Discover what your app is waiting with kernel tracing

Tracing Summit 2013

Francis Giraldeau
francis.giraldeau@gmail.com

DORSAL Lab, École Polytechnique de Montréal

apt-get install tree

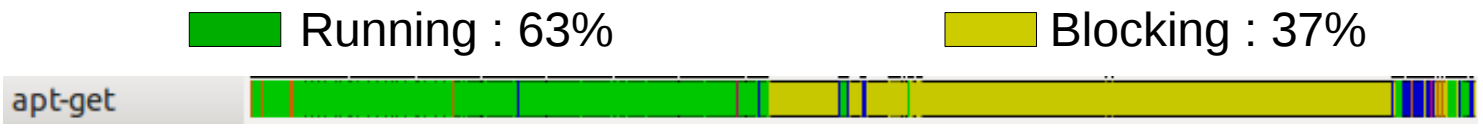

Incl.	Self	Called	Function	Location
100.00	0.00	(0)	0x0000000000001650	ld-2.17.so
99.87	0.00	1	0x000000000040a5a8	apt-get
99.87	0.00	1	(below main)	libc-2.17.so: libc-start.c
99.87	0.00	1	0x000000000040a0b0	apt-get
99.72	0.00	1	CommandLine::DispatchArg(CommandLin…	libapt-pkg.so.4.12.0
99.72	0.13	1	0x0000000000416ca0	apt-get
72.20	0.00	1	0x00000000004251e0	apt-get
63.28	0.00	1	pkgCacheFile::Open(OpProgress*, bool)	libapt-pkg.so.4.12.0
62.65	0.00	3	pkgCacheFile::BuildDepCache(OpProgress*)	libapt-pkg.so.4.12.0
62.65	0.51	1	pkgDepCache::Init(OpProgress*)	libapt-pkg.so.4.12.0
53.46	2.22	1	pkgDepCache::Update(OpProgress*)	libapt-pkg.so.4.12.0
41.23	3.66	609 598	pkgDepCache::DependencyState(pkgCac…	libapt-pkg.so.4.12.0
37.57	10.16	1 828 794	pkgDepCache::CheckDep(pkgCache::DepI…	libapt-pkg.so.4.12.0
26.49	2.24	1 130 829	debVersioningSystem::CheckDep(char con…	libapt-pkg.so.4.12.0

Traditional profiler output

What the app is waiting for?

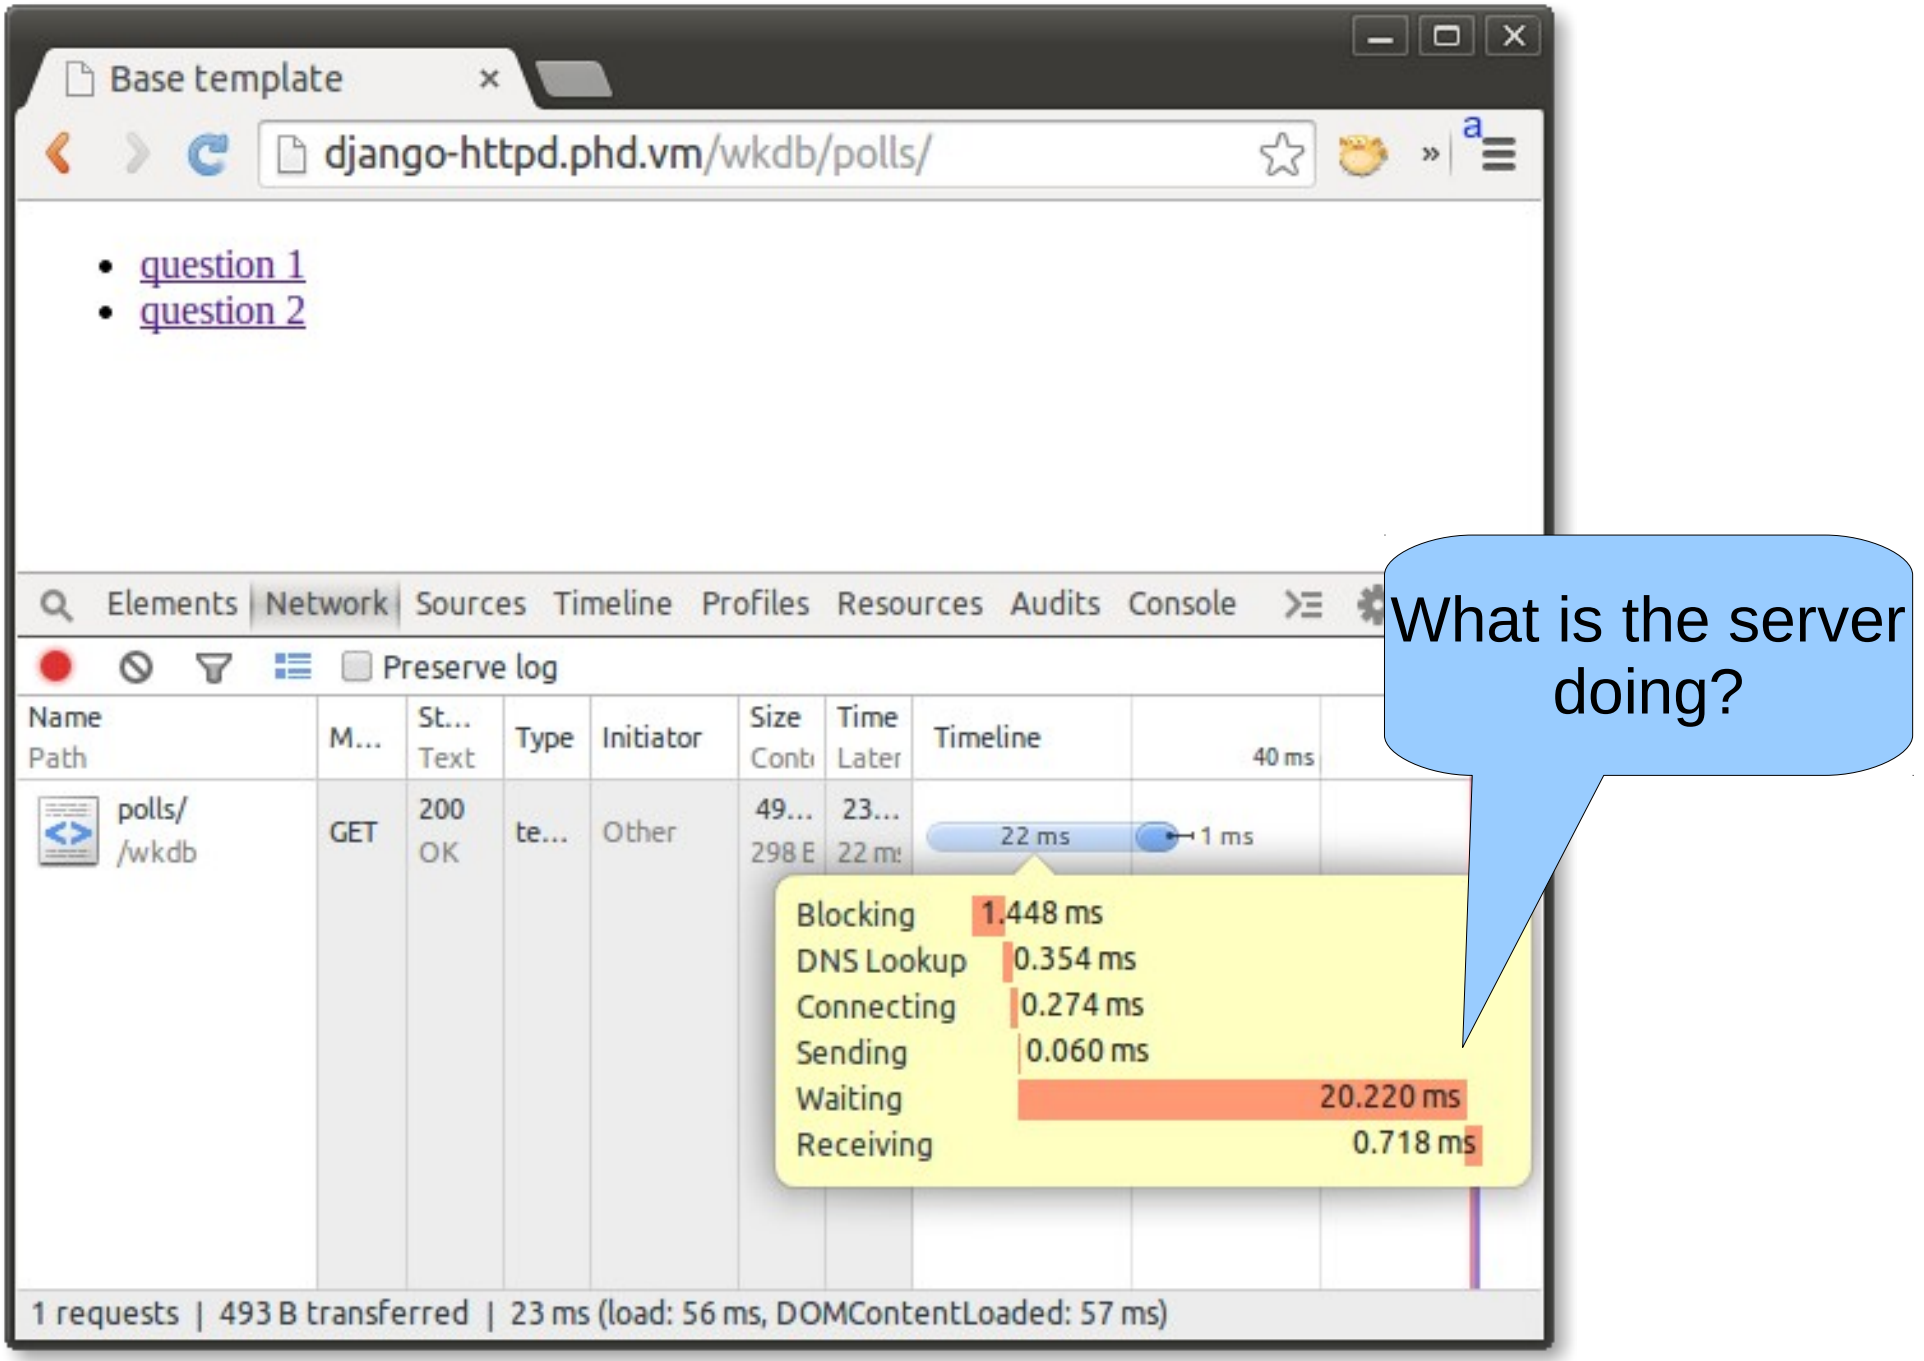
Base template
django-httpd.phd.vm/wkdb/polls/
question 1
question 2
Elements
Network
Sources
Timeline
Profiles
Resources
Audits
Console
Preserve log
Name
Path
M...
St...
Text
Type
Initiator
Size
Cont
Time
Later
Timeline
40 ms
polls/
/wkdb
GET
200
OK
te...
Other
49...
298 E
23...
22 ms
22 ms
1 ms
Blocking 1.448 ms
DNS Lookup 0.354 ms
Connecting 0.274 ms
Sending 0.060 ms
Waiting 20.220 ms
Receiving 0.718 ms
1 requests | 493 B transferred | 23 ms (load: 56 ms, DOMContentLoaded: 57 ms)
What is the server doing?

