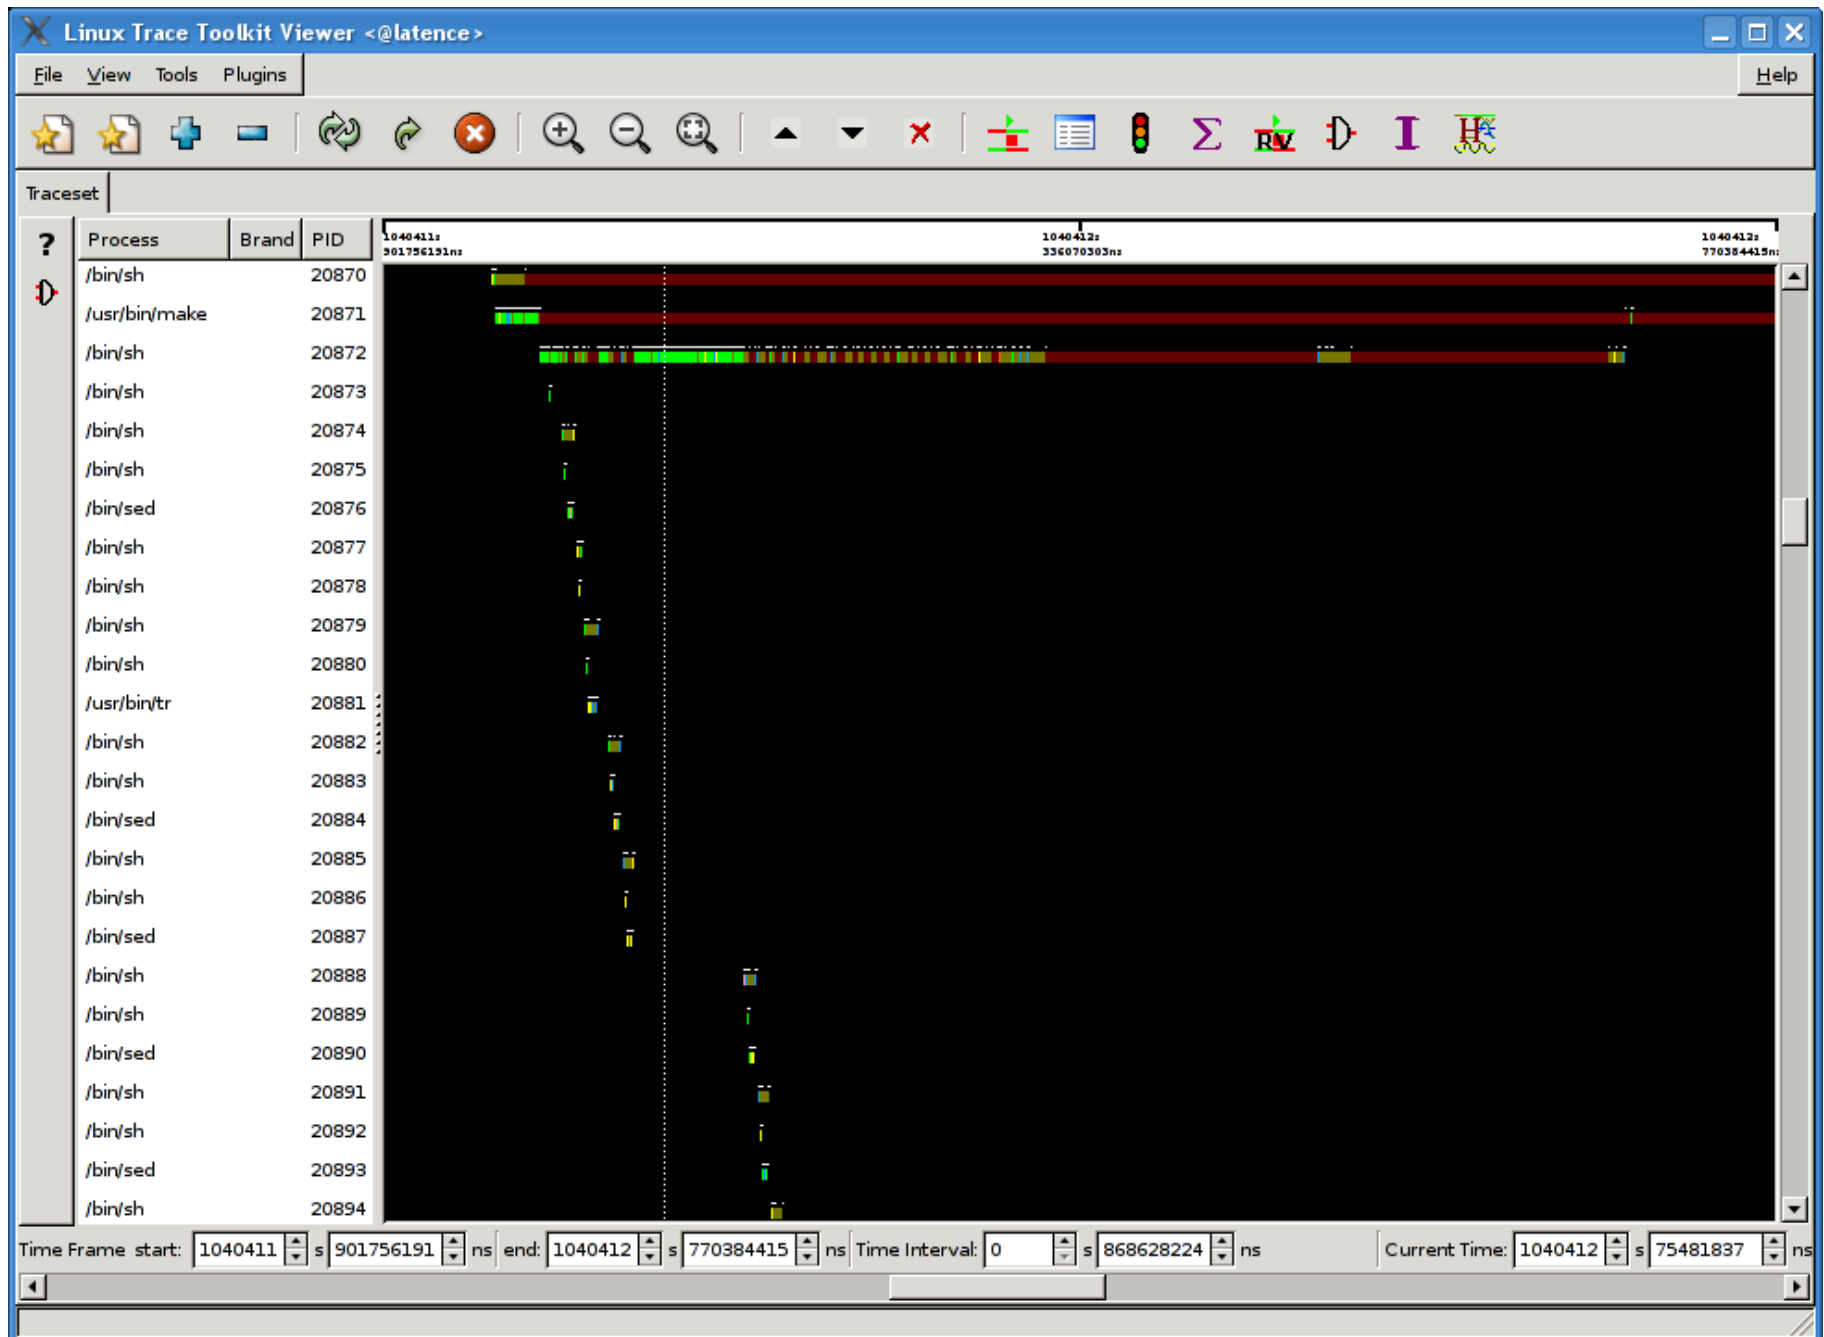
The trace table shows the following events:

