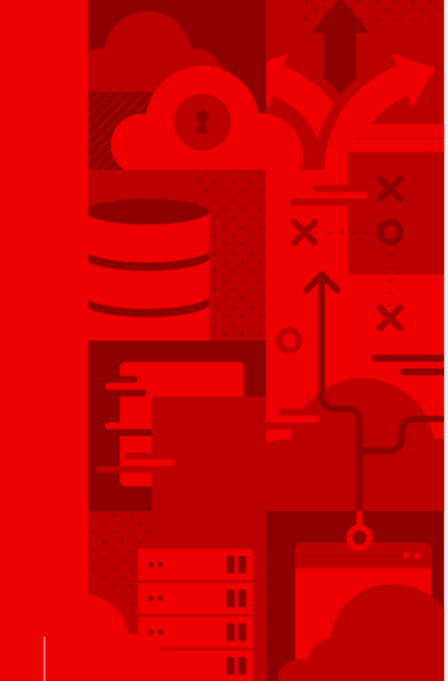

Five silly things to do

when benchmarking your BPF program

Dmitrii Dolgov

01-02-2025

Red Hat

BPF ringbuf benchmark

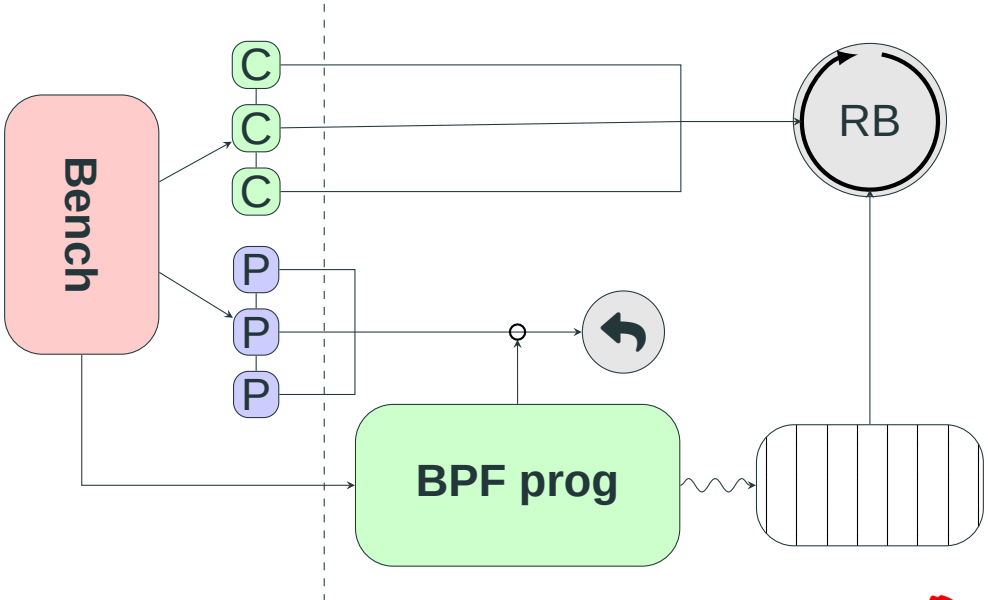

Red Hat

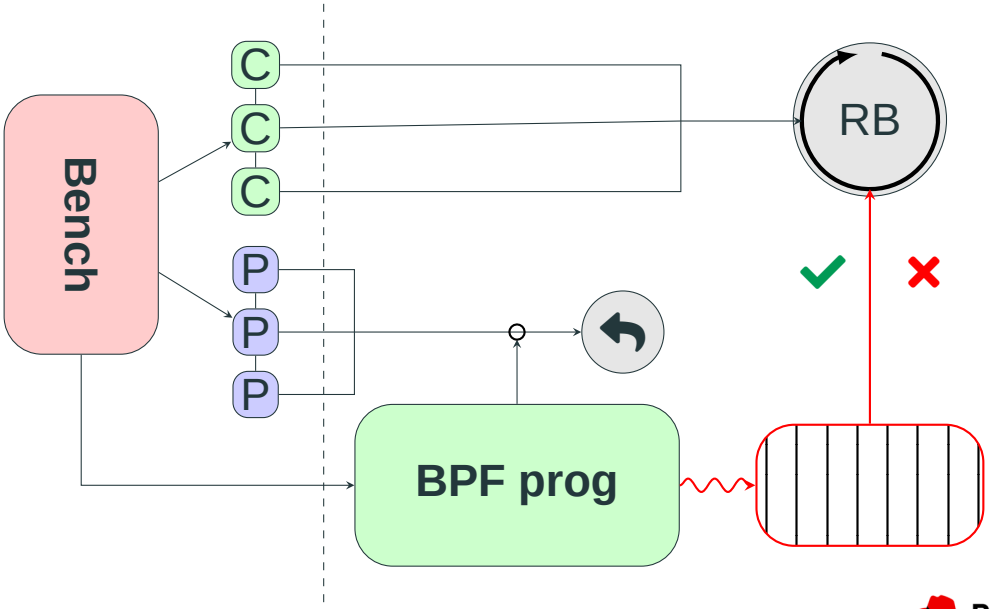
Bench
C
C
C
P
P
P
RB
BPF prog

- ➔ Concurrency
- ➔ Affinity
- ➔ Batch size

[1] Schroeder, B., Wierman, A. and Harchol-Balter, M., 2006. Open versus closed: A cautionary tale. USENIX.

Red Hat

- Concurrency
- Affinity
- Batch size
- Events distribution? [1]

[1] Schroeder, B., Wierman, A. and Harchol-Balter, M., 2006. Open versus closed: A cautionary tale. USENIX.

Red Hat

- ➔ No hyperthreading
- ➔ Scaling governor ”performance”
- ➔ No intel_pt turbo
- ➔ Poisson distribution (berserker)

