

Embedded Linux Conference Europe 2013

Linux Kernel Debugging And Profiling Tools

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax :+33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Agenda



ELC Europe 2013

- **Introduction**
- **Kernel Features**
- **Tracers**
- **Debuggers**
- **Profilers**
- **Miscellaneous**
- **Questions**

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52


Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Introduction



Motivations:

- BSP adaptations to various hardwares
- Drivers development
- Optimizations

Debugging and profiling:

What are the available features in the kernel ?
What tools exist in the community ?

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Definitions

Tracing

Specialized use of logging to collect state or behavior information about a program's execution. Done by addition of trace points in the program's source code.

Interactive Debugging

Process of monitoring the execution of a program by interactively inspecting and/or modifying its internal state (memory, stacks, registers, etc). No source code modifications but requires specialized software.

Profiling

Process of sampling the execution of a program to analyze its performance.

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Agenda



ELC Europe 2013

- Introduction
- **Kernel Features**
- Tracers
- Debuggers
- Profilers
- Miscellaneous
- Questions

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Kernel Features

printk



ELC Europe 2013

```
linux/printk.h:  
asmlinkage __printf(1, 2) __cold  
int printk(const char *fmt, ...);
```

- Famous kernel's printf like function with log levels support
- Circular log buffer, can be called from any context
- CONFIG_PRINTK, /proc/sys/kernel/printk, dmesg tool
- Documentation/printk-formats.txt
- Probably the most commonly used routine for kernel debugging
- Unconditional trace point, overheads

```
printk(KERN_DEBUG "Got there: %s:%i\n", __FILE__, __LINE__);  
printk(KERN_CRIT "I give up, bye\n");
```

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Kernel Features

DebugFS



ELC Europe 2013

```
linux/debugfs.h:  
debugfs_create_file  
debugfs_create_dir  
debugfs_create_blob  
debugfs_remove  
...
```

- Special RAM-based file system dedicated to kernel debugging
- No special rules, developers are free to put any info they want there
- CONFIG_DEBUG_FS
- Documentation/filesystems/debugfs.txt
- Typically mounted in /sys/kernel/debug

```
# mount -t debugfs none /sys/kernel/debug
```

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Kernel Features

Dynamic Debug



ELC Europe 2013

```
linux/printk.h:
#define pr_debug(fmt, ...) dynamic_pr_debug(fmt, ##__VA_ARGS__)

linux/device.h
#define dev_dbg(dev, format, ...) \
do { dynamic_dev_dbg(dev, format, ##__VA_ARGS__); } while (0)
```

- Dynamic activation/deactivation of kernel information code
- Currently, dynamic activation of `pr_debug`/`dev_dbg` per-callsite
- `CONFIG_DYNAMIC_DEBUG`
- `Documentation/dynamic-debug-howto.txt`
- Debugfs control file
`<debugfs>/dynamic_debug/control`

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Kernel Features

Probes



ELC Europe 2013

- Static probes: tracing oriented
 - Present in the source code
 - Optional code path: unlikely branch
 - Small overhead even when off
- Dynamic probes: debug oriented
 - Optional code path : breakpoint
 - No overhead when off, High overhead when on
 - Debug symbols and deep knowledge of the kernel mapping required

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax :+33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Kernel Features

Static probes



ELC Europe 2013

- Kernel Markers
 - First implementation
 - Mainline from 2.6.24 to 2.6.32
- Tracepoints
 - Dissociates instrumentation code from kernel code
 - Mainline since 2.6.28
- `TRACE_EVENT` macro
 - Tracepoints enhancement with automatic callback
 - Mainline since 2.6.32

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle


3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Kernel Features

Static probes



Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

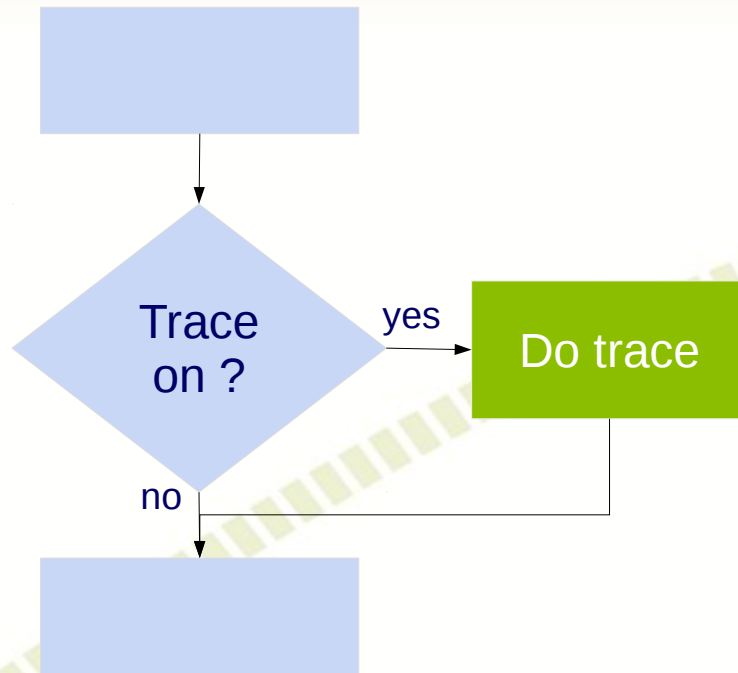
Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

