ISC's Performance Lab

DNS-OARC Workshop 26 - Madrid

Goals of the Project

- Provide a framework for automated performance testing of BIND
  - Test experimental code
  - Look at long-term trends to avoid regressions

System Features #1

- Multiple configurations
  - `git` branch or tag
  - `./configure` settings
  - `named.conf` options
  - `named` command line options
  - etc

System Features #2

- Multiple authoritative zone configurations
  - root zone
  - 1M delegations (small TLD)
  - 1M small zones (web hoster)
  - 1 zone with 1M A+AAAA
- `dnsperf` query sets for the above
- (Some) Recursive support

System Features #3

- HTML5 Web UI
  - real time updates over WebSockets
- Round-robin scheduler
  - with ad-hoc priority bump

System Features #4

- Graphs
- CSV Output
- A/B comparison with Student's t-test
- Extensible for other servers
  - NSD 4, Knot 2 already supported
  - Dummy 'echo' server support, too

Technology

- MongoDB
- NodeJS (ES6)
  - Promises
  - Express
- HTML5
  - Bootstrap
  - Angular JS (ES5)
  - WebSockets

QUICK DEMO…

Logical Architecture

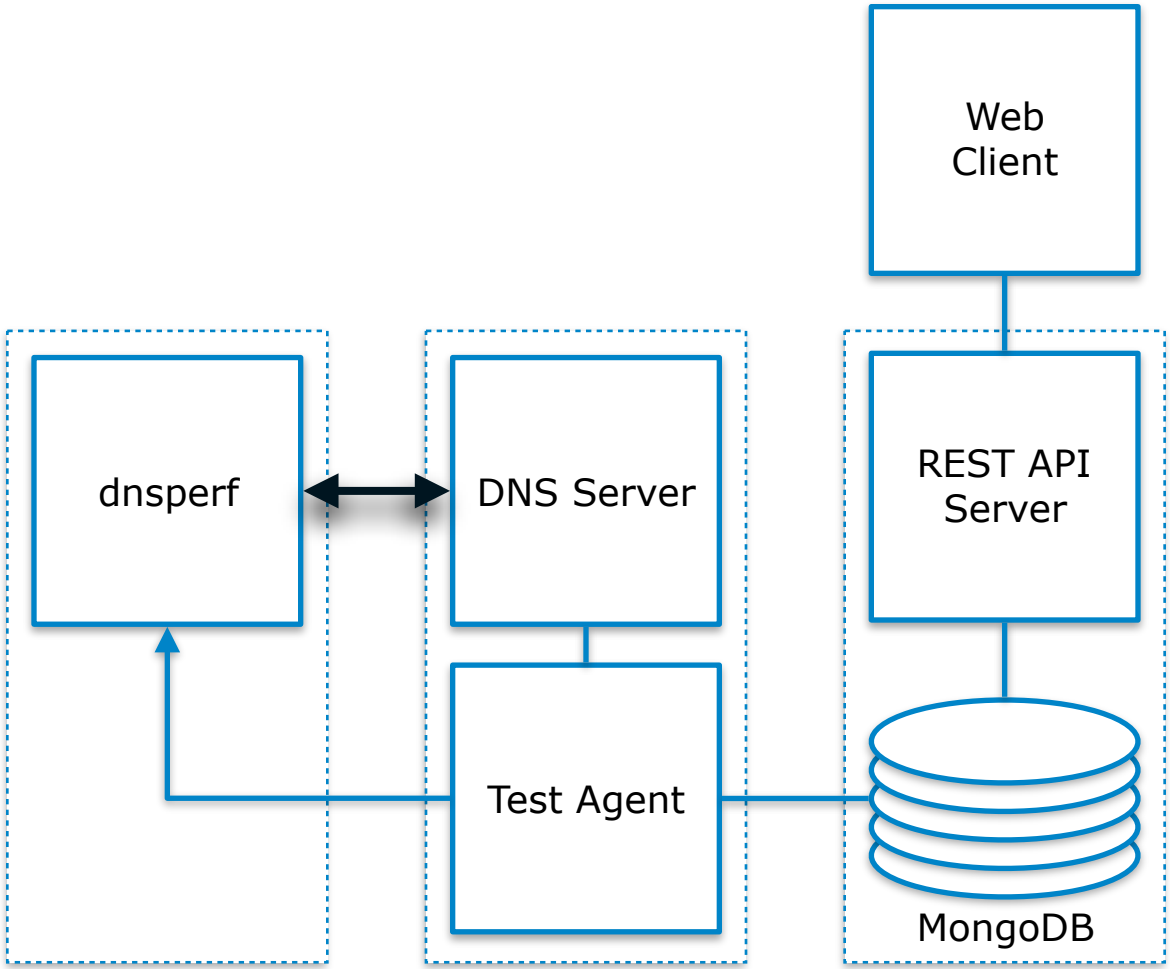

Physical Architecture

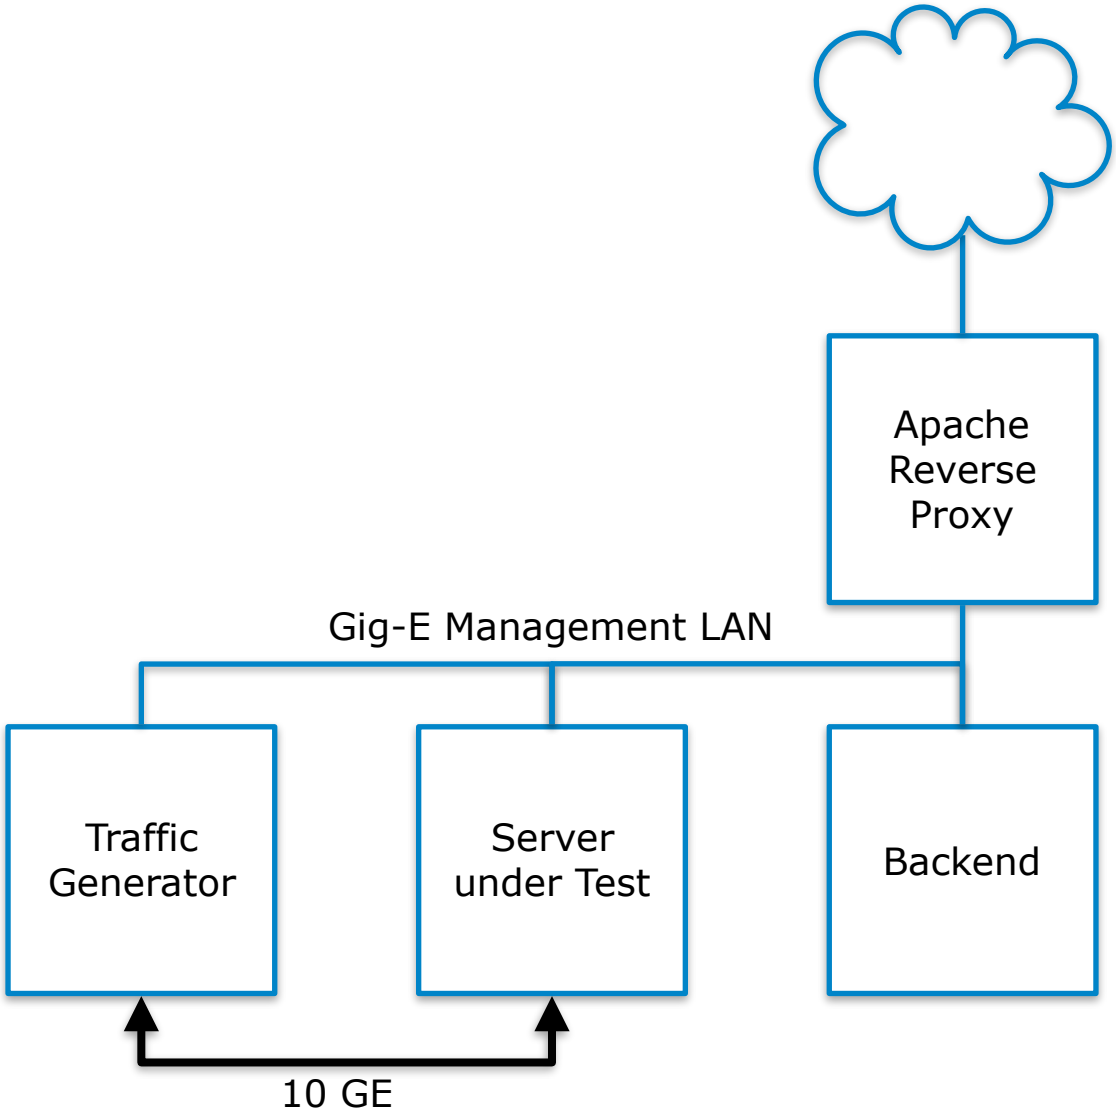

Test Methodology

- Build and start server
- Run `dnsperf` 30 times
  - 30 seconds each time
  - first run ignored for statistic (allows for cache and buffer warmup)