Trace	Tracefile	CPUID	Event	Time (s)	Time (ns)	PID	Event Description
/home/pmf/traces/latence-20071218192151	/cpu	0	kernel_irq_entry	1797	150206353	0	kernel_irq_entry: 1797.150206353 (/home/pmf/traces/latence-20071218192151)
/home/pmf/traces/latence-20071218192151	/cpu	0	kernel_timer_update_time	1797	150208687	0	kernel_timer_update_time: 1797.150208687 (/home/pmf/traces/latence-20071218192151)
/home/pmf/traces/latence-20071218192151	/cpu	0	kernel_softirq_raise	1797	150209952	0	kernel_softirq_raise: 1797.150209952 (/home/pmf/traces/latence-20071218192151)
/home/pmf/traces/latence-20071218192151	/cpu	0	kernel_softirq_raise	1797	150211105	0	kernel_softirq_raise: 1797.150211105 (/home/pmf/traces/latence-20071218192151)
/home/pmf/traces/latence-20071218192151	/cpu	0	kernel_irq_exit	1797	150212275	0	kernel_irq_exit: 1797.150212275 (/home/pmf/traces/latence-20071218192151)
/home/pmf/traces/latence-20071218192151	/cpu	0	kernel_softirq_entry	1797	150212856	0	kernel_softirq_entry: 1797.150212856 (/home/pmf/traces/latence-20071218192151)