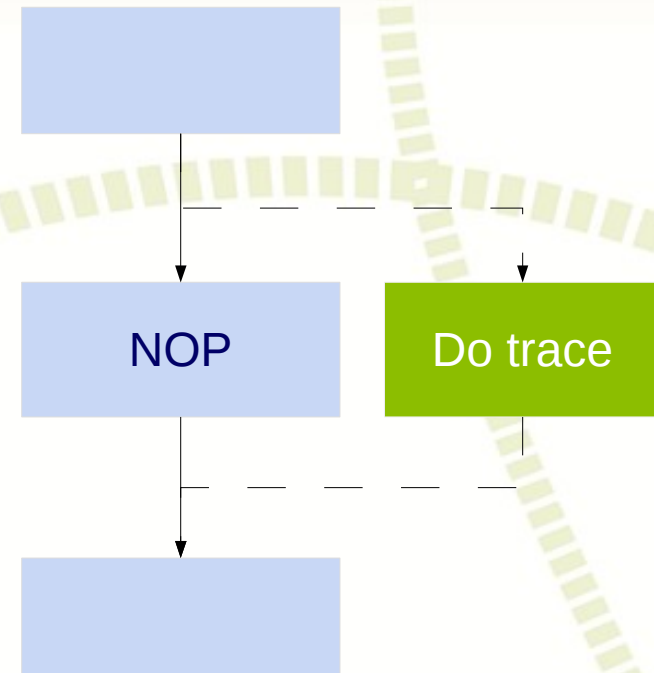
Kernel Features

Static probes

Before 2.6.37



After 2.6.37



Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

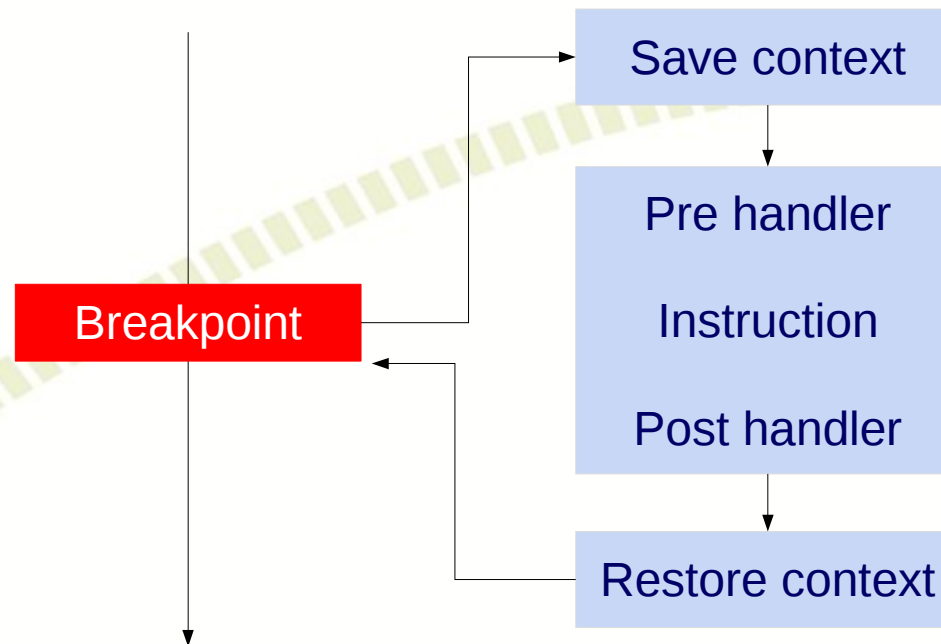
Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Kernel Features

Dynamic probes

- Kprobes
 - Software breakpoint
 - Mainline since 2.6.9 (ARM: 2.6.25)



Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Kernel Features

Perf events



ELC Europe 2013

- API for performance measurements
- Hardware events
 - Special registers on most modern CPUs
 - Limited number, round robin when needed
 - Cycles, cache misses, branch hits ...
- Software events
 - Context switches, page faults ...
- Mainline since 2.6.31 (ARM: 2.6.34)
- On ARM, access to the PMU through the CP15

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Kernel Features

Slabinfo



ELC Europe 2013

- Functionality
 - Memory usage on the slabs
 - Access through /proc/slabinfo
- Mainline since 2.6.24

```
# name          <active_objs> <num_objs> <object_size> <objperslab>
ext4_groupinfo_4k 24           24           168           24
SCTP              12           12           1280          12
rxrpc_call_jar    0            0            832           19
bsg_cmd           0            0            288           14
[...]            [...]
```

Agenda



ELC Europe 2013

- Introduction
- Kernel Features
- **Tracers**
- Debuggers
- Profilers
- Miscellaneous
- Questions

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Tracers

Ftrace (Function Tracer)



ELC Europe 2013

- **Functionality**
 - Various tracers (function, graph function, irqs-off, wakeup, branch)
 - Various time sources (local and global clocks, atomic counter)
 - Function and event filters
 - Early boot tracing
 - Dump on Oops
 - No user tool required, debugfs access
- Mainline since 2.6.27 (ARM: 2.6.34)
- Based on mcount and debugfs, optionally on tracepoints and kprobes