Challenge - Test Variability

- The results from 30x dnsperf tests against a single run of BIND are generally consistent
- Testing a new run of BIND often shifts the mean significantly, by several standard deviations

Challenge - Test Variability

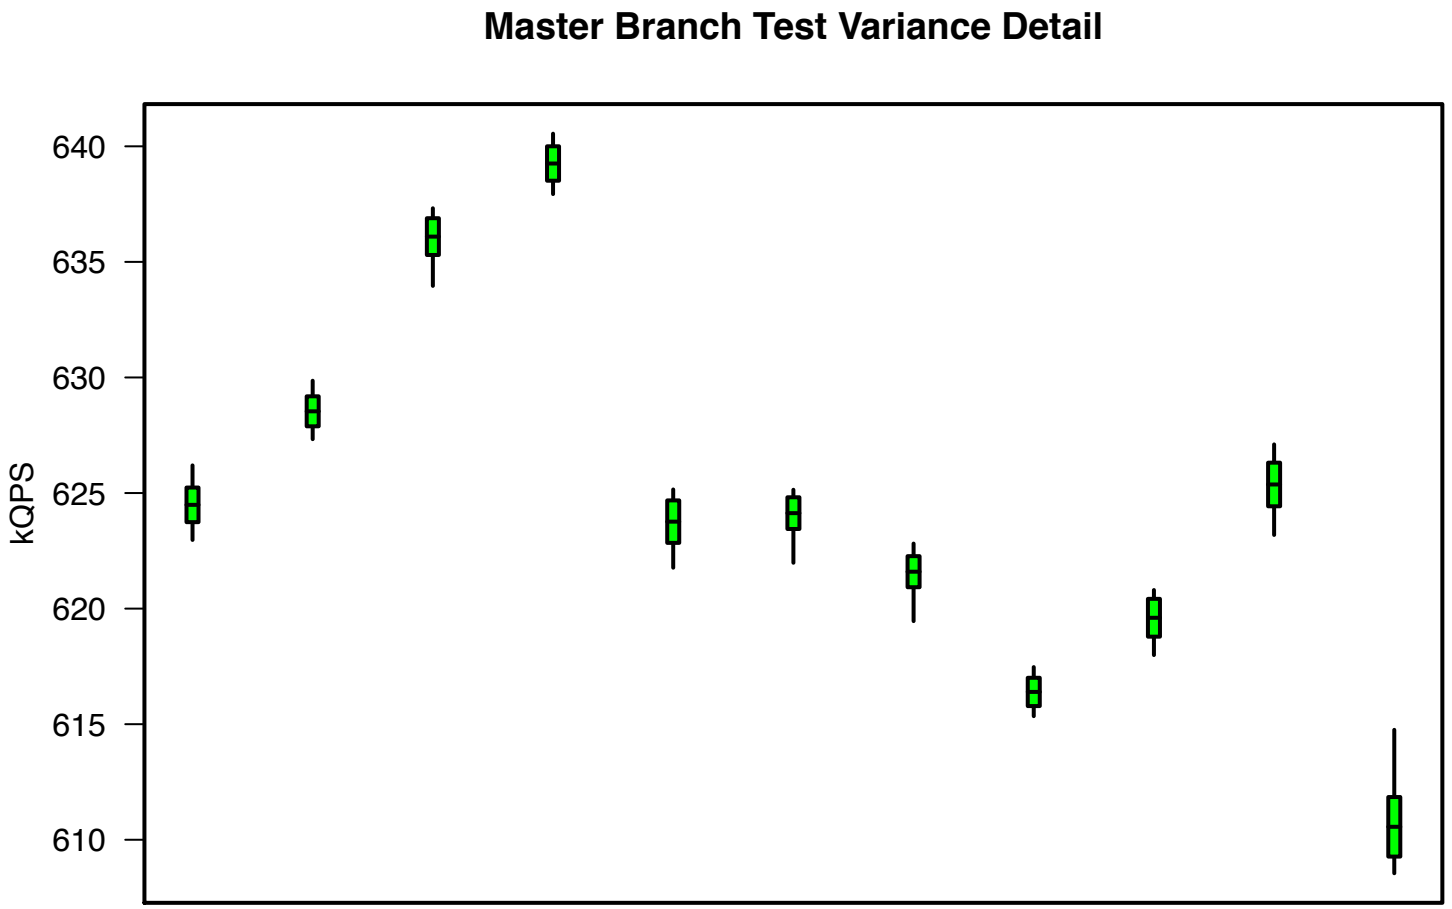

(Some) Mitigation

- Tune for stability, not peak performance:
  - Disable Hyperthreading
  - Lock CPU clock rate
    - Disable Intel SpeedStep
    - Disable Turbo mode
  - Lock Client and Server CPU core affinity
  - Lock NIC RX/TX queue CPU core affinity
  - Set NIC queue flow hashing to use deterministic ports

