

Linux Plumbers Conference 2012

LTTngTop: Human-Readable Trace Viewer

Julien Desfossez
<jdesfossez@efficios.com>

> Presenter

- Julien Desfossez
- EfficiOS Inc.
 - <http://www.efficios.com>
- Part of the LTTng Team :
 - `lttng-tools`, `lttng-modules`, `lttng-ust`, `babeltrace`
- Author of LttngTop

> Content

- LTTng 2.0 features
- LTTngTop overview
- Work in progress
- Demo
- Future

> LTTng 2.x features

- Low-impact kernel and user-space tracer
- TRACE_EVENT support
- Perf PMU counters support
- Kprobes support

> LTTngTop Overview

- Sysadmin-oriented kernel trace viewer
- Gather statistics only from kernel events (no constant /proc hammering)
- Lightweight and console-based (ncurses)
- top-like look and feel
- CPU usage, I/O statistics, Perf PMU counters evolution per-process

> LTTngTop Overview

- Replay the recorded trace at the same rate it happened
- Pause and navigate back and forth in the history
- Display the state of the system at any point in time (CPU usage, opened files, bandwidth, process creation/termination, etc)

> LTTngTop use-cases

- Sysadmin
 - Hard to reproduce bugs (happens sometimes, disappears by itself)
 - Detailed statistics per-process at any point in time (including opened files)
- Developer
 - Quickly isolate interesting events from a given trace and then read its text dump?

> LTTngTop Work in Progress

- Live tracing/viewing !
- Containers support (LXC) :
 - vpid/vtid/vppid
 - nesting and hierarchy support
- Viewer-side filtering :
 - selected processes (with/without childs)
 - selected container (with/without nesting)

> LTTngTop Work in Progress

- “Attach” to a live process (graphical and textdump like strace without ptrace overhead)
- CPU hotplug support
- On-the-fly Perf PMU counters activation
- On-the-fly kprobe registration (and hit stats)

> LTTngTop Demo

```
sinkpad:/home/julien 82x42
Statistics for interval [20:12:14.225117603, 20:12:15.225118300[
  CPUs      2      (max/cpu : 50.00%)
  Threads   367    (+1, 0)
  FDs       1651   (+5, -8)          5KB/sec

CPU Top
CPU(%)  PID    TID    NAME
0.31    4129   4129   firefox-bin
0.28    7709   7709   ifconfig
0.22    4196   4196   wicd
0.12    4971   4971   kworker/1:2
0.12    7447   7447   kworker/0:1
0.11    7373   7373   /usr/bin/x-term
0.07    2580   2580   Xorg
0.07    2441   2441   dbus-daemon
0.07    4227   4227   wicd-monitor
0.03    4075   4075   tor
0.03    6021   6021   xscreensaver
0.01    7298   7298   kworker/u:0
0.01    6808   6808   kworker/u:2
0.01    2498   2498   acpid
0.01    5957   5957   awesome
0.00    4114   4114   uml_switch
0.00    2585   2585   wpa_supplicant
0.00    7682   7682   lttngtop
0.00    18675  18675  migration/1
0.00    28967  28967  watchdog/0
0.00    28969  28969  watchdog/1
0.00    7682   7683   lttngtop
0.00    7682   7684   lttngtop
0.00    7588   7648   lttng-sessiond
0.00    7588   7588   lttng-sessiond
0.00    0      0      swapper/1
0.00    7274   7274   kworker/0:2

Status
Going forward in time
Manually moving forward
Manually moving forward
Manually moving forward

F2:CPUtop  F3:PerfTop  F4:IOtop  Enter:Details  Space:Highlight  q:Quit  r:Pref
```

> LTTngTop Future

- More analysis modules (feedbacks and suggestions welcomed !)
 - Separate disk vs network I/O
 - Disk latency analysis
- Clean integration of the live tracing feature with lttng-tools and babeltrace

> LTTngTop Future

- Separate the analysis modules from the core to allow the use in “desktop widgets”, snmpd, etc.
- Large-scale data center use-cases :
 - Remote traces live analysis (monitoring, trending and debugging)
 - Distributed analysis computation
 - High-level to fine-grained view

> Questions ?

LTTngTop available at <http://lttng.org>

```
git clone -b lttngtop-live git://git.dorsal.polymtl.ca/~jdesfossez/lttng-tools  
git clone -b lttngtop-live git://git.dorsal.polymtl.ca/~jdesfossez/babeltrace  
git clone -b live git://git.lttng.org/lttngtop.git
```

*Effici*OS

<http://www.efficios.com>