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

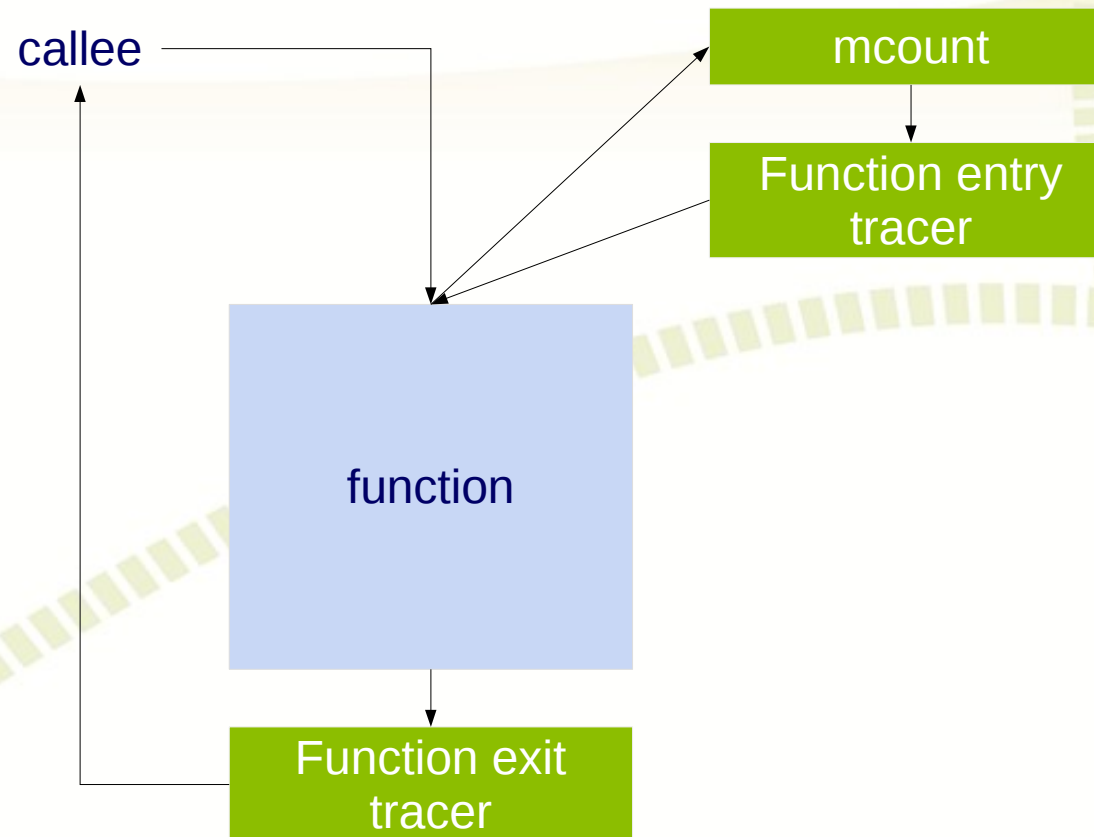
Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Tracers Ftrace



Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Tracers

Ftrace



ELC Europe 2013

Function tracer :

```
# TASK-PID CPU#  TIMESTAMP FUNCTION
#      ||      |           |
<idle>-0   [001] 6075.461563: mwait_idle <-cpu_idle
bash-16939 [000] 6075.461563: __fsnotify_parent <-vfs_write
[...]
```

Branch tracer :

```
# TASK-PID CPU#  TIMESTAMP CORRECT FUNC:FILE:LINE
#      | |      |           |           |
Kworker/0:1-12 [000] d... 7.440216: [ ok ] console_unlock:printk.c:2108
<idle>-0 [000] d.h. 7.440383: [ MISS ] note_interrupt:spurious.c:284
```

Stack tracer :

```
Depth  Size  Location (34 entries)
-----
0) 2952 16 mempool_alloc_slab+0x15/0x17
1) 2936 144 mempool_alloc+0x52/0x104
```

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Tracers

Ftrace



ELC Europe 2013

Irqs-off :

```
#          _-----=> CPU#
#          /  _-----=> irqs-off
#          | /  _----=> need-resched
#          || /  _---=> hardirq/softirq
#          ||| /  _--=> preempt-depth
#          |||| /                delay
# cmd pid  ||||| time  | caller
#  \ /    |||||  \   | /
<idle>-0  0d ...   0us +: __dev_pm_qos_read_value <-rpm_check_suspend_allowed
<idle>-0  0d ...   31us : __rpm_callback <-rpm_suspend
[...]
```