The bottom of the interface shows the Time Frame settings: start: 1793 s 259329151 ns, end: 1809 s 289533377 ns, Time Interval: 16 s 30204226 ns, Current Time: 1797 s 146705925 ns.

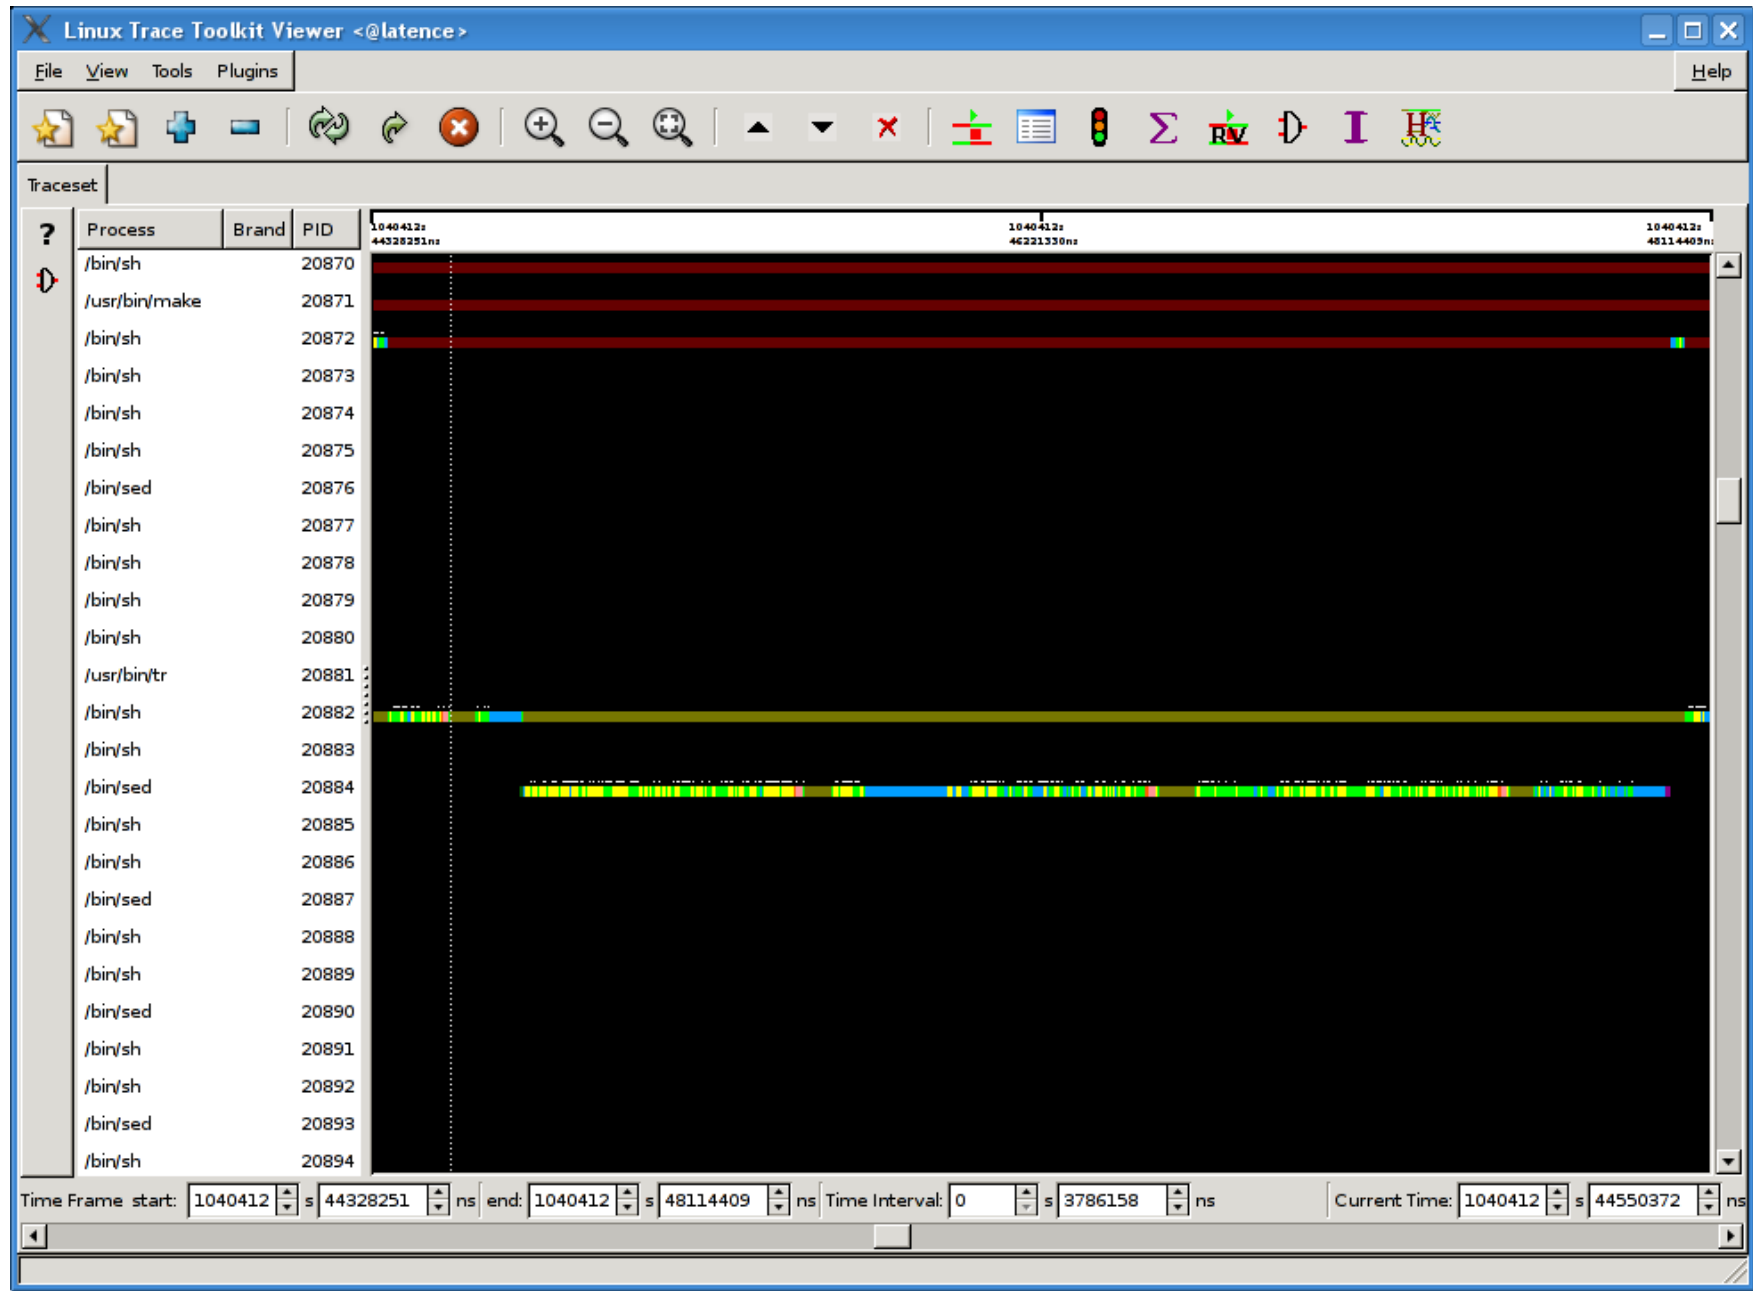
Control flow view

High level view of a compilation



Control flow view (2)

Zoom into a process: What's that long system call? (blue)



Control flow view (3)

It's execve()!

