

Ftrace

Neependra Khare

Fedora Mini-Conf - FOSS.in'10

December 16, 2010

- General tracing
- Tracing for specific events
 - Filtering
- Flight recording mode.
 - Contiguous logging of some selected events

What to Trace?

- Key is to select right tool/event.
 - Understanding of the problem
 - Can take few iteration of general tracing to select most appropriate one
- Use the existing profiler, if available.
- Filtering

Available Tracing Tools

- SystemTap
- Ftrace
- Perf
- Linux Trace tool kit

Specialized tracers

- Blktrace
- Wireshark, tcpdump
- Usbmon
- Bootchart
- ltrace, strace

- Ftrace operates by adding tracepoints to the Linux kernel
- Tracepoints
 - Explicitly declared tracepoints,
 - Implicit tracepoints.

- Kernel Build system add -pg flag.
 - Compiler add code to prologue of each function
 - Calls a specially assembly routine mcount

File system to hold the control files as well as the files to display output.

- Put following in the `/etc/fstab` file

```
debugfs /sys/kernel/debug debugfs defaults 0 0
```

Or you can mount it at run time with:

- `mount -t debugfs nodev /sys/kernel/debug`

Control and Output files

- `current_tracer`
- `available_tracers`
- `tracing_enabled`
- `trace`
- `trace_pipe`
- live tracing, consumer. reads data
- `tracing_max_latency`
- `set_ftrace_filter`
- `set_ftrace_pid`

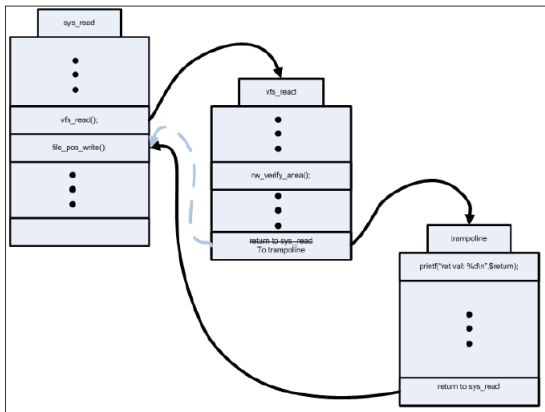
The Tracers

- function
- function_graph
- wakeup and wakeup_rt
- irqsoff, preemptoff, preemtirqsoff
- mmiotrace
- sched_switch
- nop

Reading the Output

- trace
- trace_pipe

The Function Graph Tracer



- binary tool to read Ftrace's buffers
 - Records into a trace.dat file for later reads
 - Reads the trace.dat file
 - Will automatically mount debugfs if it is not mounted

- A front end reader of the trace-cmd trace.dat file
- Graph view
- List view
- Simple and Advance filtering

- <http://lwn.net/Articles/365835/>
- <http://lwn.net/Articles/366796>
- <http://people.redhat.com/srostedt/>
- <http://www.redbooks.ibm.com/redpapers/pdfs/redp4469.pdf>
- <http://ming.vg/pub/tracing>

Question ?