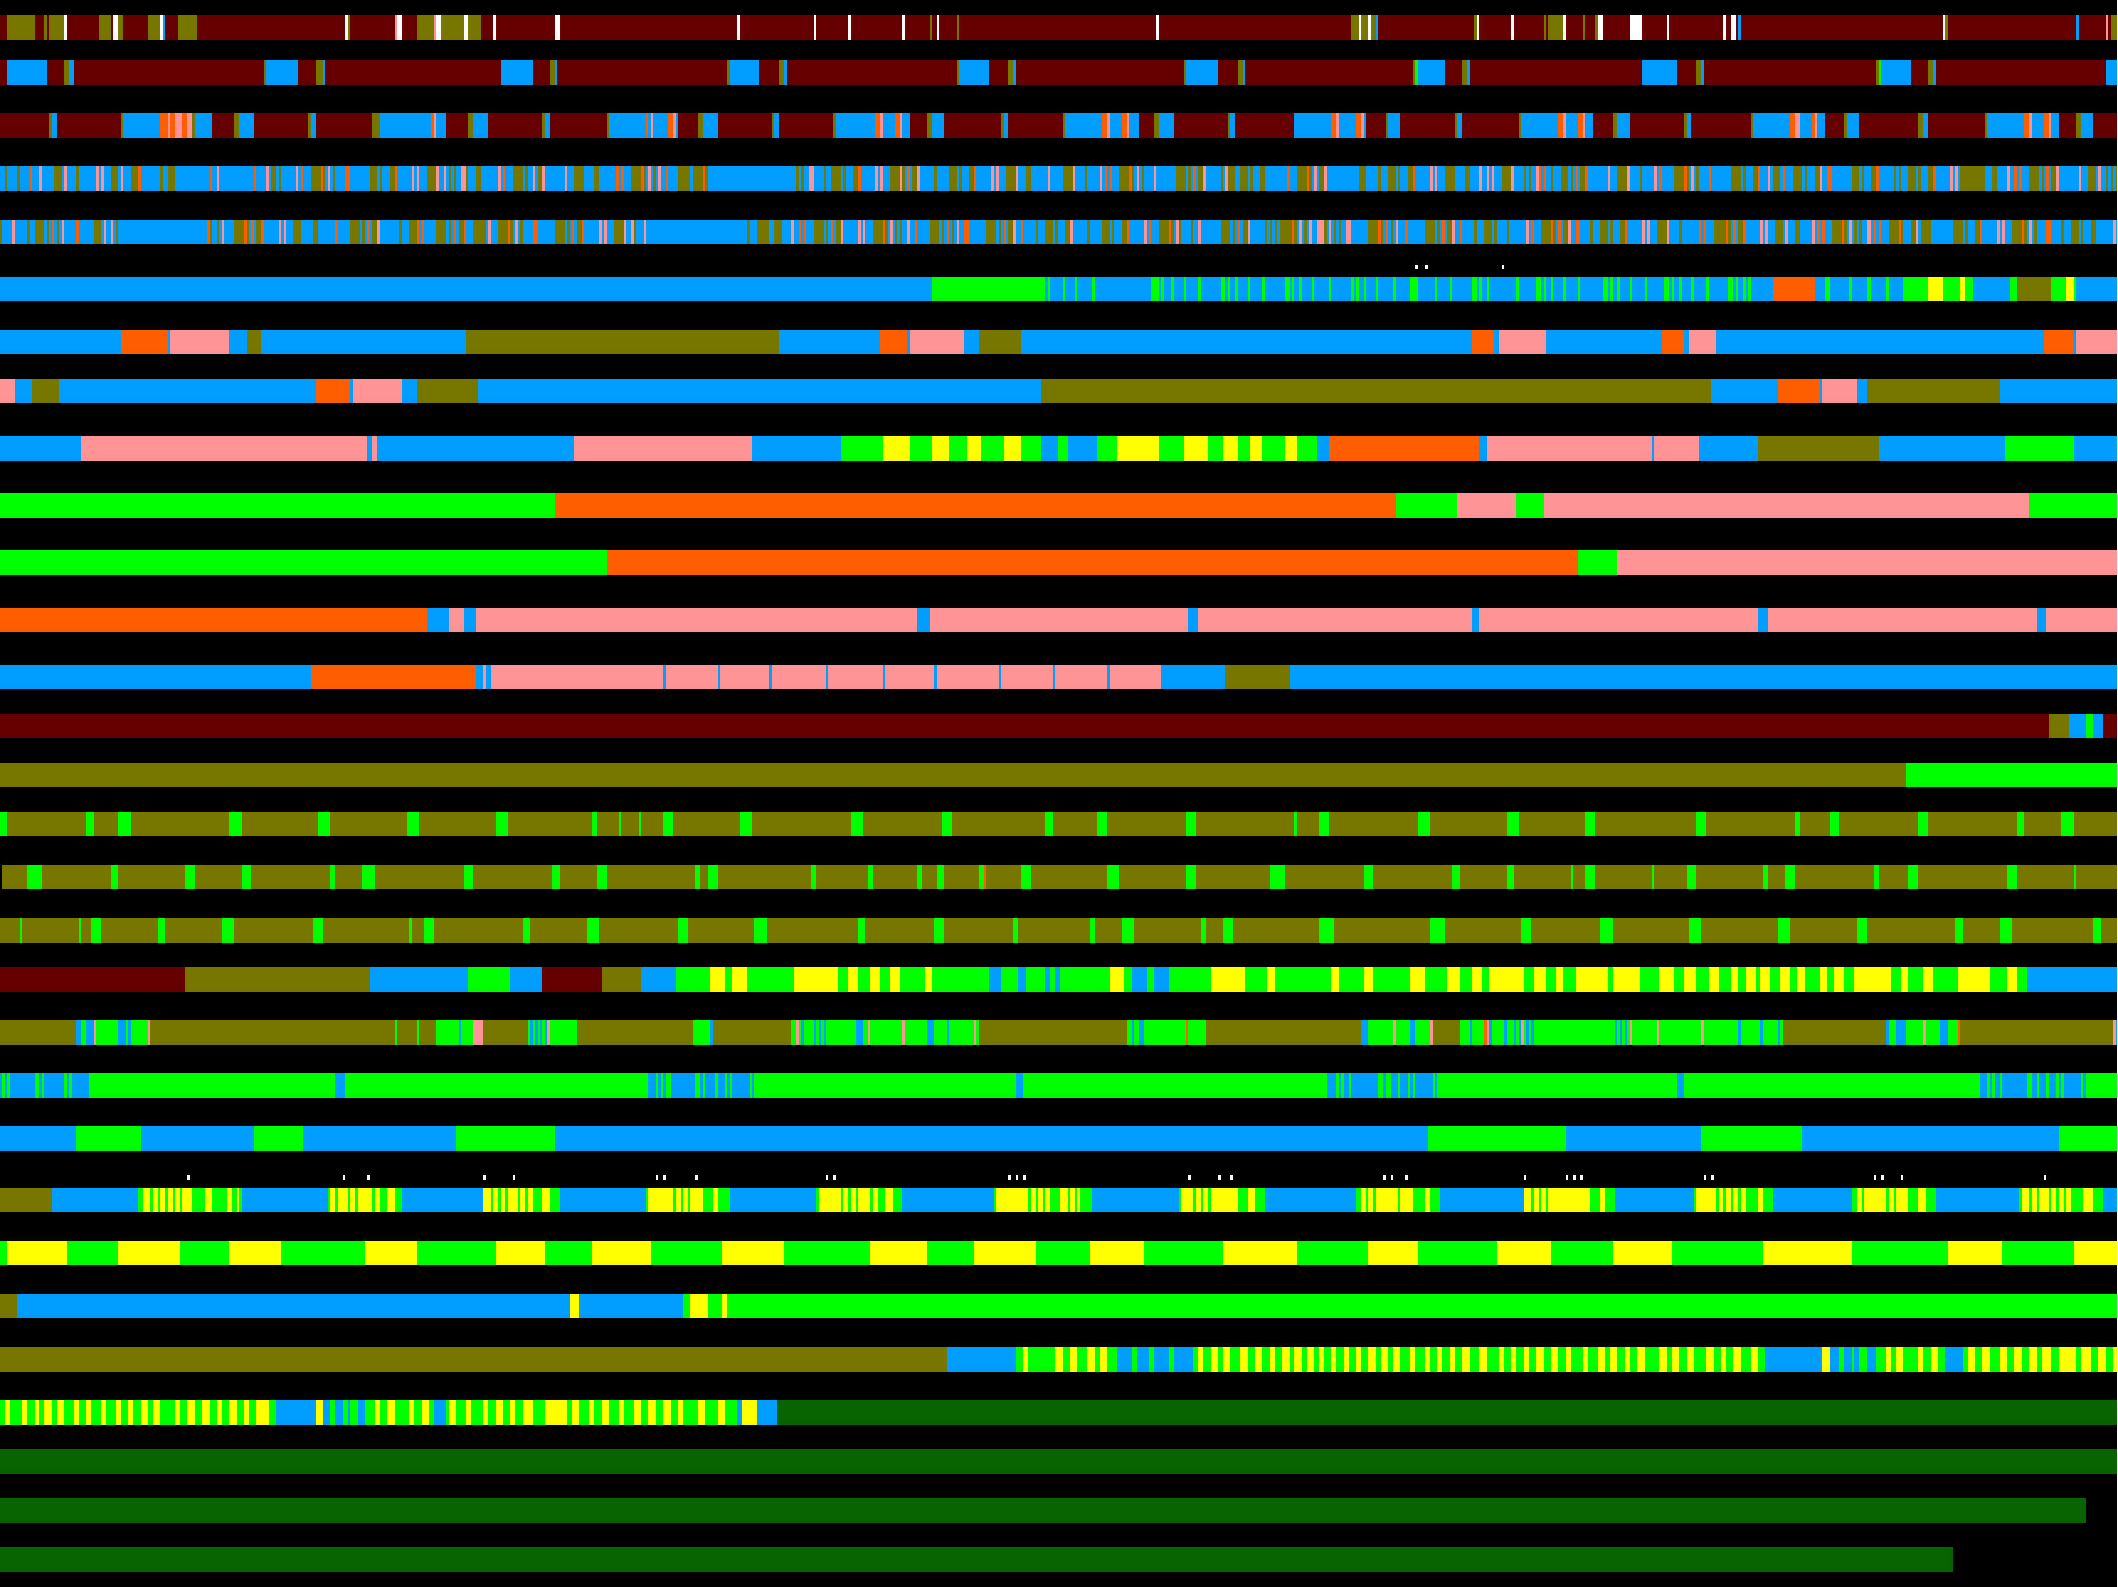

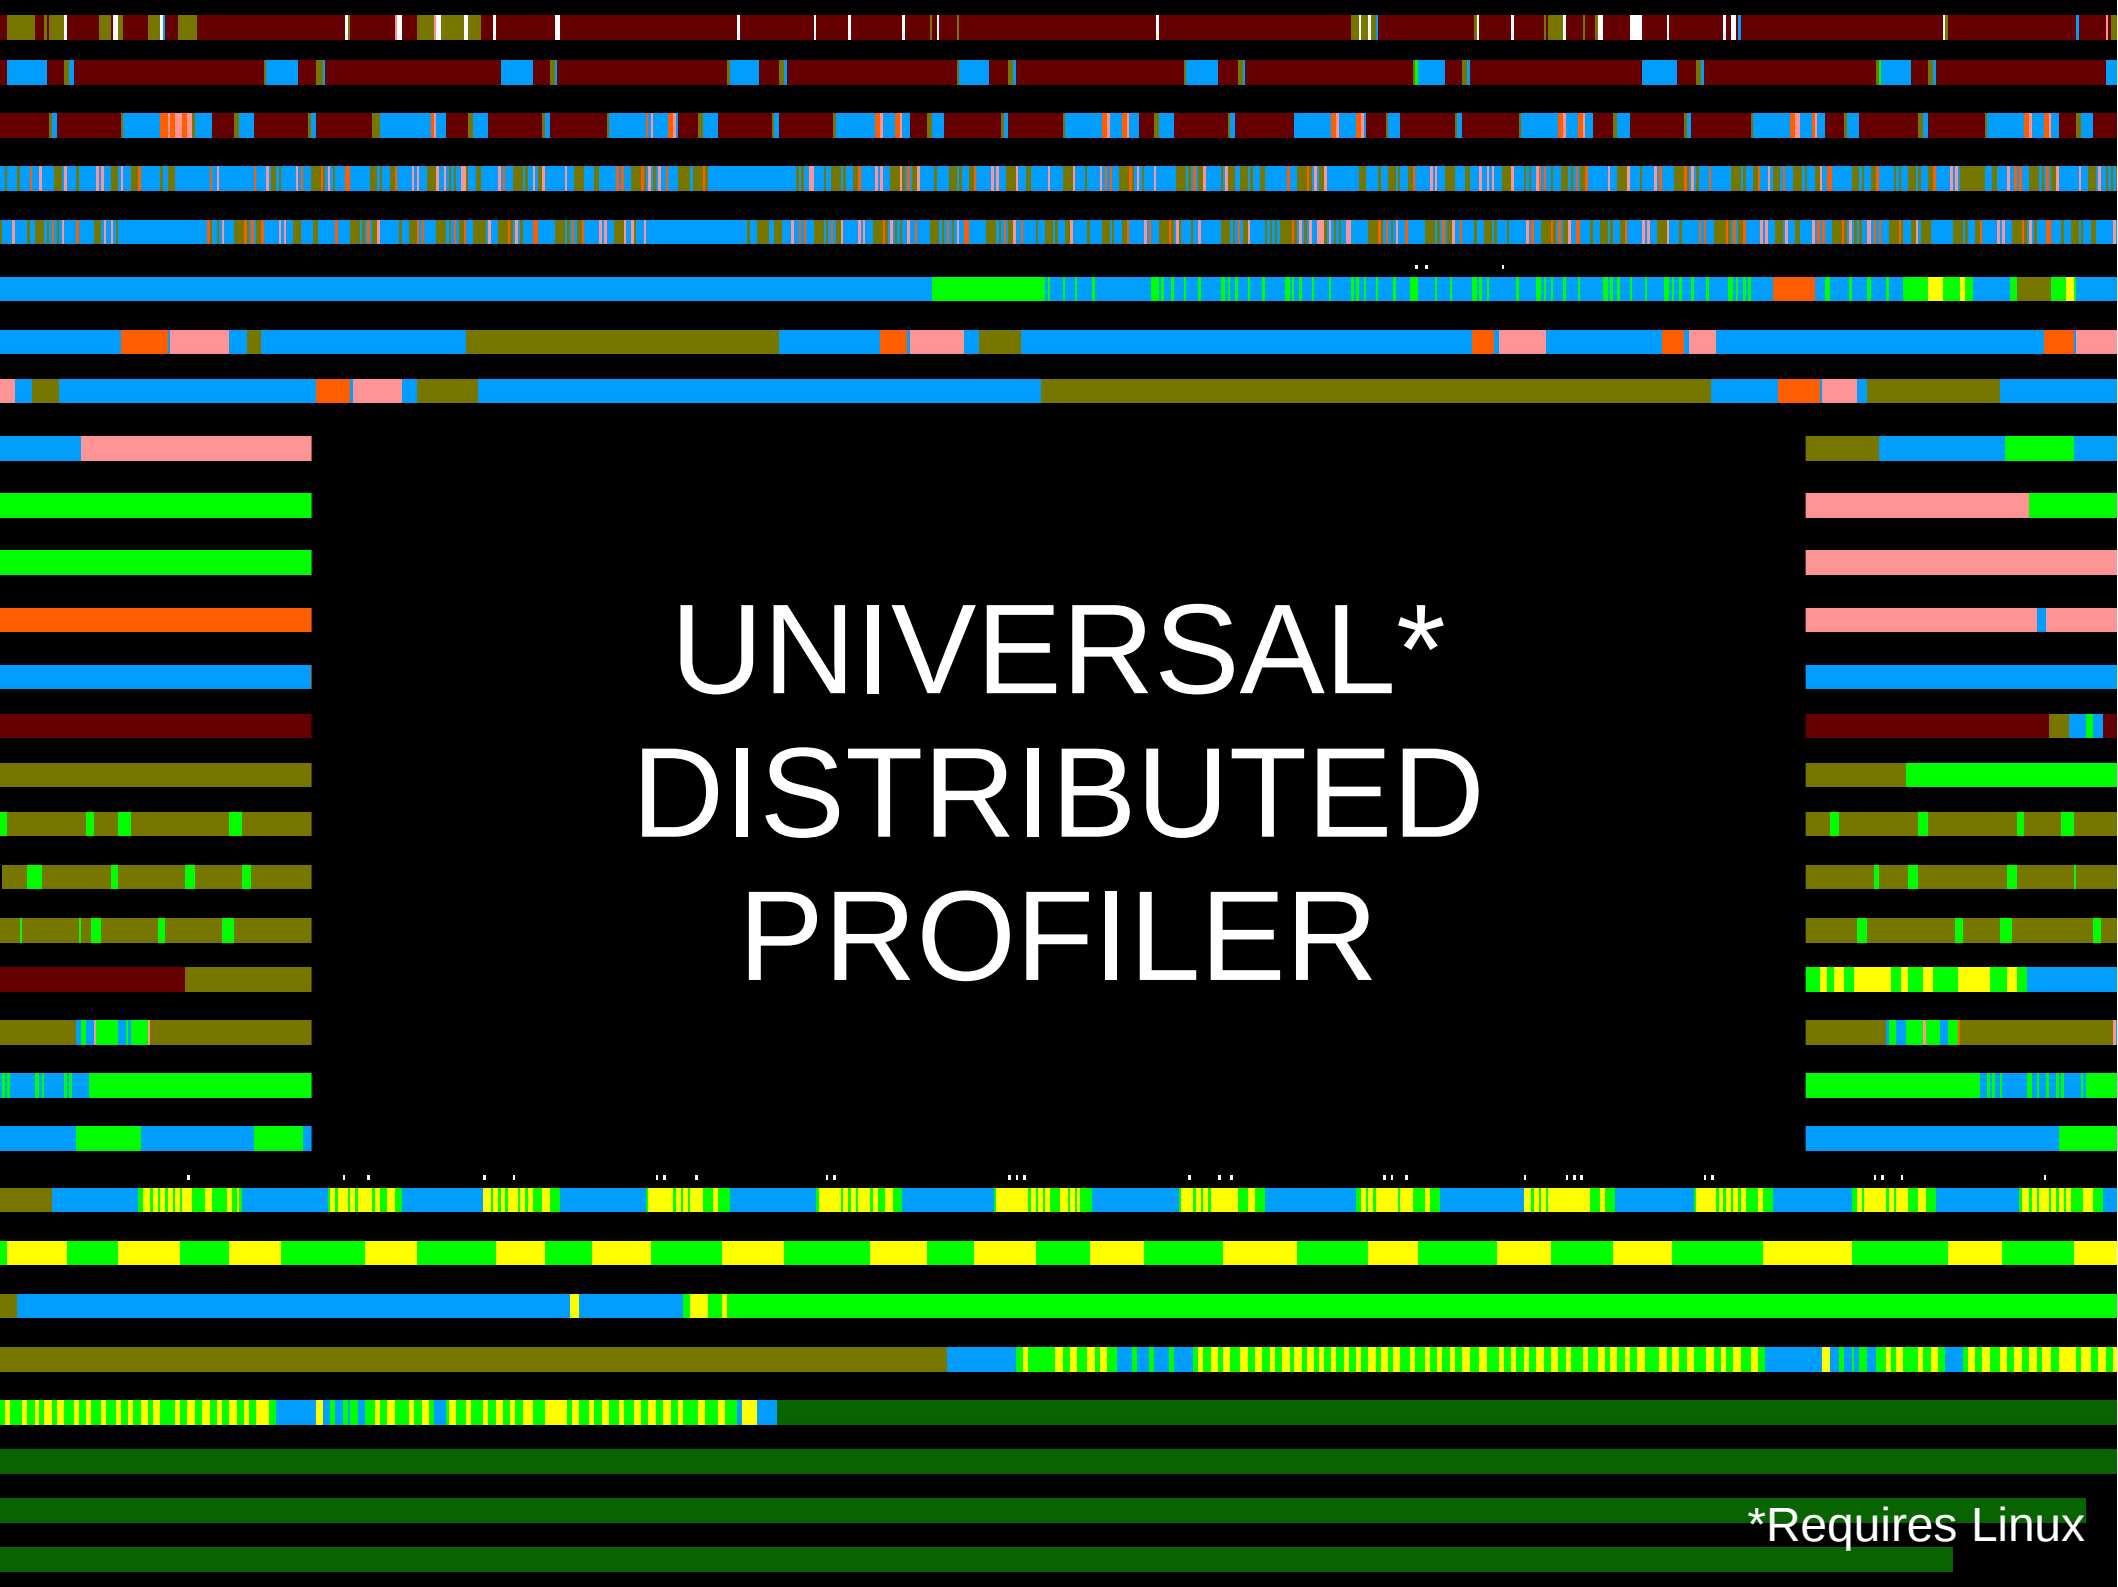
UNIVERSAL*
DISTRIBUTED
PROFILER
*Requires Linux