Linux Trace Toolkit Viewer <@latence>

File View Tools Plugins Help

Traceset

Process	Brand	PID
/bin/sed		20876
/bin/sh		20877
/bin/sh		20878
/bin/sh		20879
/bin/sh		20880
/usr/bin/tr		20881
/bin/sh		20882
/bin/sh		20883
/bin/sed		20884
/bin/sh		20885
/bin/sh		20886
/bin/sed		20887

Event Description

```
kernel_arch_trap_exit: 1040412.045721157 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, USER_MODE
kernel_arch_syscall_entry: 1040412.045723463 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL { syscall_id = 11 [sys_execve+0x0/0x7e],
mm_handle_fault_entry: 1040412.045779393 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL { address = 3221225459, ip = 0xffffe410, v
mm_page_alloc: 1040412.045785053 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL { order = 0, pfn = 18948 }
mm_page_alloc: 1040412.045791685 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL { order = 0, pfn = 90545 }
mm_handle_fault_exit: 1040412.045792815 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL
mm_page_free: 1040412.045847742 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL { order = 0, pfn = 90544 }
mm_page_free: 1040412.045851296 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL { order = 0, pfn = 52494 }
mm_page_free: 1040412.045852511 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL { order = 0, pfn = 17475 }
mm_page_free: 1040412.045853828 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL { order = 0, pfn = 9812 }
mm_page_free: 1040412.045855030 (/home/pmf/traces/latence-20080125164625/cpu_0), 20884, 20884, /bin/sh, , 20882, 0x0, SYSCALL { order = 0, pfn = 9815 }
```

Time Frame start: 1040412 s 45601031 ns end: 1040412 s 46074300 ns Time Interval: 0 s 473269 ns Current Time: 1040412 s 45723463 ns