Challenge - Test Variability #2

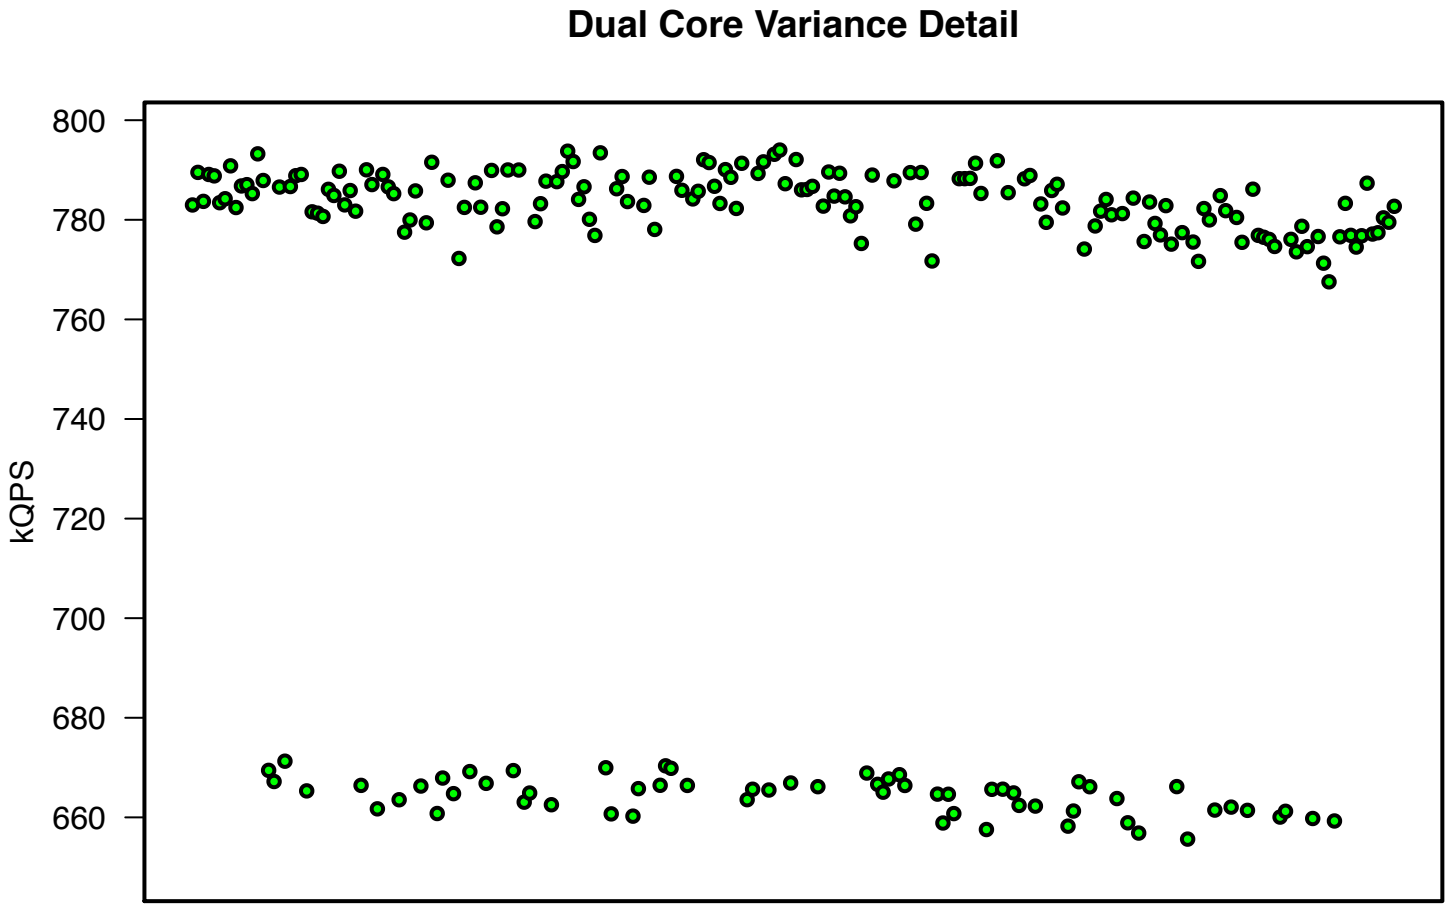

echo server two threads locked to CPU#0 and #1