Task state

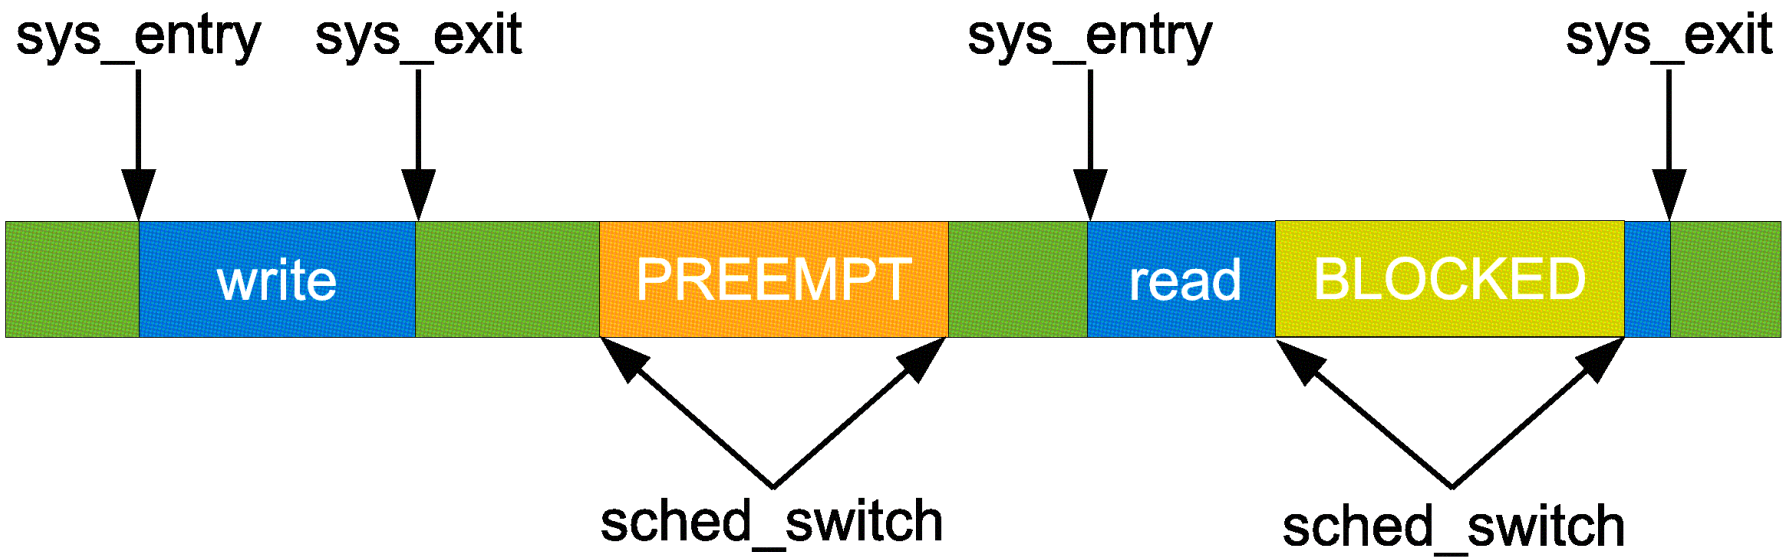

TASK_INTERRUPTIBLE
TASK_UNINTERRUPTIBLE

Device wake-up

Task wake-up

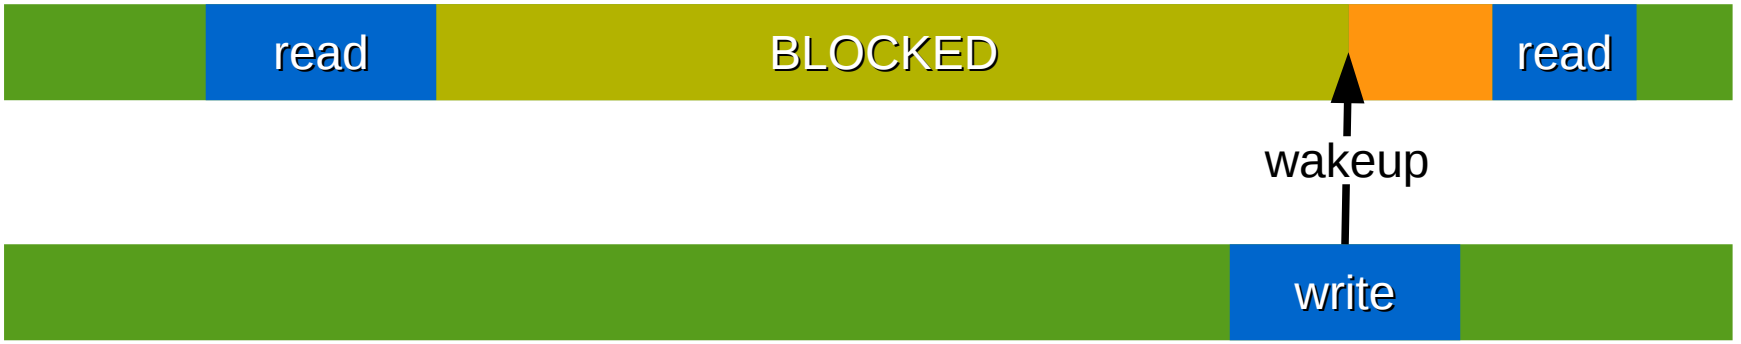

Remote wake-up

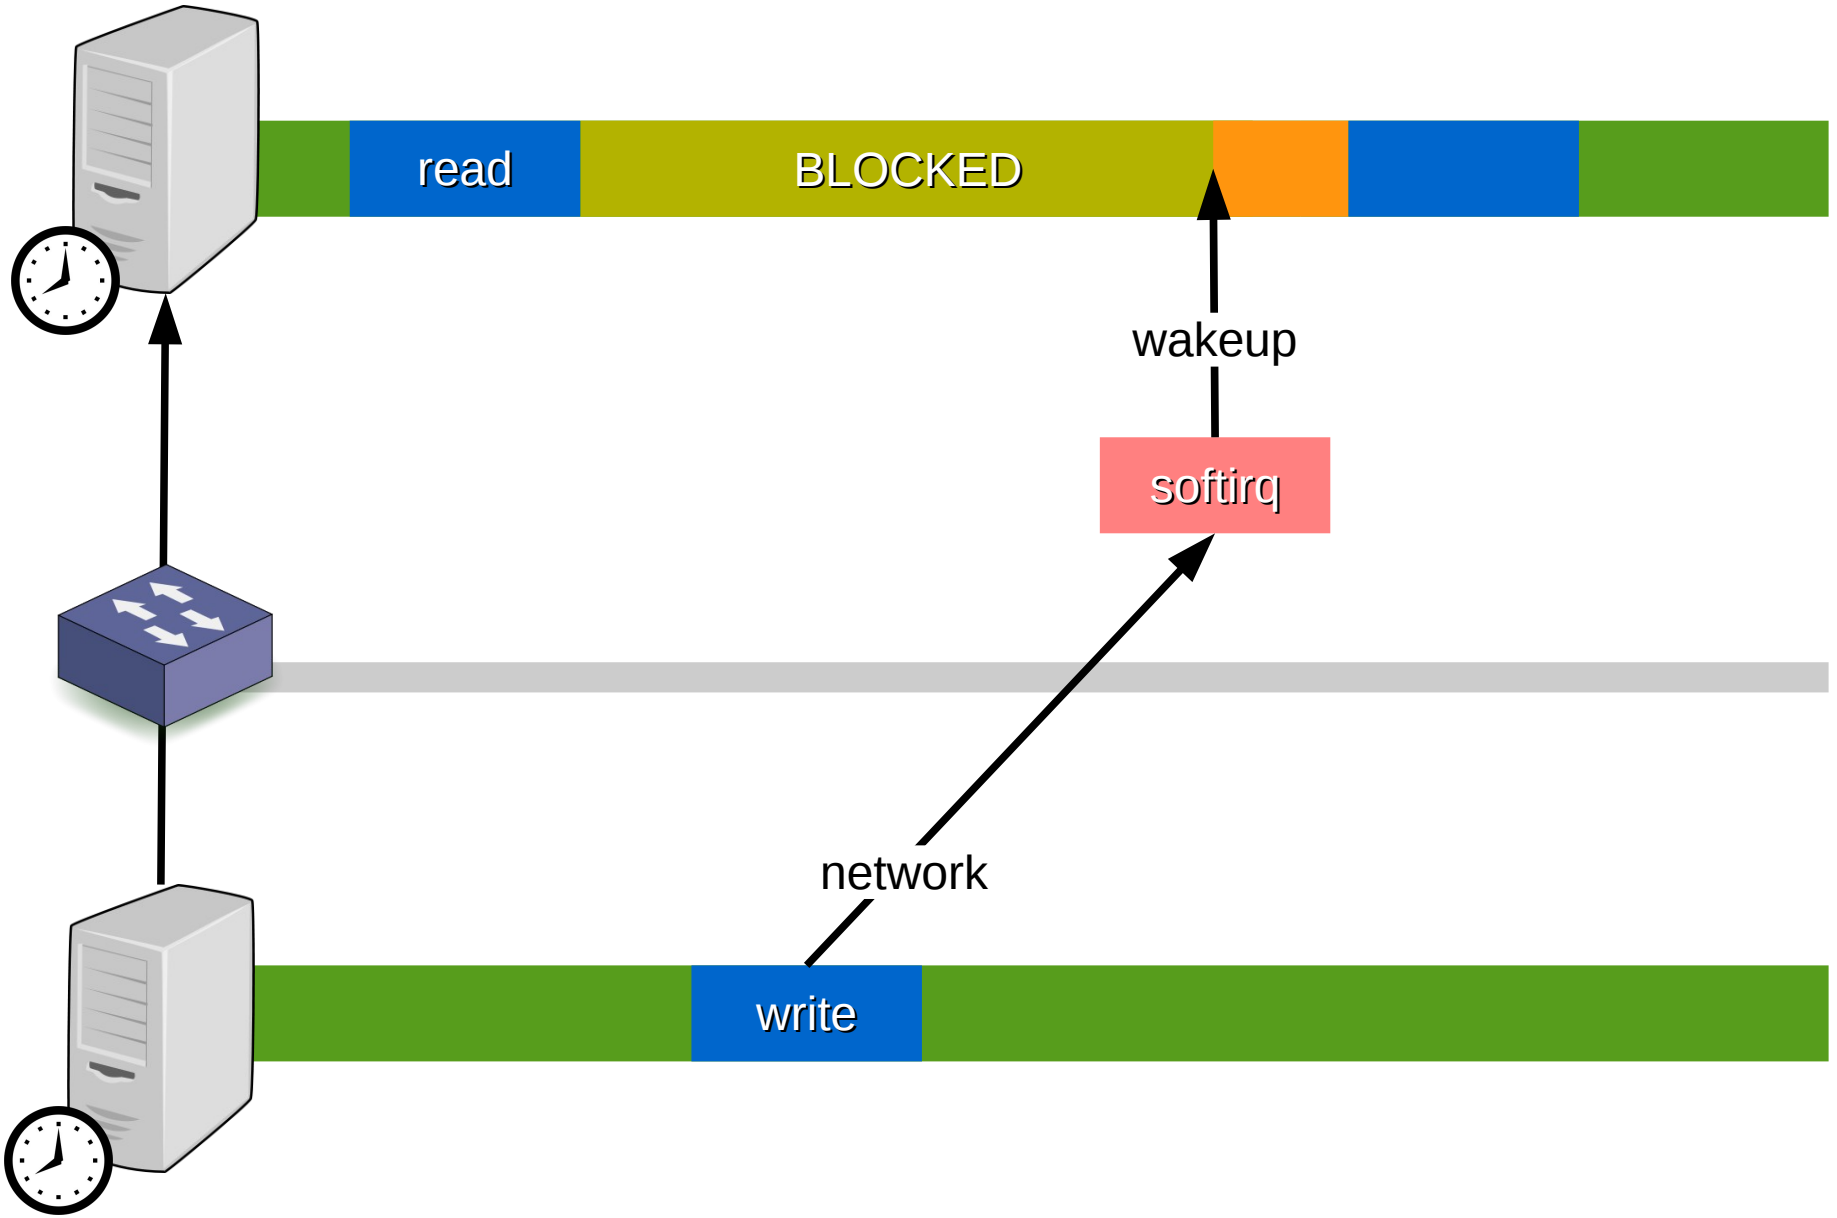

Required events

- sched_switch
- sched_wakeup
- irq_handler_{entry, exit}
- hrtimer_expire_{entry, exit}
- softirq_{entry, exit}
- inet_sock_local_{in, out}

System calls not required!

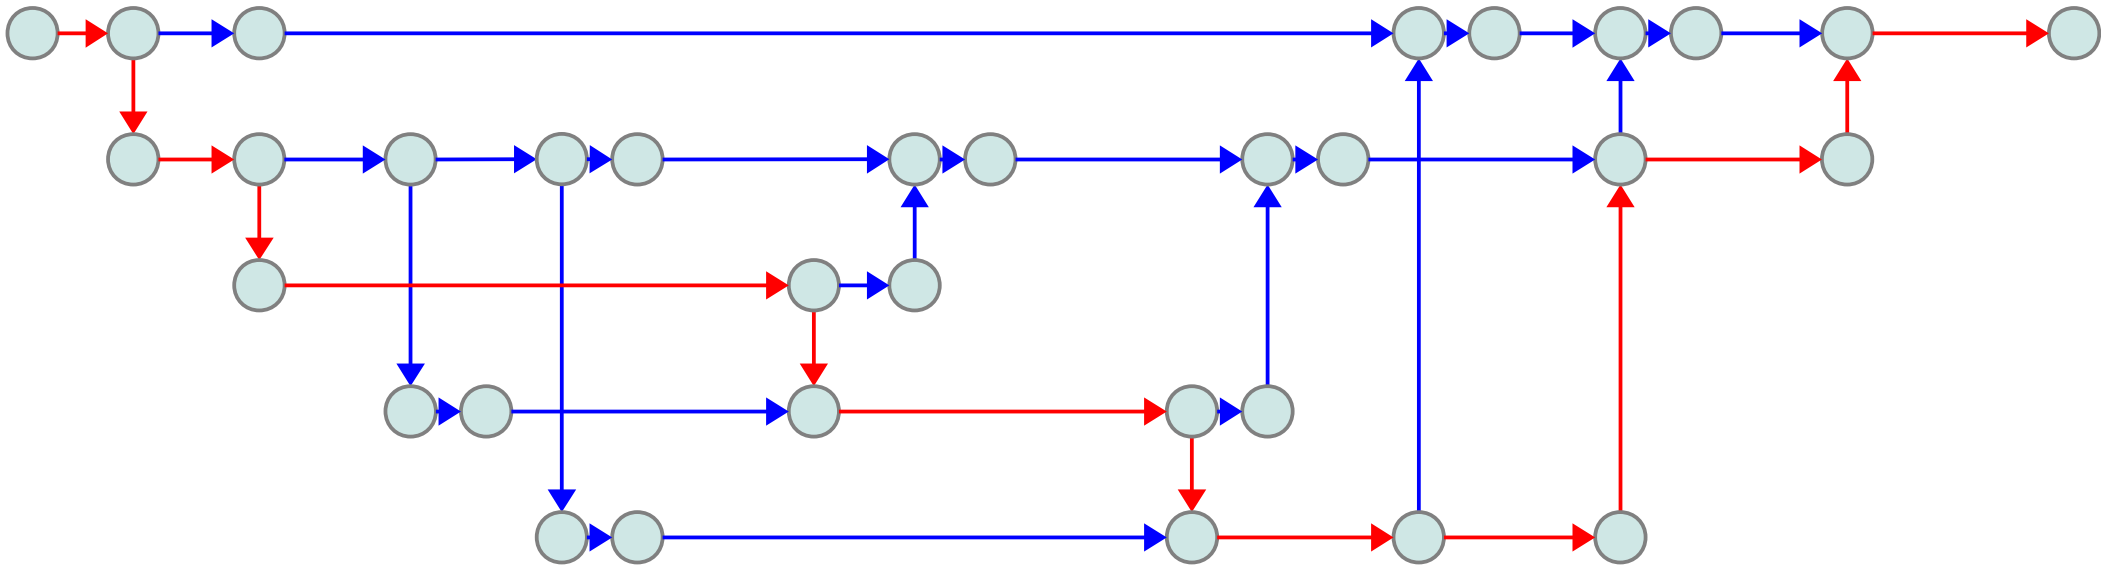

Analysis overview

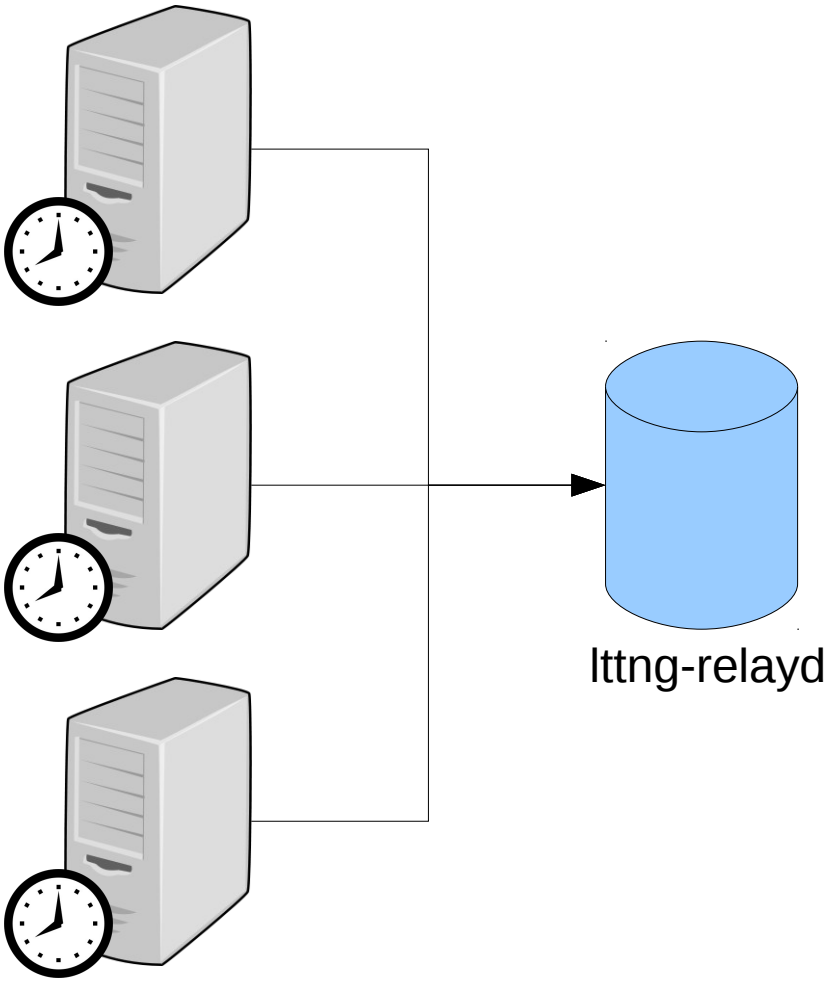

Distributed tracing
1) Start relayd
2) Start tracing on each machine
3) Execute the load
4) Stop tracing

Off-line analysis in TMF
1) Create experiment
2) Synchronize traces
3) Compute the execution graph
4) Compute the critical path

Demo

Conclusion

- Crossing machines boundaries works