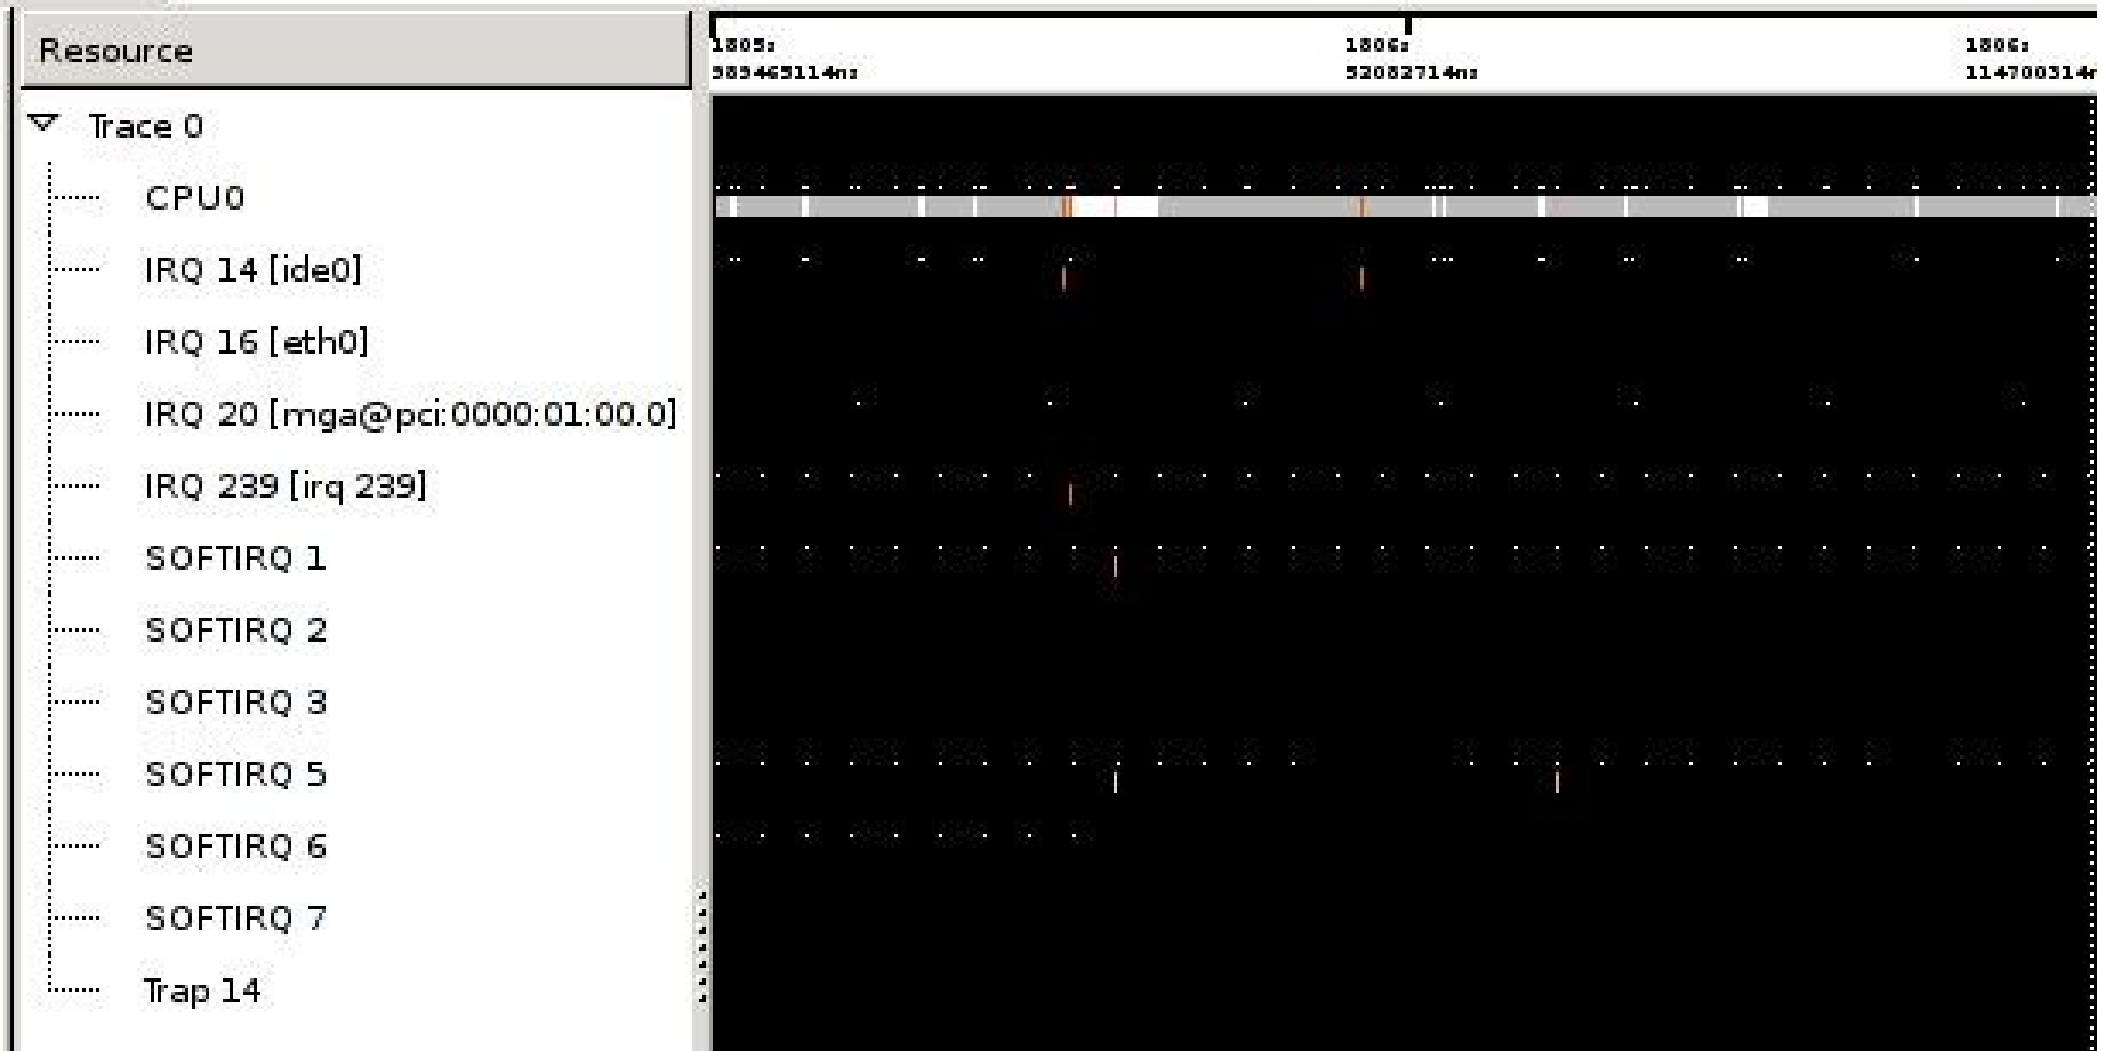
What's new?