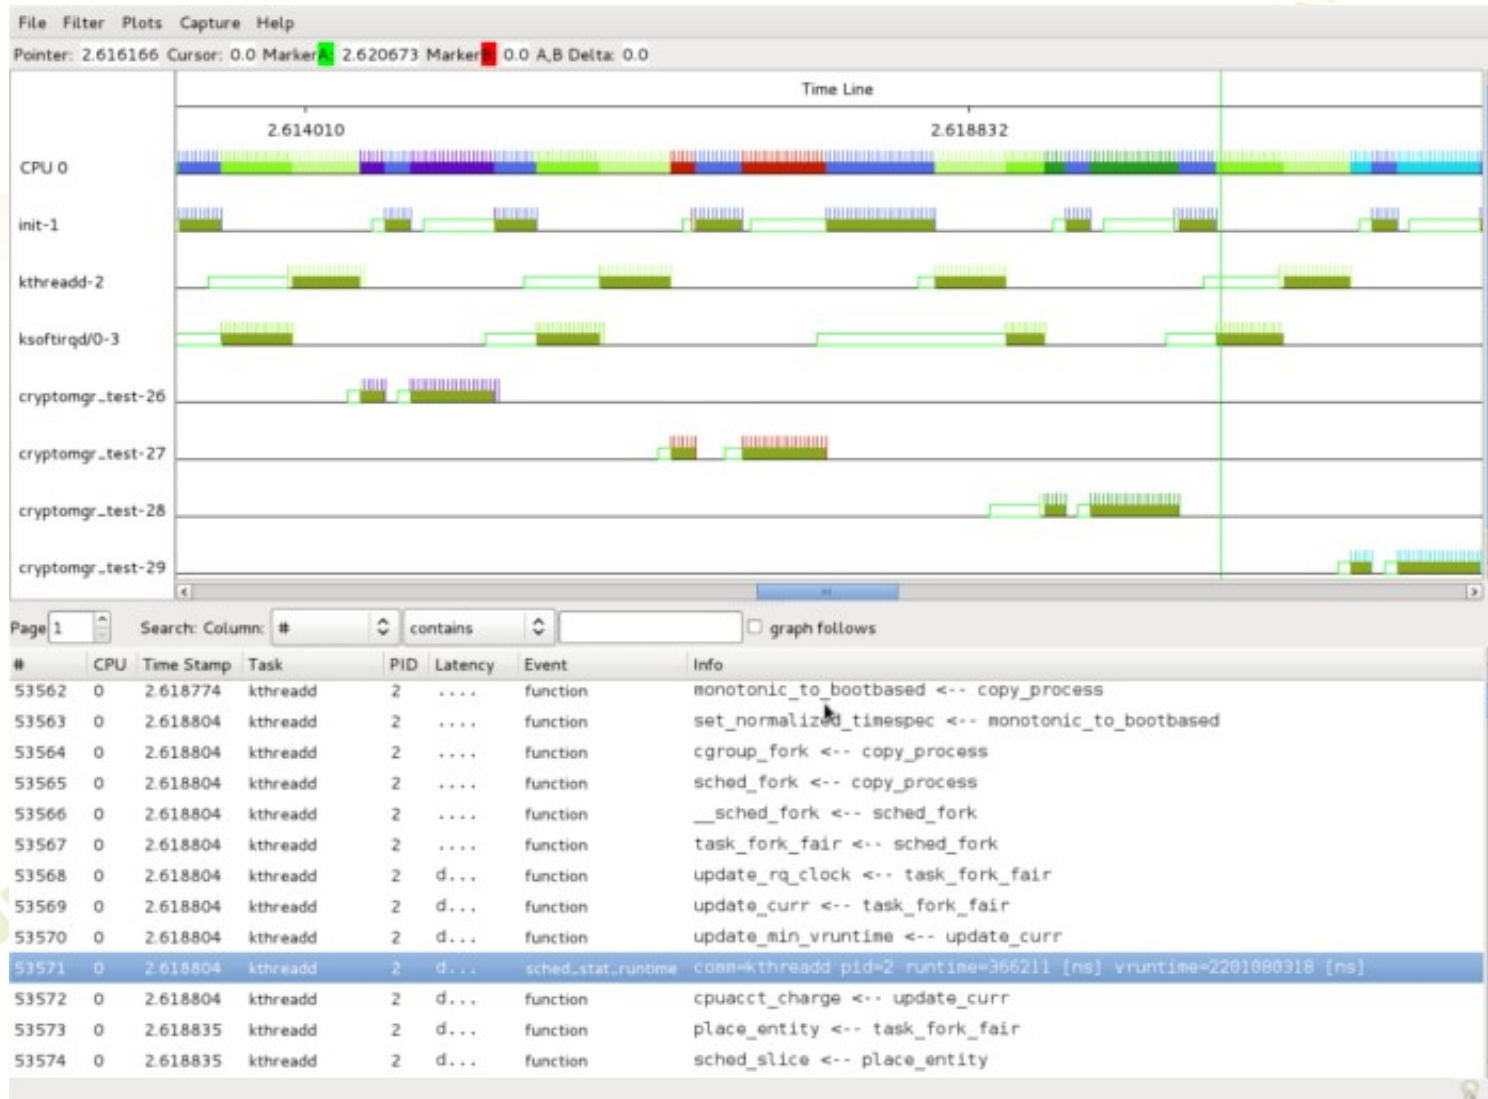
. =normal context s= soft irq h=hard irq H=hard irq occured in a soft irq
= <1μs += <10μs != >100μs

Tracers

Trace-cmd & Kernelshark



ELC Europe 2013



Tracers Lttng



ELC Europe 2013

- **Functionality:**
 - Tracepoint tracer
 - Kprobes support
 - Mcount support
- Not in mainline (modules)
- Available since 2.6.9 (ARM: 2.6.20)
- Based on tracepoints, optionally on kprobes and mcount

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Tracers Lttng



ELC Europe 2013

The screenshot shows the LTTng Kernel - Eclipse Platform interface. The top bar indicates the date and time: "mer. 17 avril, 14:22". The main window displays the "Control Flow" view for a test named "test1". The timeline shows CPU 0 activity from 02:40:03.434 to 02:40:03.446. The "Events - test1" table is visible below the timeline.

Timestamp	Channel	Event Type	Content
02:40:03.435468	channel0_0	sched_stat_runtime	tid=163, comm=lttng-consumerd, vruntime=8643580247, runtime=4394531, context._procname=lttng-sessiond, context._pid=111, context._perf_cs=1597, context._perf_cpu=0
02:40:03.435468	channel0_0	sched_wakeup_new	tid=260, target_cpu=0, comm=lttng-consumerd, prio=120, success=1, context._procname=lttng-sessiond, context._pid=111, context._perf_cs=1597, context._perf_cpu=0
02:40:03.435689	channel0_0	timer_cancel	timer=3658045992, context._procname=lttng-sessiond, context._pid=111, context._perf_cs=1597, context._perf_cpu=0

The "Histogram" view at the bottom shows a bar chart of event counts over time, with a current event count of 946690803.43546875 and a window span of 0.016460904 seconds.

Agenda



ELC Europe 2013

- Introduction
- Kernel Features
- Tracers
- **Debuggers**
- Profilers
- Miscellaneous
- Questions

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Debuggers

KGDB



ELC Europe 2013

- Functionality
 - GDB server
 - Serial transport / kdb front-end
 - Breakpoints, watchpoints and step-by-step debugging
 - Early boot debug
- Mainline since 2.6.26 (KDB merge: 2.6.35)

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Debuggers KGDB



ELC Europe 2013

The screenshot shows the Eclipse IDE interface for debugging a kernel module. The main window is titled "Debug - linux/drivers/tty/serial/kgdboc.c - Eclipse Platform". The "Debug" view shows the call stack for the function `configure_kgdboc()` at address `0xc046d920`. The "Disassembly" view shows the assembly code for `configure_kgdboc+628` to `configure_kgdboc+656`. The "Source" view shows the C code for `kgdboc.c` with the `do_register` function. The "Registers" view shows the state of the registers `r0`, `r1`, and `r2`.