- Trace synchronization appropriate
- Shed light into complex, actual executions
- Precise and universal tool, independent from user-space runtime

Future work

- Performance counters of the active path
- Filter lttng-relayd packets for streaming
- More scalable analysis using state history

Thanks to Professor Michel Dagenais and our partners EfficiOS and Ericsson.

Special thanks to Geneviève Bastien for her excellent work on Luna Dorsal.

More information on:

http://step.dorsal.polymtl.ca/~fgiraldeau

Annexes

RPC Server

Commands: hog

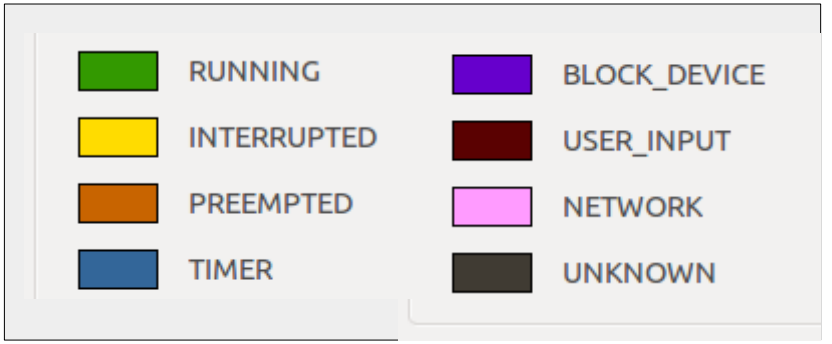

Default control flow view

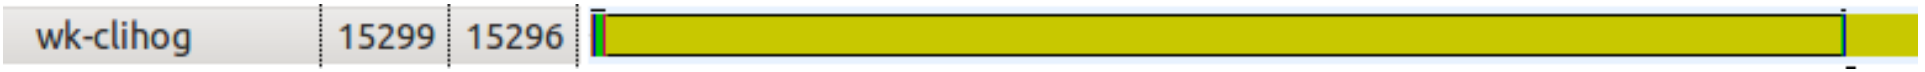

Critical Flow View : request hog()

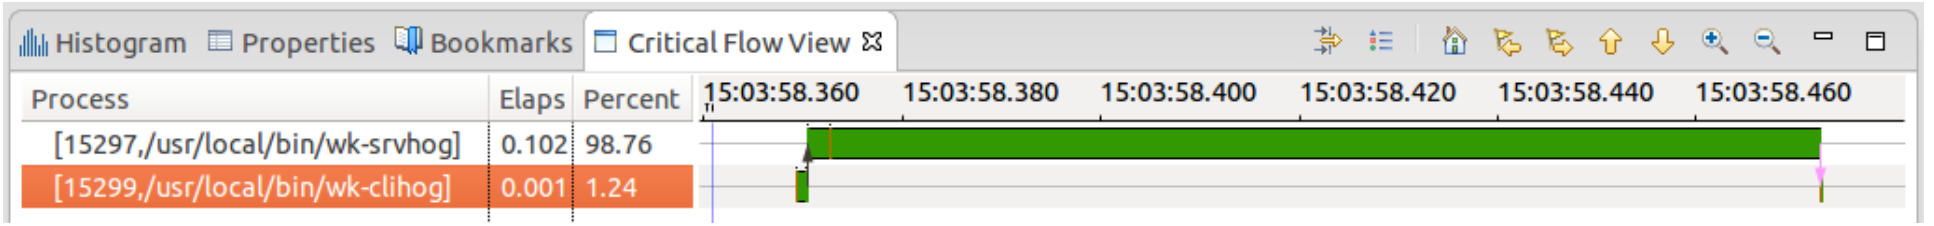

Critical Flow View : request hog()