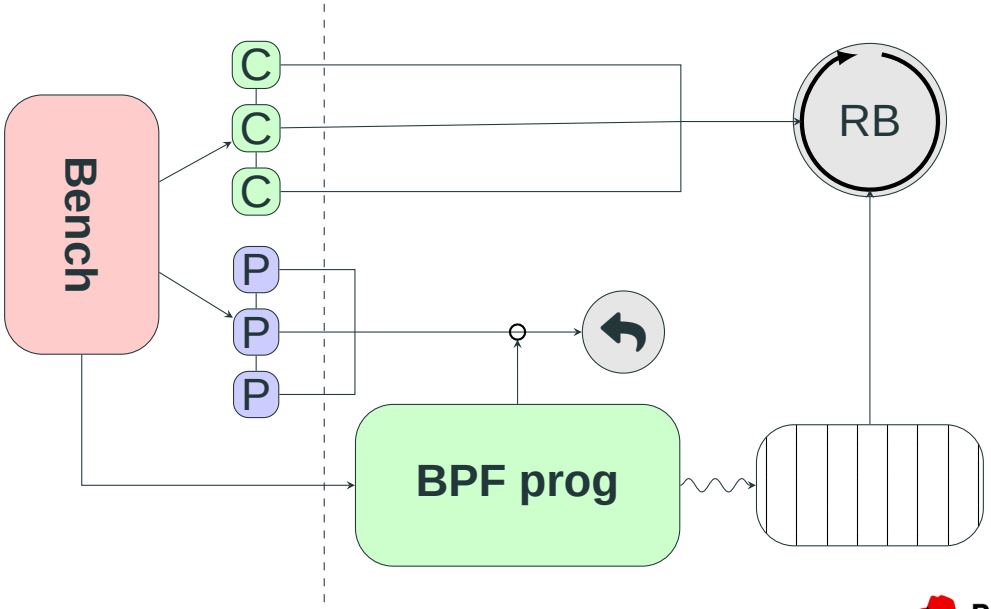
Bench
C
C
C
P
P
P
RB
BPF prog

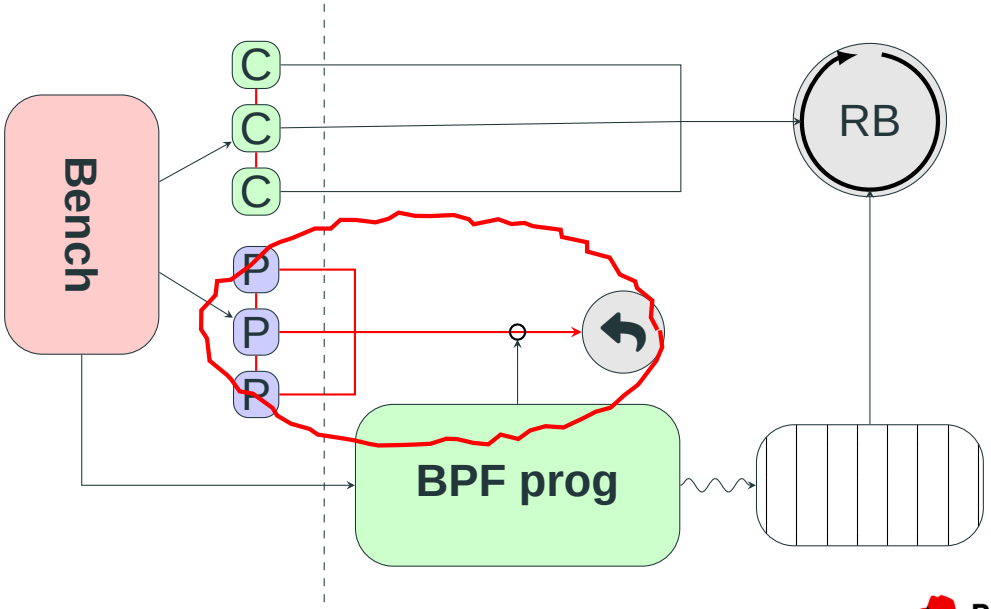

Red Hat

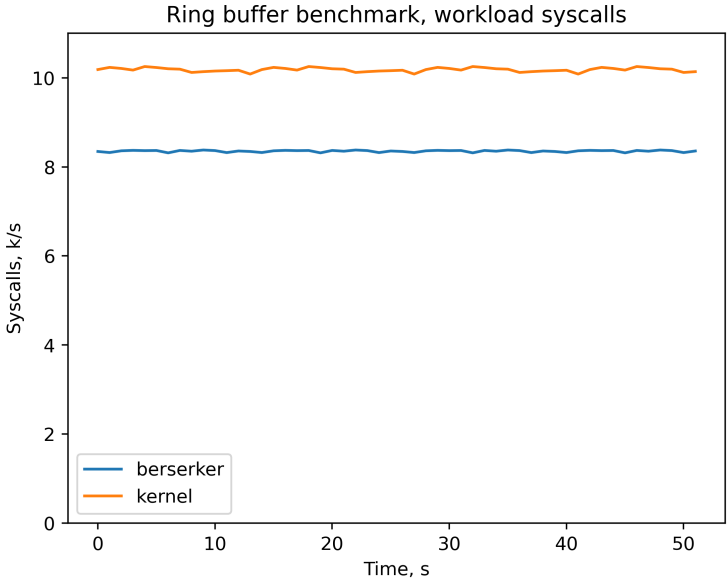
Ring buffer benchmark, workload syscalls
Syscalls, k/s
Time, s
berserker
kernel
0
2
4
6
8
10
0
10
20
30
40
50