- Resource viewer
 - CPUs
 - IRQs
 - softIRQs
 - traps
- View traces of virtual machines and their host




Resource view

“Zoomed out”



Resource view (2)

Zoomed in: activity after a timer interrupt



Resource view (3)

Multiple CPUs

The screenshot displays the Linux Trace Toolkit Viewer interface. The main window shows a resource view for multiple CPUs (CPU0, CPU1) and various IRQs (IRQ 14, IRQ 23, IRQ 239, SOFTIRQ 1). The interface includes a menu bar (File, View, Tools, Plugins), a toolbar with various icons, and a Traceset panel on the left. The main display area shows a timeline of events for the selected resources. Below the main display, there is a table listing processes and their attributes, and another table listing trace events with their details.

Process	Brand	PID	TGID	PPID	CPU	Birth sec	Birth nsec	TRACE
sensor		2083	2083	1	1	1125914	6409577	0
Xorg		2138	2138	2136	1	1125914	6419137	0
gdmgreeter		2216	2216	2136	1	1125914	6449015	0
sshd		6256	6256	6253	1	1125914	6468076	0
ltd		6339	6338	1	1	1125914	6480371	0
/usr/bin/find		6340	6340	6259	1	1125924	884567936	0
./loop		6341	6341	6259	0	1125924	889248356	0

Trace	Tracefile	CPUID	Event	Time (s)	Time (ns)	PID	Event Description
/home/pmf/traces/amd64-20080125143547 /cpu		0	kernel_irq_entry	1125933	644468854	6341	kernel_irq_entry: 1125933.644468854 (/home/pmf/traces/amd64-20080125143547 /cpu)
/home/pmf/traces/amd64-20080125143547 /cpu		0	kernel_timer_update_time	1125933	644473588	6341	kernel_timer_update_time: 1125933.644473588 (/home/pmf/traces/amd64-20080125143547 /cpu)
/home/pmf/traces/amd64-20080125143547 /cpu		0	kernel_softirq_raise	1125933	644475149	6341	kernel_softirq_raise: 1125933.644475149 (/home/pmf/traces/amd64-20080125143547 /cpu)
/home/pmf/traces/amd64-20080125143547 /cpu		0	kernel_irq_exit	1125933	644476812	6341	kernel_irq_exit: 1125933.644476812 (/home/pmf/traces/amd64-20080125143547 /cpu)
/home/pmf/traces/amd64-20080125143547 /cpu		0	kernel_softirq_entry	1125933	644477878	6341	kernel_softirq_entry: 1125933.644477878 (/home/pmf/traces/amd64-20080125143547 /cpu)
/home/pmf/traces/amd64-20080125143547 /cpu		0	kernel_softirq_exit	1125933	644479978	6341	kernel_softirq_exit: 1125933.644479978 (/home/pmf/traces/amd64-20080125143547 /cpu)

Time Frame start: 1125928 s 582889110 ns end: 1125938 s 701177386 ns Time Interval: 10 s 118288276 ns Current Time: 1125933 s 642033248 ns

Virtual machine and host