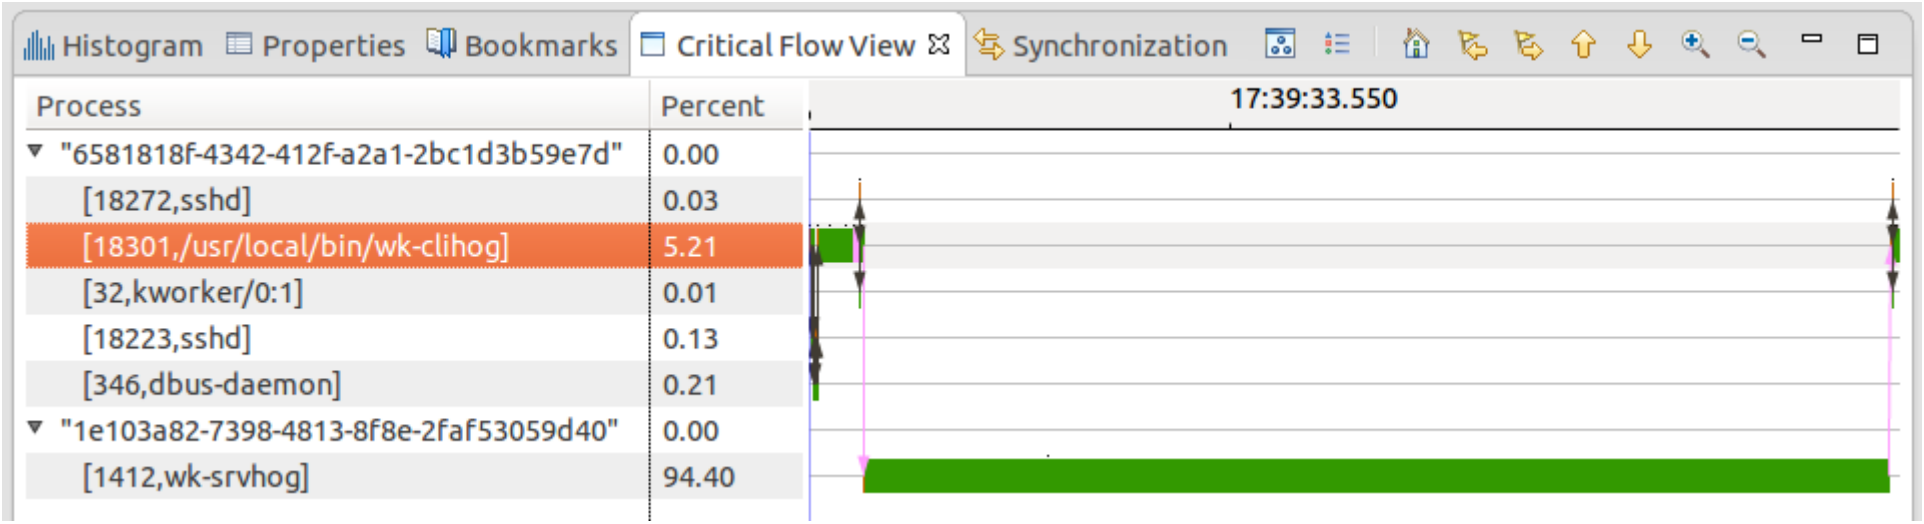

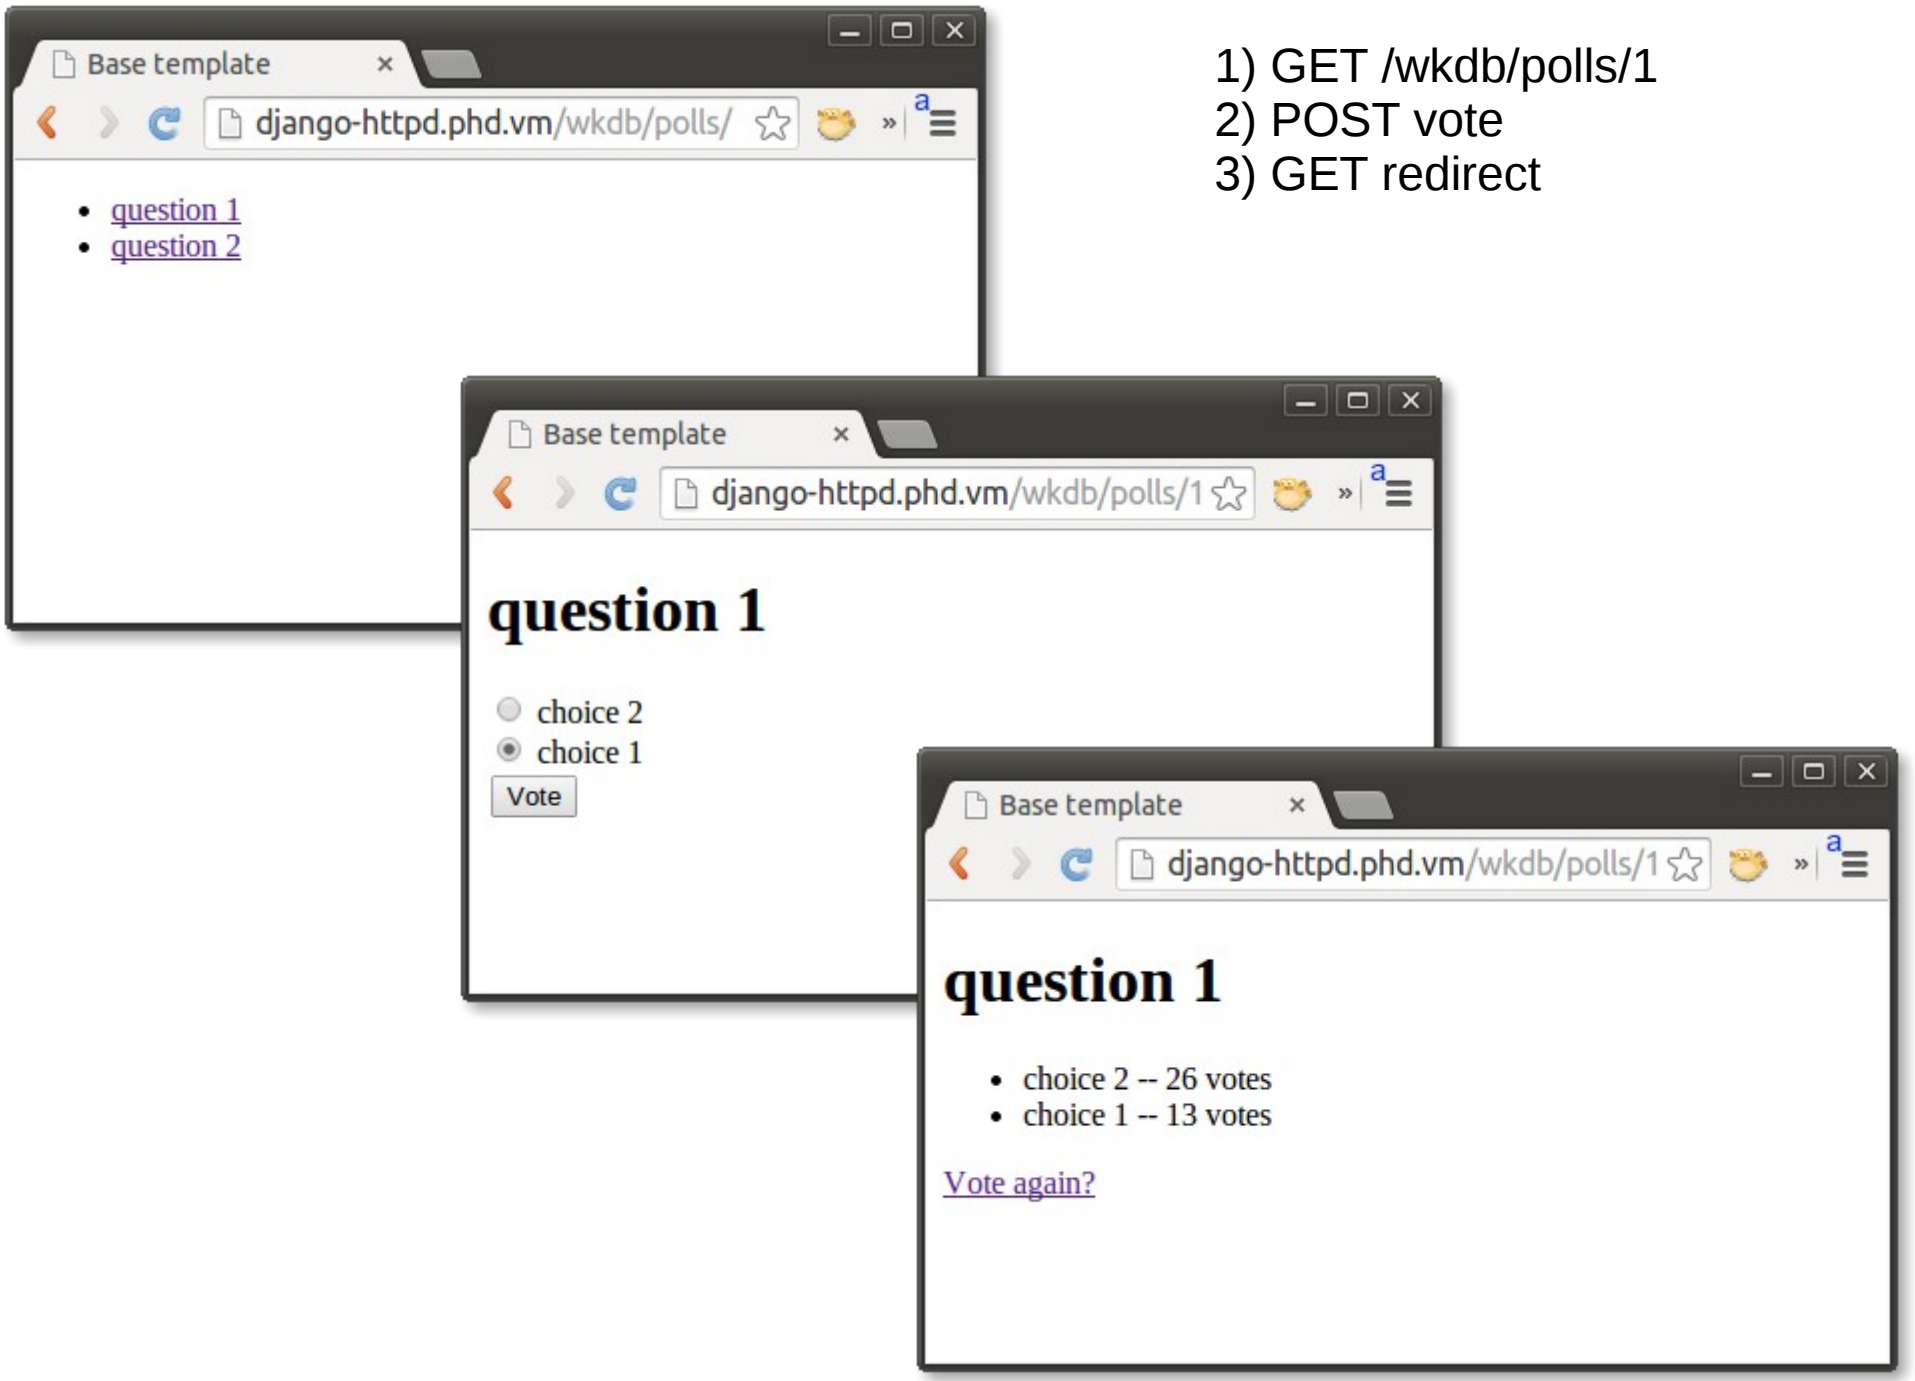

1) GET /wkdb/polls/1
2) POST vote
3) GET redirect

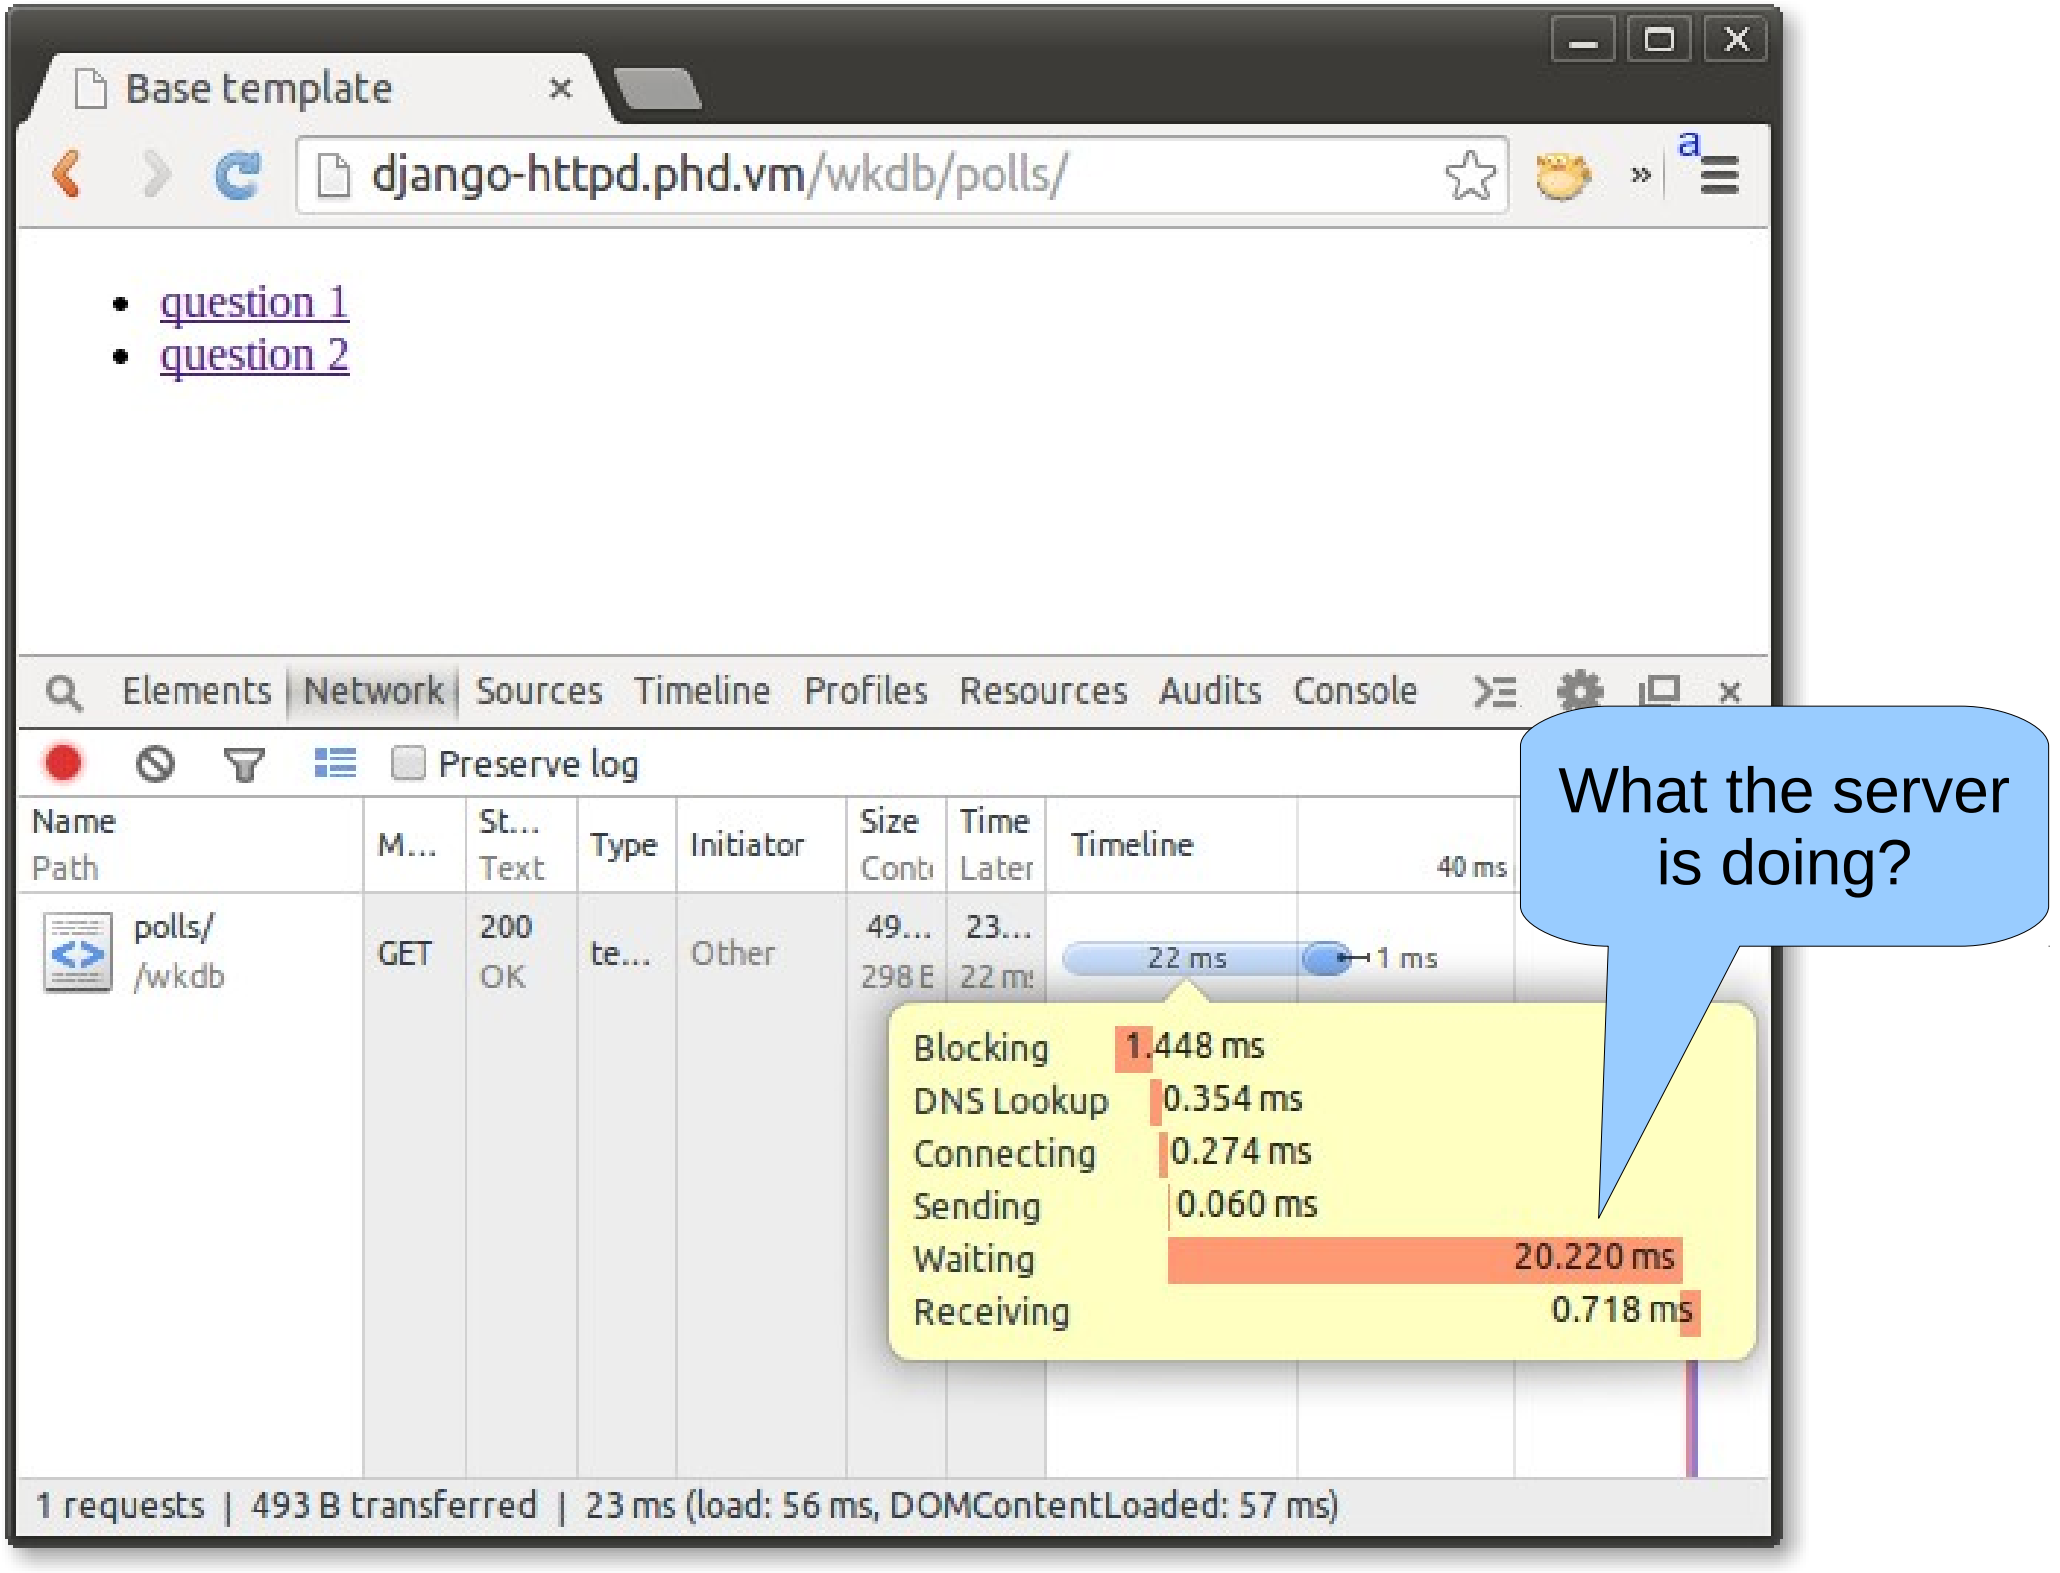
Base template
django-httpd.phd.vm/wkdb/polls/
question 1
question 2
Elements
Network
Sources
Timeline
Profiles
Resources
Audits
Console
Preserve log
Name
Path
M...
St...
Text
Type
Initiator
Size
Cont
Time
Later
Timeline
40 ms
polls/
/wkdb
GET
200
OK
te...
Other
49...
298 B
23...
22 ms
22 ms
1 ms
Blocking
1.448 ms
DNS Lookup
0.354 ms
Connecting
0.274 ms
Sending
0.060 ms
Waiting
20.220 ms
Receiving
0.718 ms
What the server is doing?
1 requests | 493 B transferred | 23 ms (load: 56 ms, DOMContentLoaded: 57 ms)