Name	Type	Value
tty_line	int	2
err	int	0

Name	Value	Description
General Registers		
r0	0	General Purpose and FPU Register Group
r1	0	
r2	3740996804	

Debuggers

KGTP



ELC Europe 2013

- Functionality:
 - GDB server / kprobe
 - tracepoint support
 - Serial and Ethernet transports
- Not in mainline (module)
- Available since 2.6.18
- Based on kprobes, optionally on tracepoints

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Agenda



ELC Europe 2013

- Introduction
- Kernel Features
- Tracers
- Debuggers
- **Profilers**
- Miscellaneous
- Questions

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Profilers Perf



ELC Europe 2013

- An API but also a set of tools
- Functionality:
 - Various profilers
 - Event based sampling
- Mainline since 2.6.31 (ARM: 2.6.34)
- Based on perf_events, optionally on tracepoints
- On ARM, Oprofile uses perf as backend

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax :+33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Profilers Perf



ELC Europe 2013

Perf Top :

PerfTop: 279 irqs/sec kernel:96.1% exact: 0.0% [4000Hz cycles], (all, 1 CPU)

```
60.77% [kernel] [k] cpuidle_wrap_enter
10.29% [kernel] [k] finish_task_switch.constprop.59
8.99% [kernel] [k] tick_nohz_idle_enter
3.29% [kernel] [k] tick_nohz_idle_exit
2.33% [kernel] [k] __do_softirq
0.96% [kernel] [k] ftrace_likely_update
0.92% [kernel] [k] rcu_idle_enter
[...]
```

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Profilers Perf



ELC Europe 2013

Perf report :

Samples: 8K of event 'cycles'

Event count (approx.): 1745355260

#

Overhead Command Shared Object Symbol

.....

#

9.33% hackbench [kernel.kallsyms] [k] ftrace_likely_update

9.21% hackbench [kernel.kallsyms] [k] finish_task_switch.constprop.59

6.37% hackbench [kernel.kallsyms] [k] memchr_inv

4.75% hackbench [kernel.kallsyms] [k] unix_stream_recvmsg

2.78% hackbench [kernel.kallsyms] [k] __wake_up_sync_key

[...]

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Profilers Perf



ELC Europe 2013

Perf annotate :

Percent | Source code & Disassembly of libpthread-2.15.so

```
-----  
:  
:  
: Disassembly of section .text:  
:  
:  
: 0000f7a0 <__read>:  
7.32 : f7a0: ldr ip, [pc, #-40] ; f780 <__local_syscall_error+0x28>  
3.25 : f7a4: ldr ip, [pc, ip]  
6.50 : f7a8: teq ip, #0  
0.00 : f7ac: push {r7} ; (str r7, [sp, #-4]!)  
0.00 : f7b0: bne f7cc <__read+0x2c>  
[...]
```