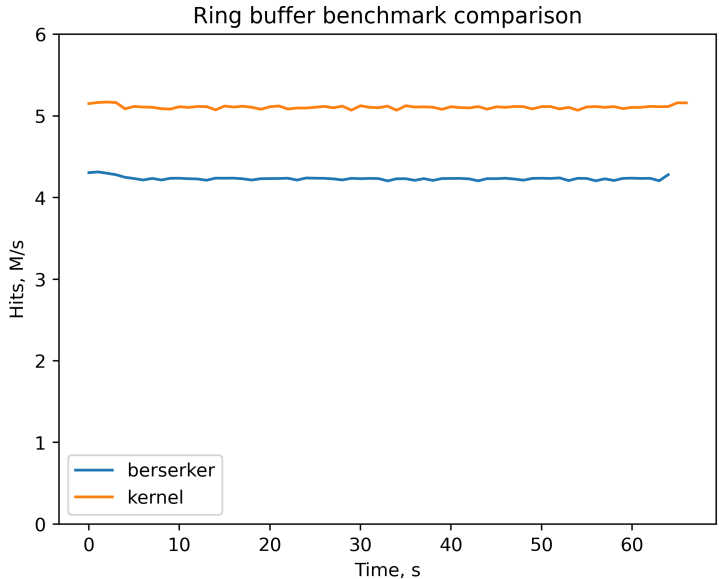
Ring buffer benchmark comparison
Hits, M/s
Time, s
berserker
kernel
0
1
2
3
4
5
6
0
10
20
30
40
50
60

```
# patched version to collect historgramms
bpftool prog profile cycles
```

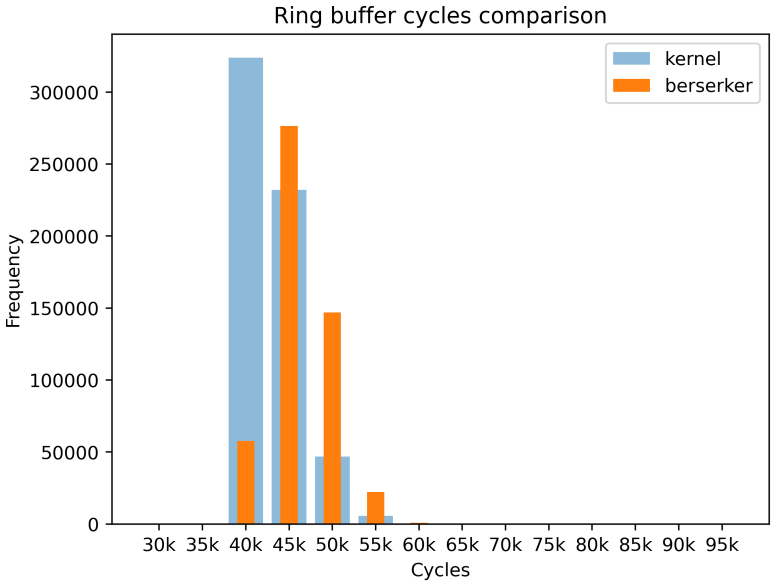
Ring buffer cycles comparison
kernel
berserker
300000
250000
200000
150000
100000
50000
0
Frequency
30k 35k 40k 45k 50k 55k 60k 65k 70k 75k 80k 85k 90k 95k
Cycles

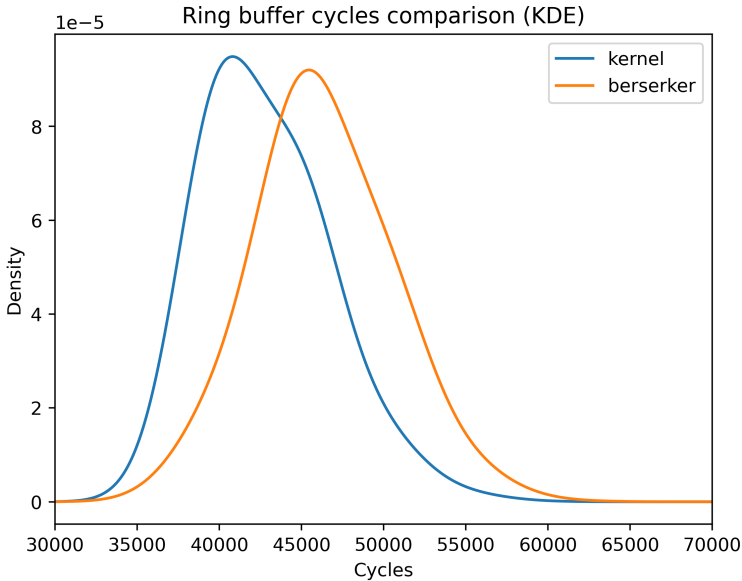
Ring buffer cycles comparison (KDE)
1e−5
kernel
berserker
Density
0
2
4
6
8
30000
35000
40000
45000
50000
55000
60000
65000
70000
Cycles