Django Polls App

mechanize + Apache WSGI + postgresql

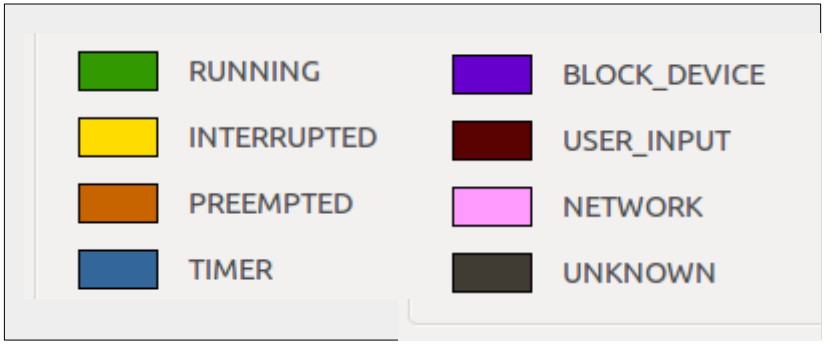

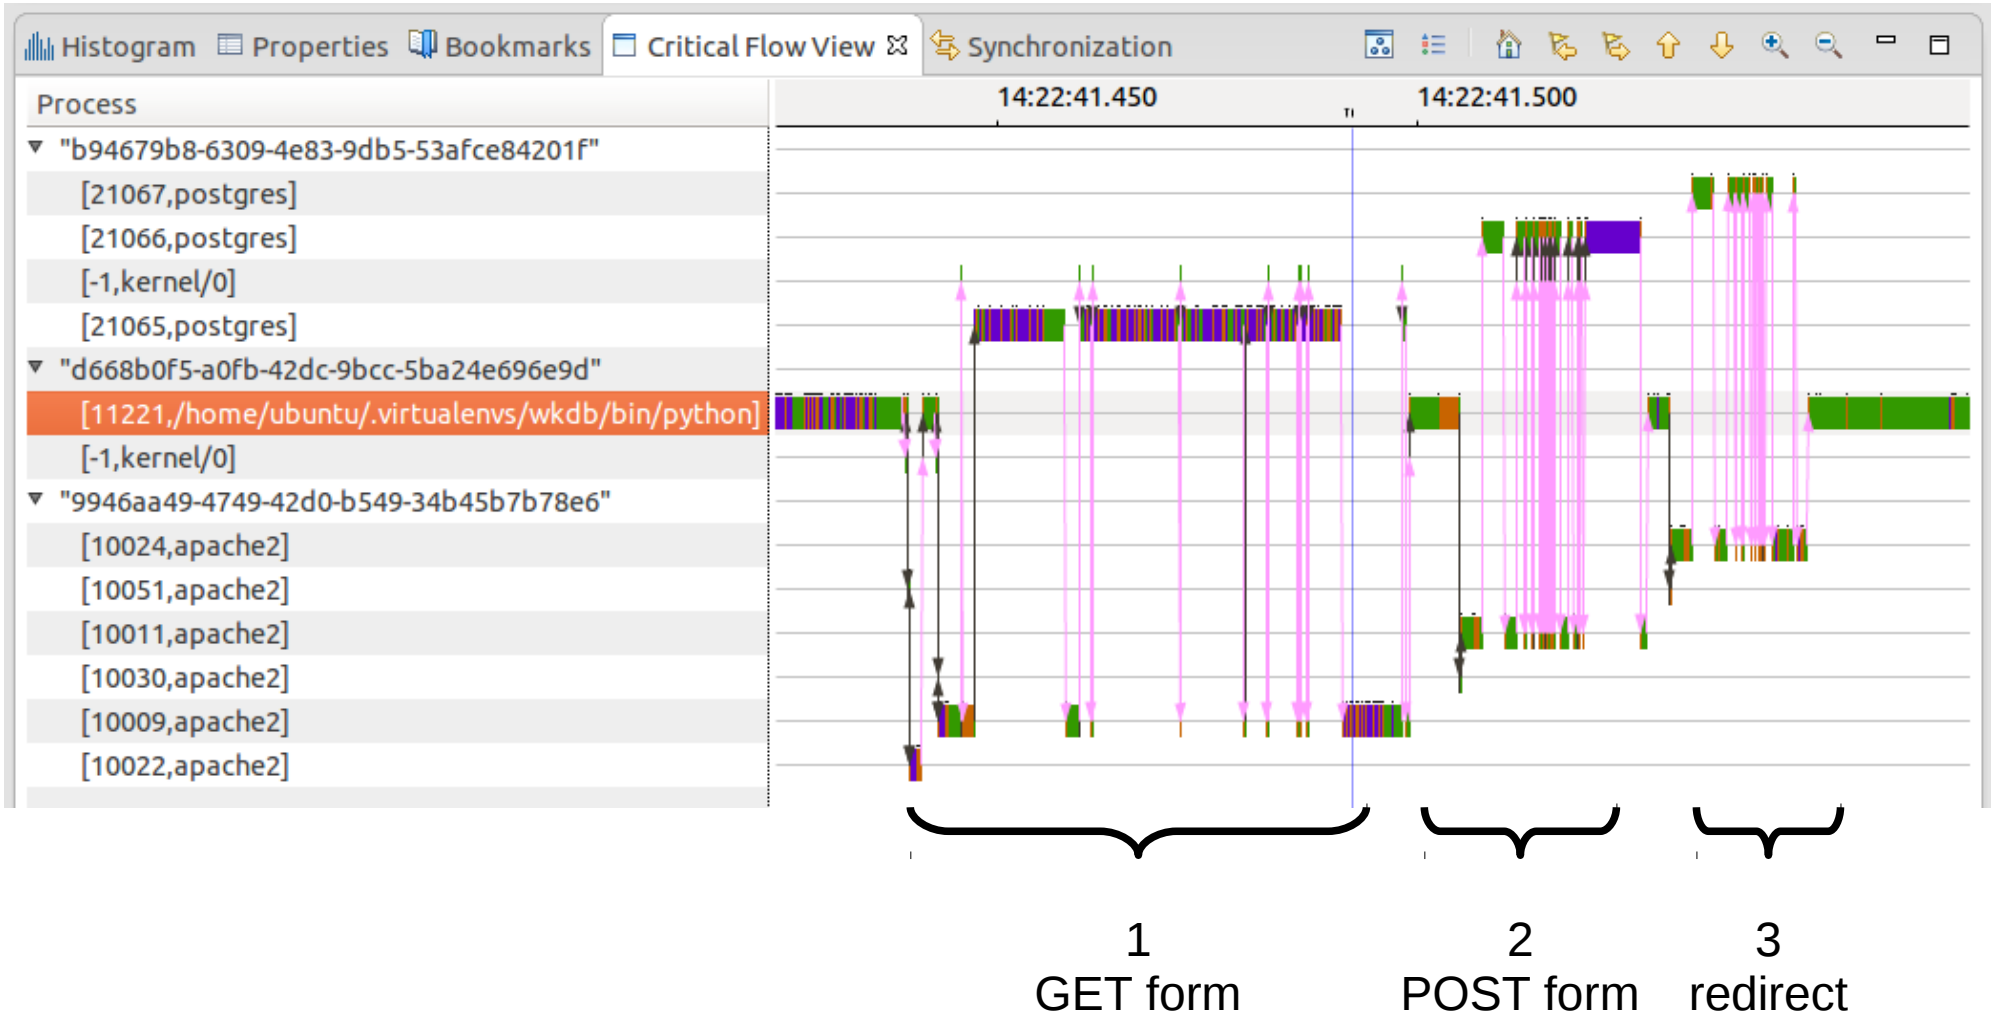

Django Polls App

mechanize + Apache WSGI + postgresql

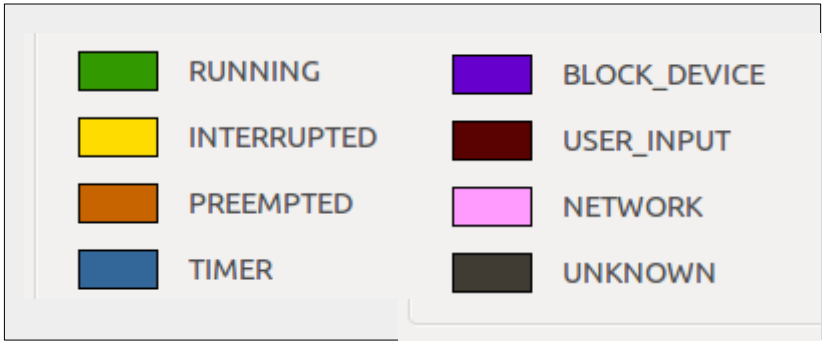

1) GET form

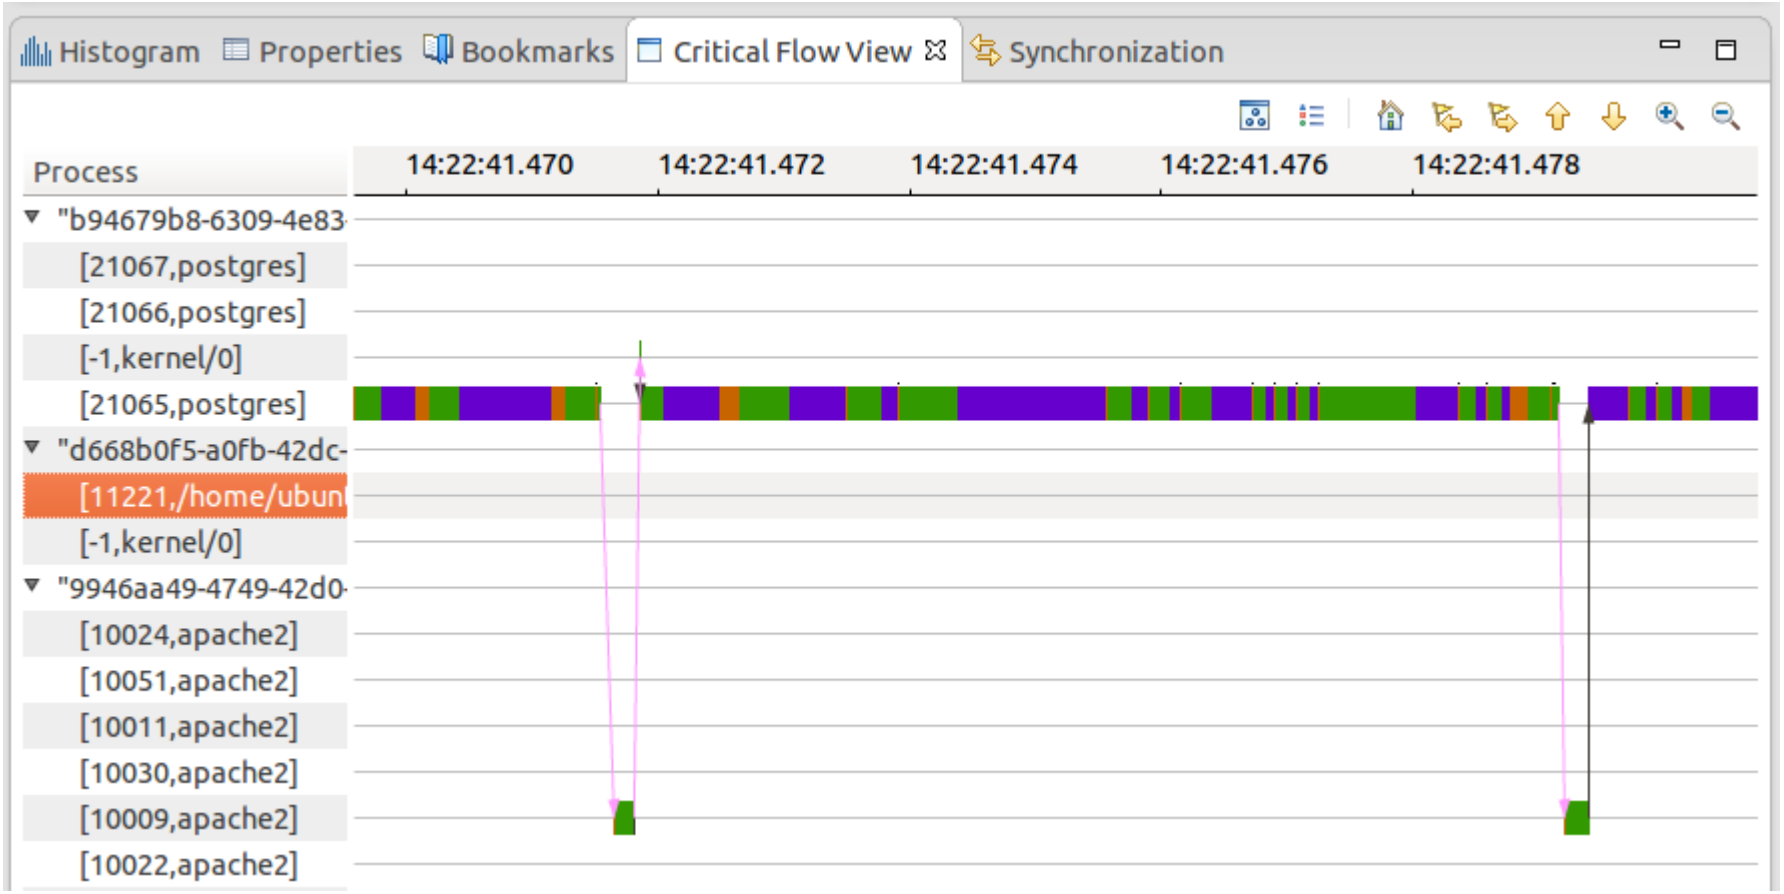

Django Polls App

mechanize + Apache WSGI + postgresql

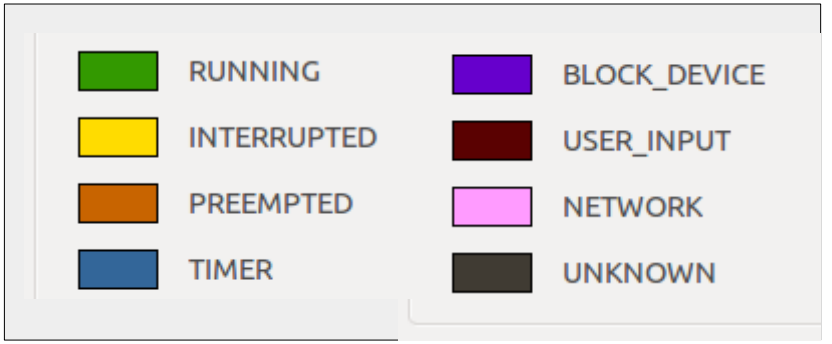