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911


Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

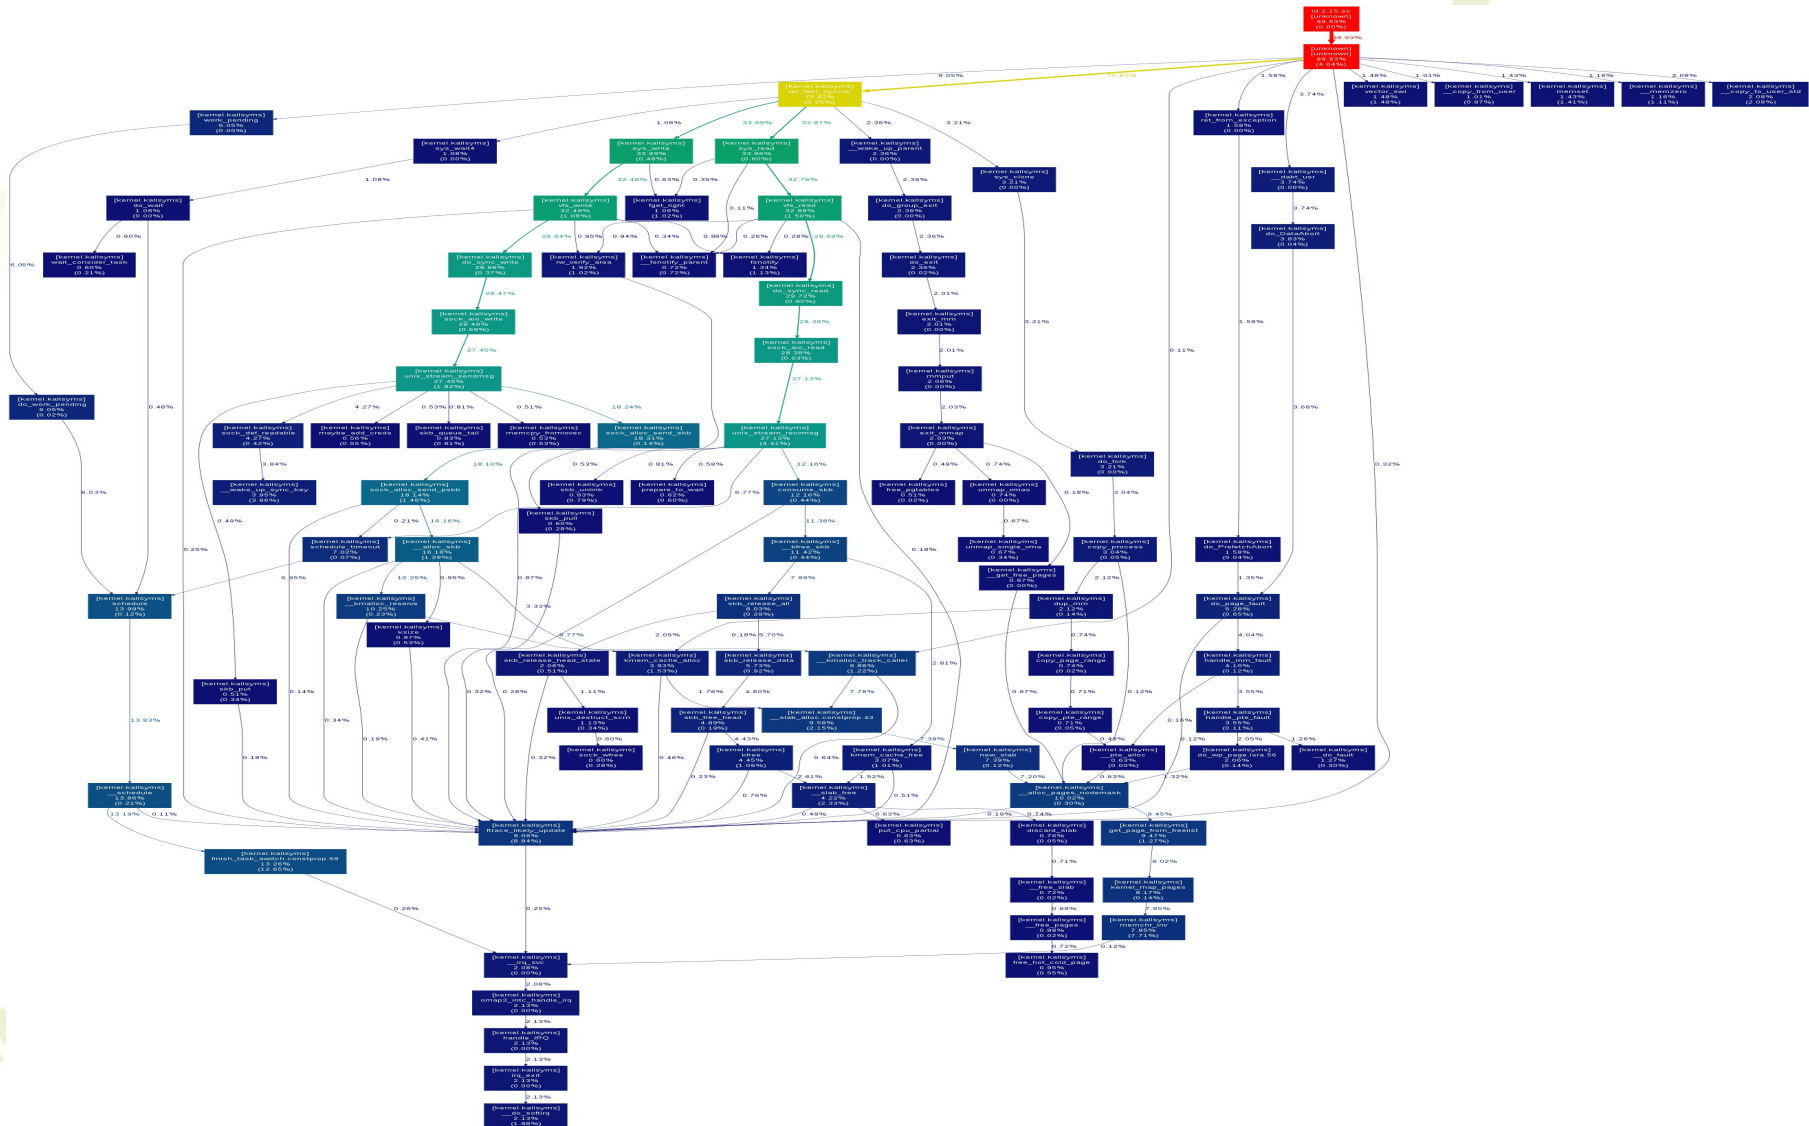
Profilers Perf



ELC Europe 2013



Profilers Gprof2dot

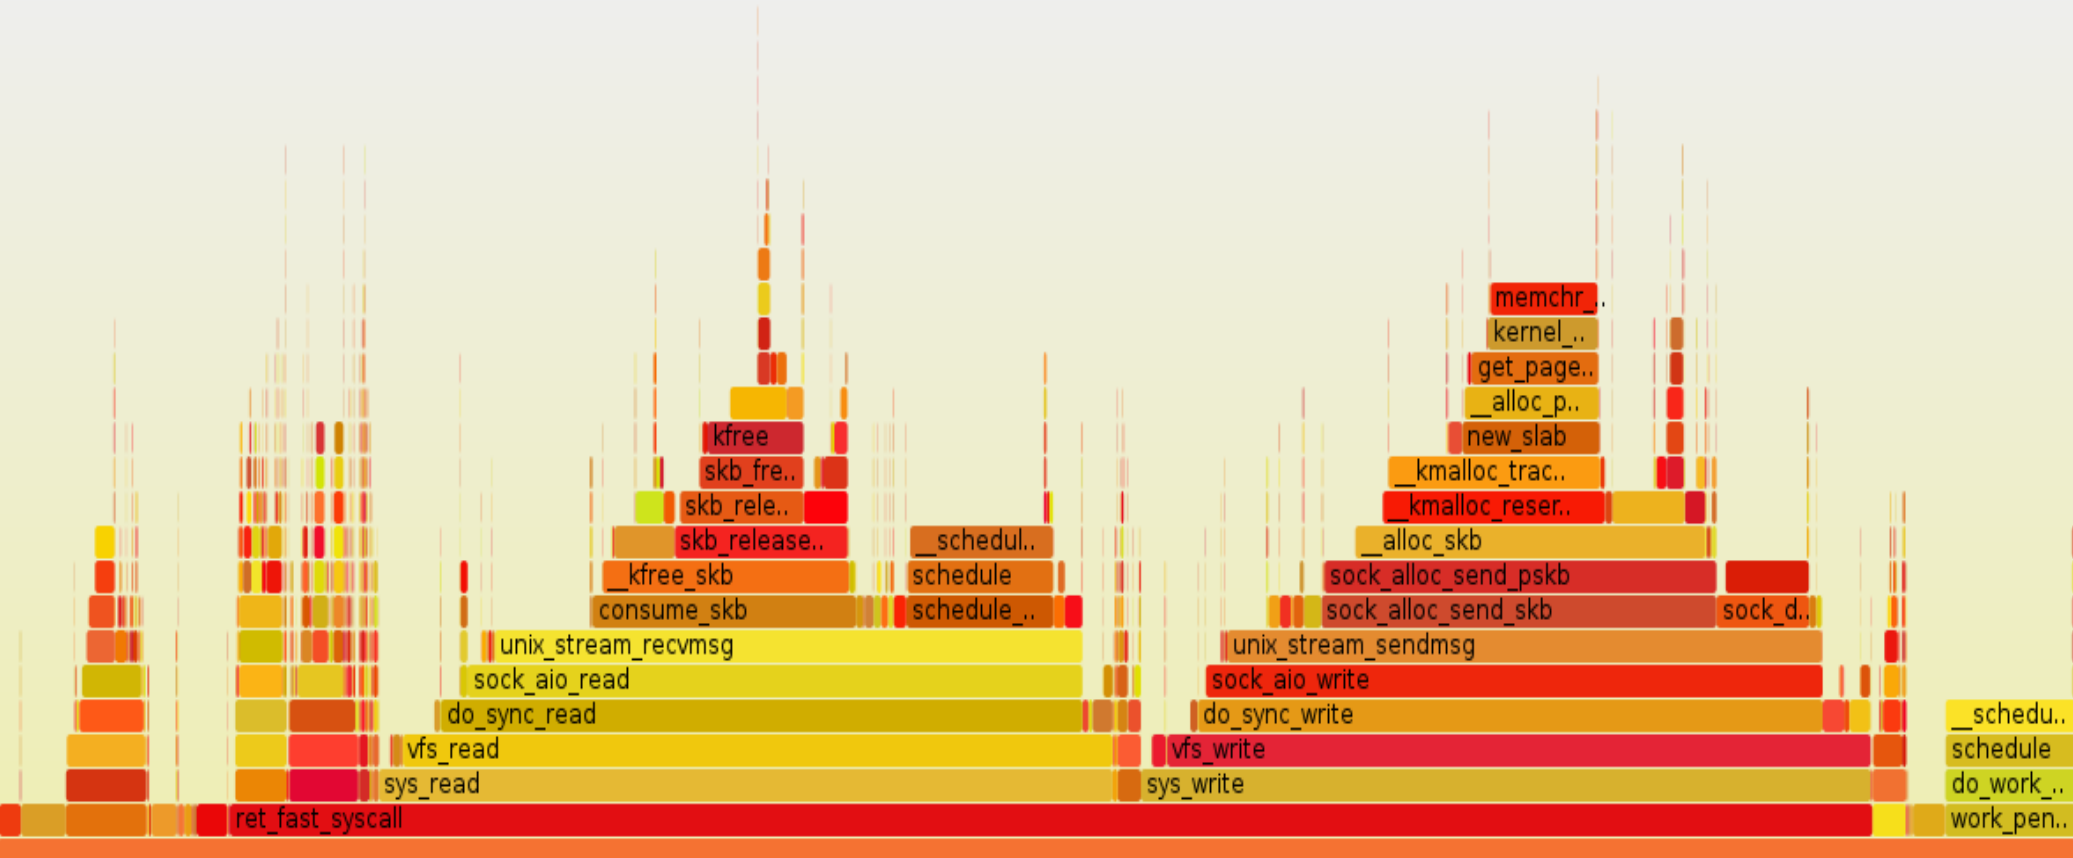


Profilers FlameGraph



ELC Europe 2013

Flame Graph



Function:

Headquarters
2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris
3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle
3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt
Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Profilers

Kmemleak



ELC Europe 2013

- **Functionality**
 - Garbage collector based memory profiler
 - Kmalloc, vmalloc and friends
 - Log pointer, size and backtrace, unlog at free
 - Periodic scan (default: 10 min)
- **Mainline since 2.6.31**

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Profilers

Kmemleak



ELC Europe 2013

```
# cat /sys/kernel/debug/kmemleak
unreferenced object 0xef42d000 (size 28):
 comm "khubd", pid 189, jiffies 4294937550 (age 2543.370s)
 hex dump (first 28 bytes):
 00 01 10 00 00 02 20 00 08 d0 42 ef 08 d0 42 ef
 00 00 00 00 00 00 00 00 ff ff ff ff
 backtrace:
 [<c0080fe1>] create_object+0xa1/0x1ac
 [<c007eac5>] kmem_cache_alloc+0x8d/0xdc
 [<c01a966d>] isp1760_urb_enqueue+0x2f9/0x358
 [<c019bbbd>] usb_hcd_submit_urb+0x75/0x574
 [<c019d8f1>] usb_start_wait_urb+0x29/0x80
 [<c019daad>] usb_control_msg+0x89/0xac
```

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Agenda



ELC Europe 2013

- Introduction
- Kernel Features
- Tracers
- Debuggers
- Profilers
- **Miscellaneous**
- Questions