- No hyperthreading
- Scaling governor ”performance”
- ~~No intel_pt turbo~~
- Poisson distribution

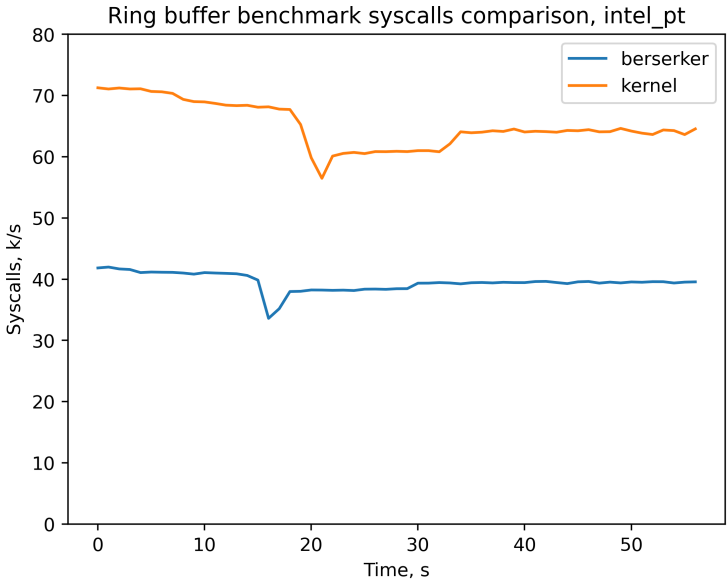
Ring buffer benchmark syscalls comparison, intel_pt
berserker
kernel
Syscalls, k/s
Time, s
0
10
20
30
40
50
60
70
80

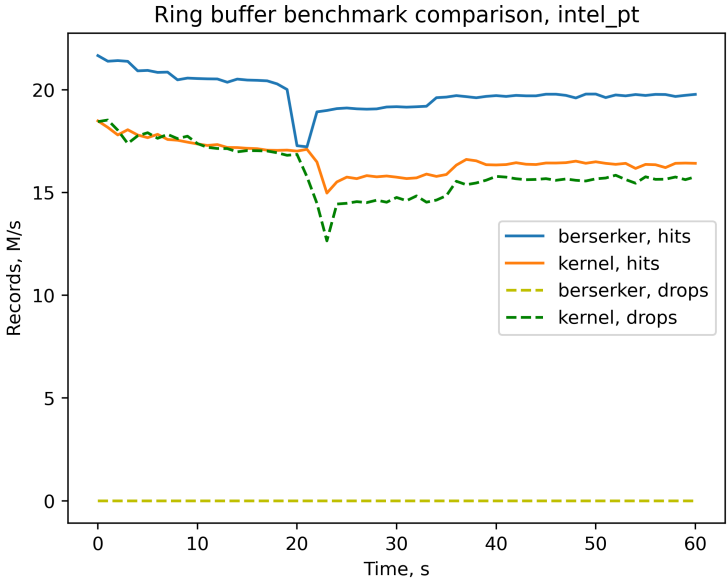
Ring buffer benchmark comparison, intel_pt
Records, M/s
Time, s
0
5
10
15
20
0
10
20
30
40
50
60
berserker, hits
kernel, hits
berserker, drops
kernel, drops

```
perf stat -M tma_backend_bound_group -b prog_id
# kernel: -14.2 %  tma_core_bound
# berserker: -34.2 %  tma_core_bound

perf stat
    -e CYCLE_ACTIVITY.STALLS_MEM_ANY
    -b prog_id
# kernel: 52,985,696,179
# berserker: 17,918,457,903
```

github.com/stackrox/berserker

- ➔ Distributions
- ➔ Userspace networking
- ➔ BPF programs contention
- ➔ ...

Why not stress-ng?

Questions?

@erthalion@fosstodon.org

ddolgov at redhat dot com

Red Hat