



Frank Ch. Eigler
Software Developer, Red Hat

problem: troubleshooting a live system

- which process is hogging the network?
- which process is swapping?
- what is the maximum/average latency between A and B?
- what fraction of time is some flag set?

problem: troubleshooting a live system

- need more data than /proc-based tools show (vmstat, ...)
- need to integrate data from many sources
- need to be able to answer a specific question
- need not be too painful to use
- need to run on linux
- must not impact subject system noticeably

our solution: systemtap

- simple: awk-like scripting language for events & actions
- expressive: full control structures, \$context data access
- safe: strict error checking, bounded time & space
- fast: compiles to C then to native code, runs within kernel
- organic: scripts reusable, composable in “tapset” layers
- free software from thriving community

a little script

```
probe syscall.* {
    if (execname()=="bash") println("%s %s", name, argstr)
}
probe program("/lib/glibc.so").function("malloc") { # soon
    if ($size > 1024*1024) println("hog!")
}
probe timer.profile {
    hotfuncs[execname(), probefunc()] <<< 1
}
global hotfuncs
```

recently added features

- caching recently compiled scripts => run faster
- cross- and client-server compilation => install simpler
- debuginfo-less operation => install smaller
- kernel marker probing => probe faster
- utrace user-space probing => probe more/faster
- semi-privileged users run approved scripts => more users
- production support under RHEL

coming soon

- proper symbolic user-space probing => probe much more
- unprivileged users => casually probe own programs
- hardware performance counters => detailed profiling

try it

```
# yum install systemtap kernel-devel
```

```
# debuginfo-install kernel
```

(or boot from demonstration live-cd)

```
# cd /usr/share/doc/systemtap*/examples
```

further information

- questions?
- <http://sources.redhat.com/systemtap>
tutorials, code, fedora-based live-cd
- .../wiki : samples, tips
- systemtap@sources.redhat.com

systemtap