2) POST form

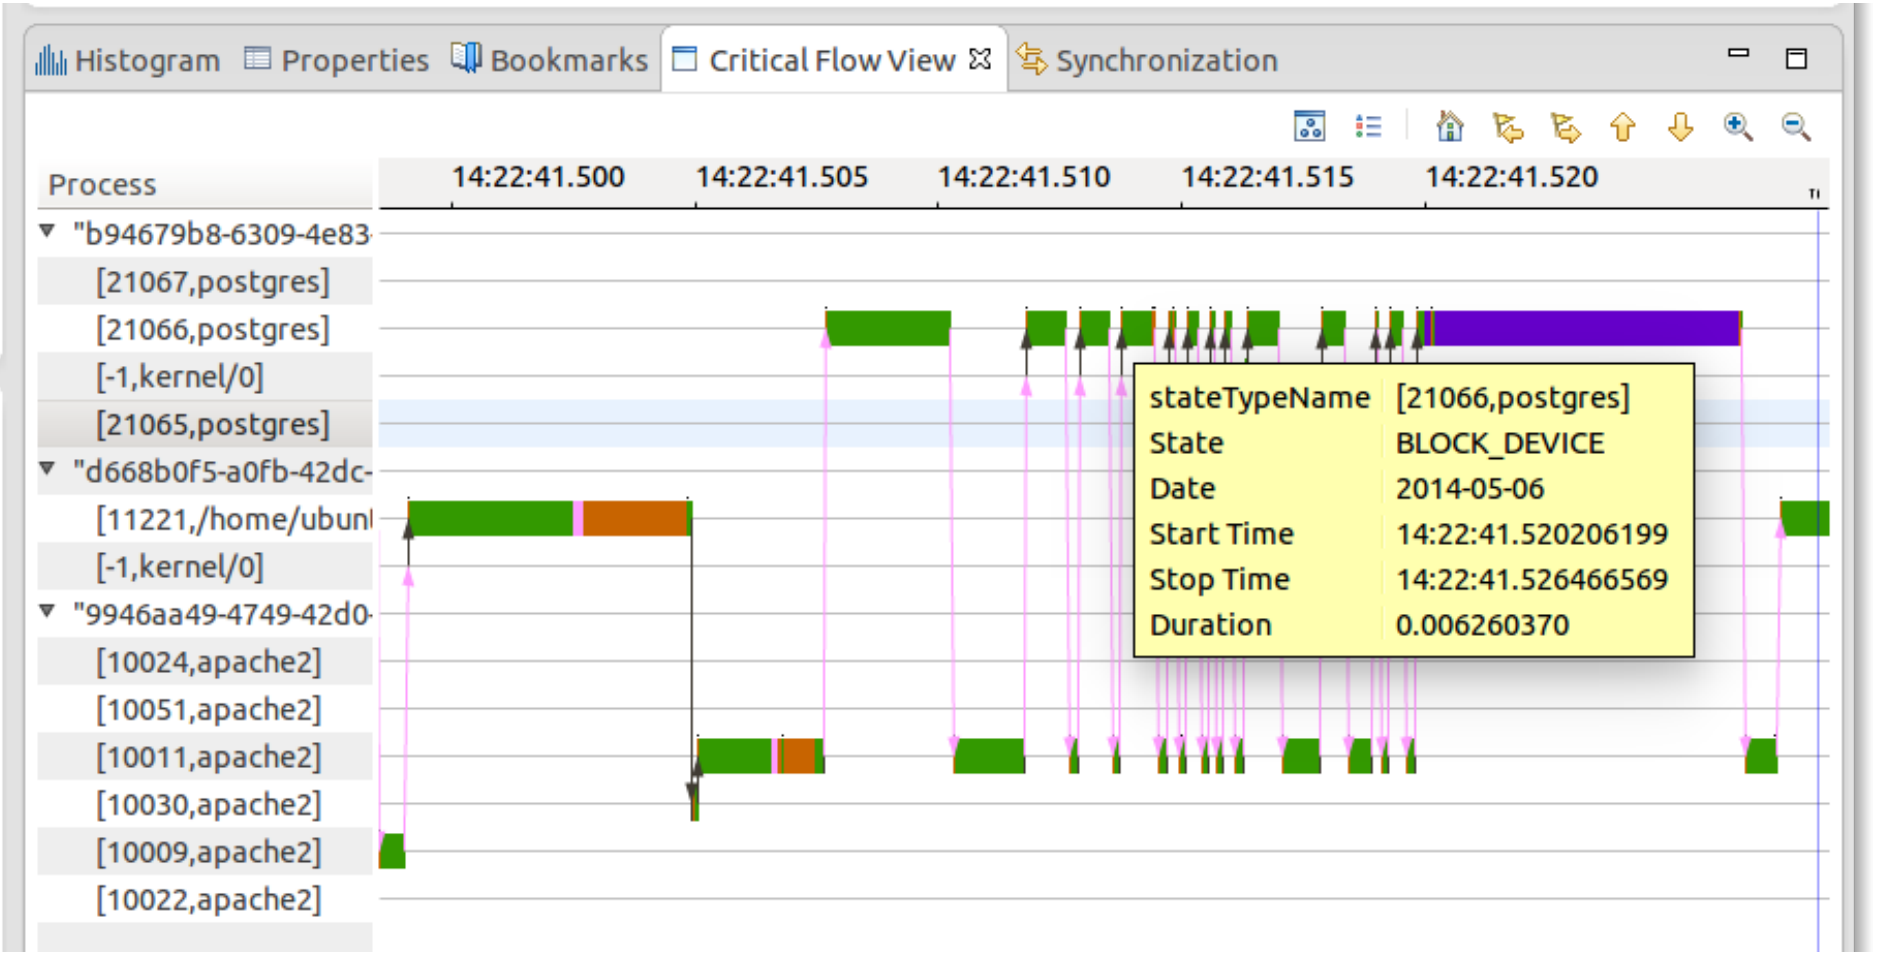

Django Polls App

mechanize + Apache WSGI + postgresql

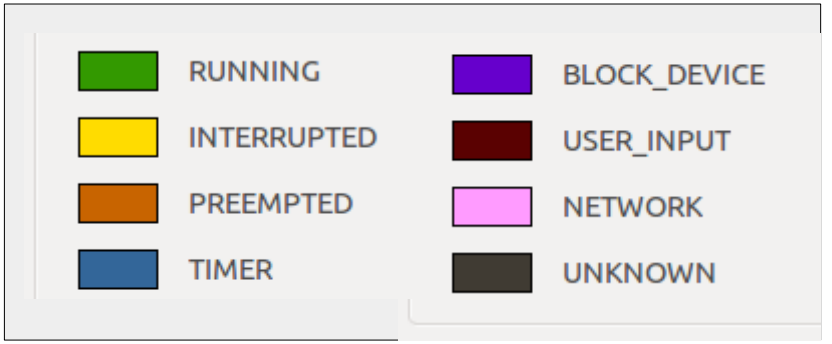

3) GET redirect

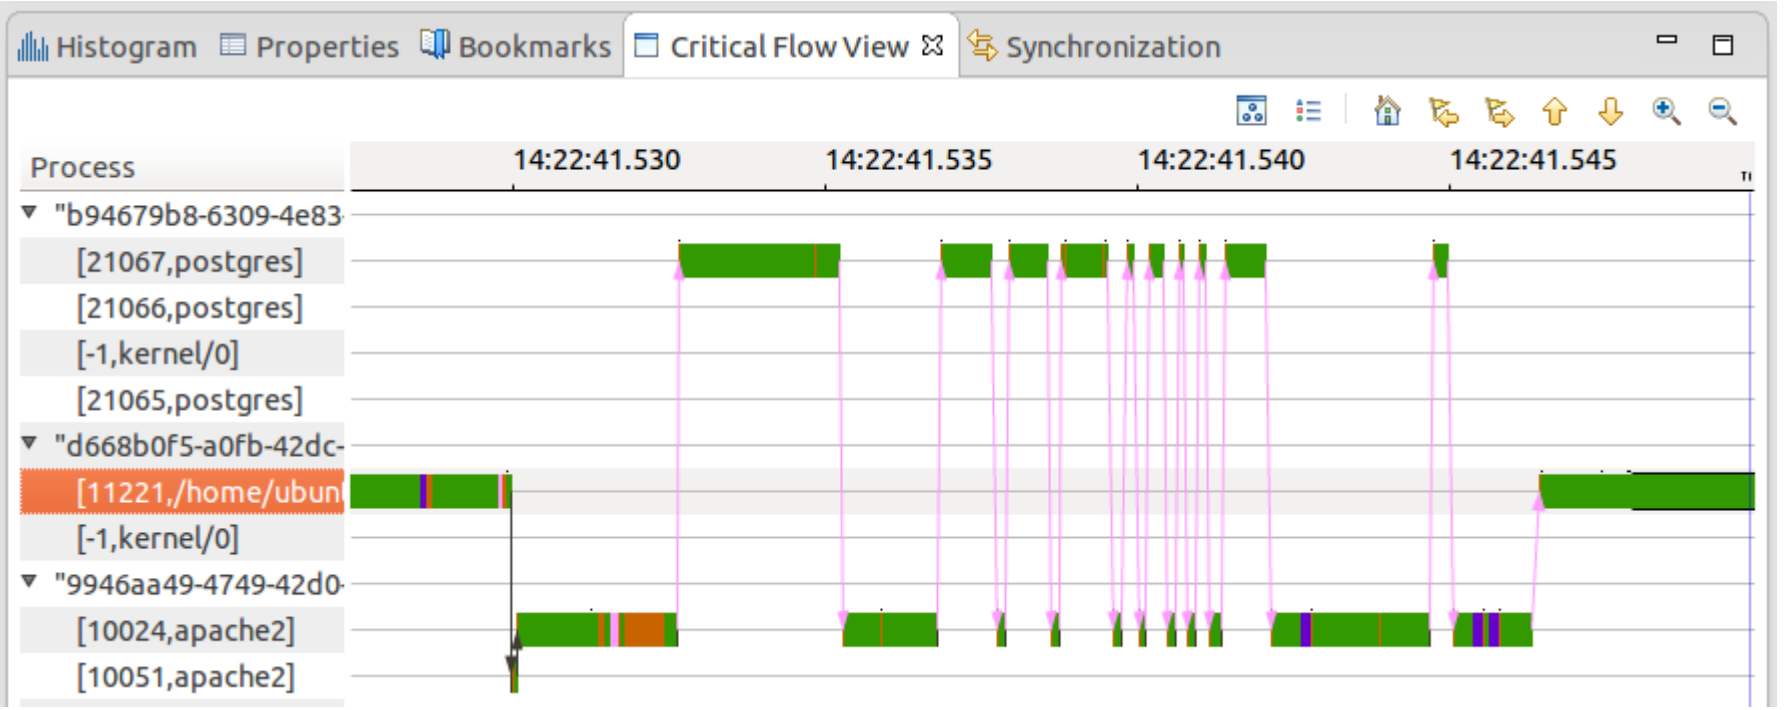

Network share

mount.cifs ~330ms

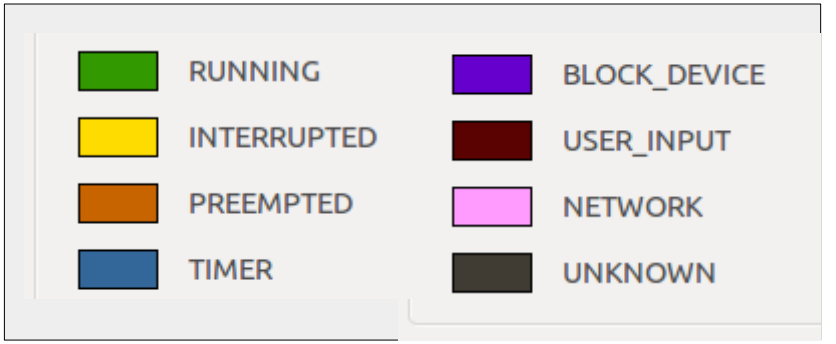

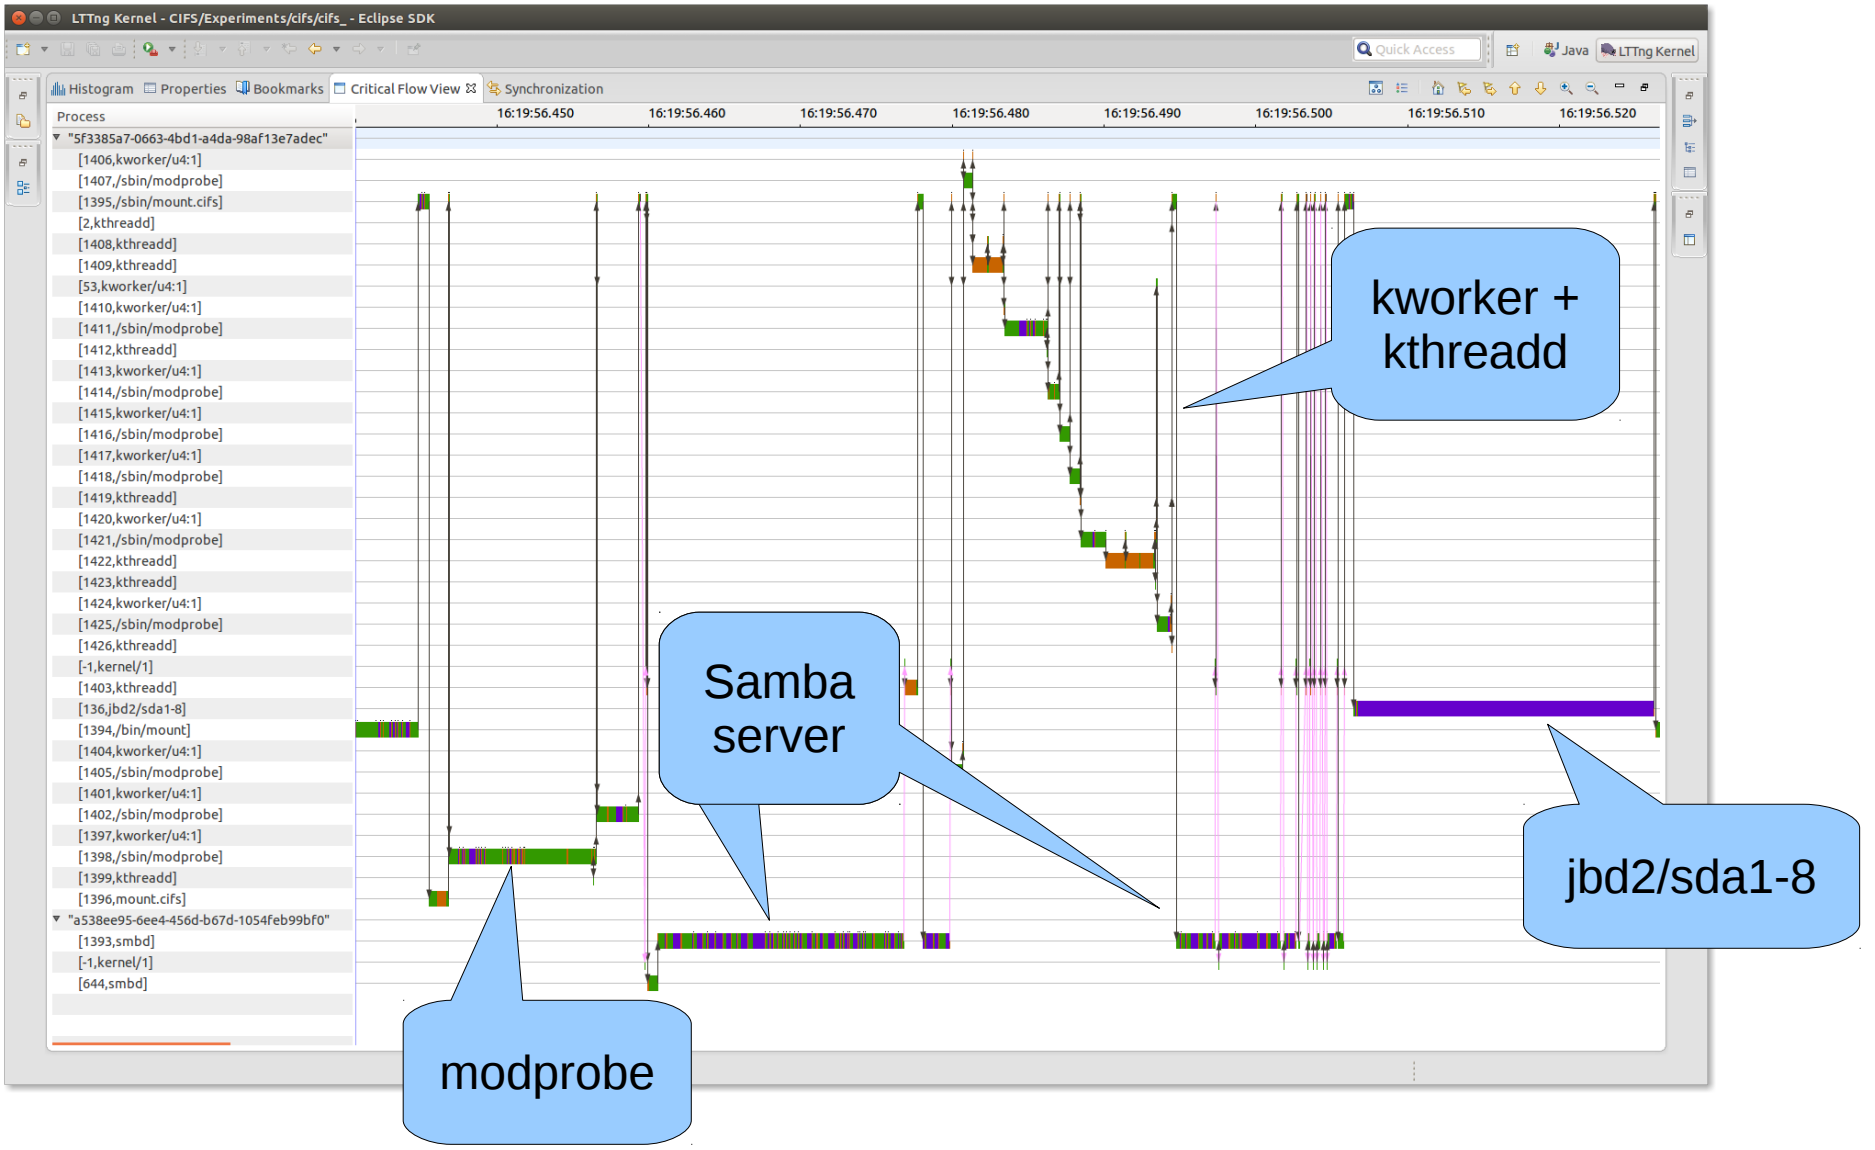

Network share

directory listing