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Miscellaneous SystemTap



ELC Europe 2013

- **Functionality:**
 - Lightweight debug
 - Kprobes front-end using a scripting language
 - Trace dump on Oops
 - Tracepoint support
 - Perf support
- Not in mainline (module)
- Available since 2.6.9
- Based on kprobe, optionally on tracepoint and perf_event

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

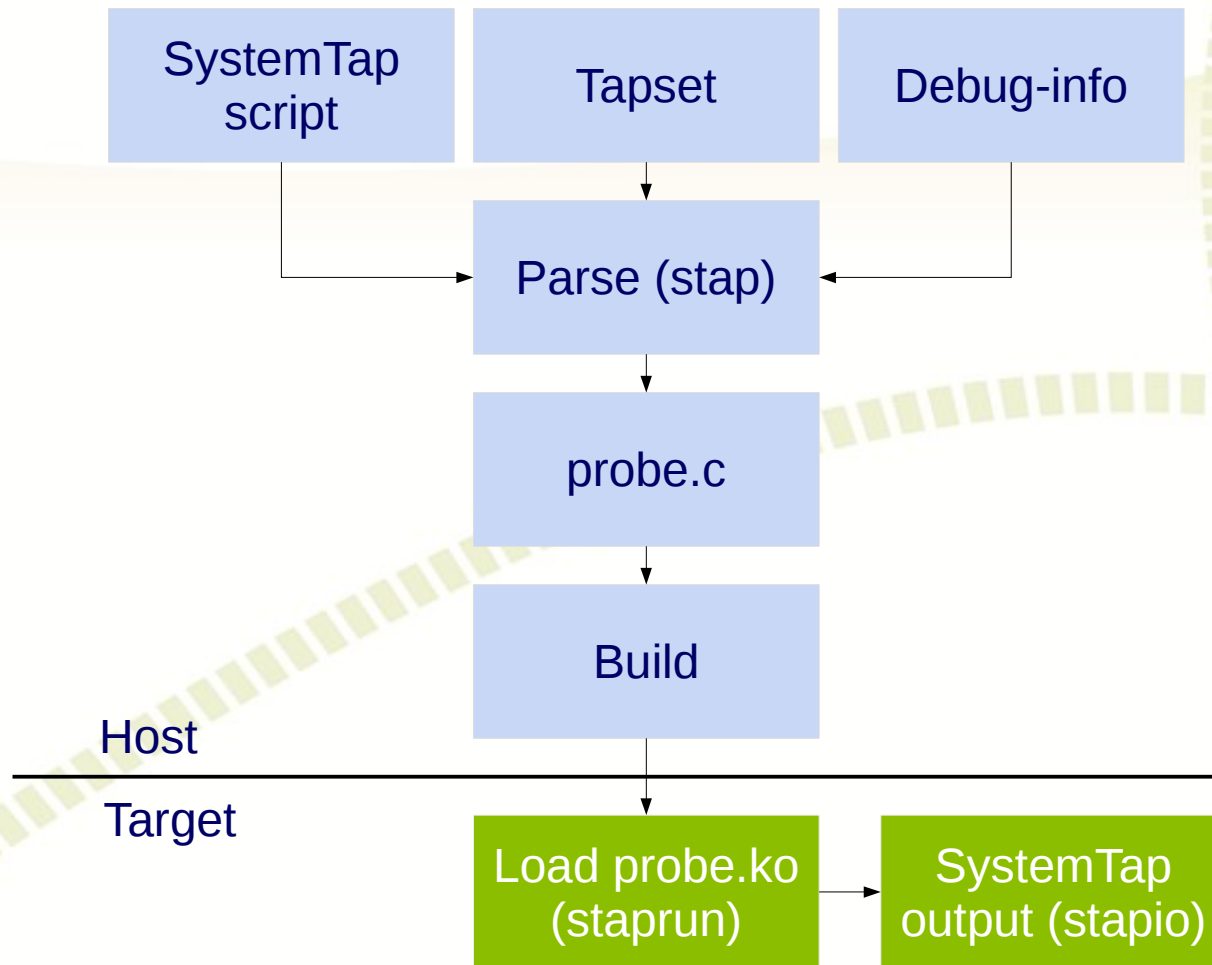
Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Miscellaneous SystemTap



Miscellaneous SystemTap



ELC Europe 2013

```
%{  
#include<linux/netdevice.h>  
%}  
  
function change_mtu (dev:long, val:long) %{  
    int ret;  
    struct net_device *netdevice;  
    netdevice = (struct net_device *) (long) THIS->dev;  
    dev_set_mtu(netdevice, THIS->val);  
%}  
  
probe module("tg3").function("tg3_get_stats") {  
    change_mtu($dev, 2000)  
    printf("Changing tg3 MTU\n")  
}
```

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Miscellaneous Kexec & Crash



ELC Europe 2013

- Functionality
 - Post-mortem analysis
 - Boot on a new kernel instance on panic
 - Old kernel memory dump through `/proc/vmcore`
- Mainline since 2.6.13

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Agenda



ELC Europe 2013

- Introduction
- Kernel Features
- Tracers
- Debuggers
- Profilers
- Miscellaneous
- **Questions**

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax : +33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070

Questions & Answers



ELC Europe 2013

Thank you for your attention

Any Question ?

Headquarters

2 chemin du Ruisseau – 69134 Ecully, France
Phone : +33 4 26 49 25 39 / Fax :+33 4 72 18 08 41

Adeneo Embedded Paris

3 rue Galvani – 91300 Massy, France
Phone : +33 1 80 75 01 52

Adeneo Embedded Seattle

3150 Richards Road, Suite 210 – Bellevue, WA 98005, USA
Phone : +1 425 749-4335 / Fax : +1 425 818-1911

Adeneo Embedded Frankfurt

Pfingstweide 51 – 61169 Friedberg, Deutschland
Phone : +49 6031 693 7070