Challenge - Test Variability #2

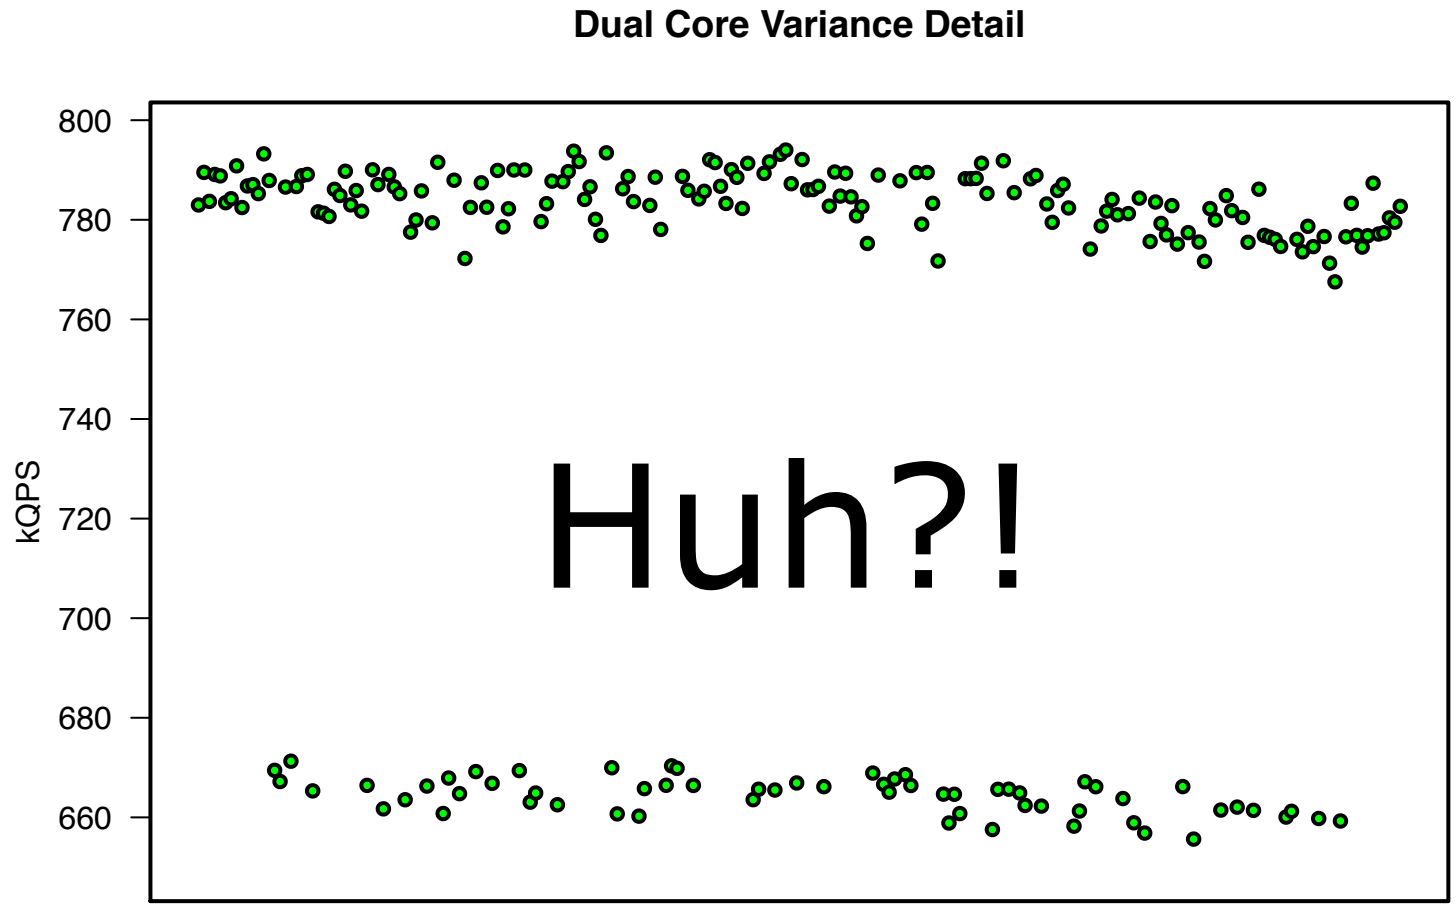

echo server two threads locked to CPU#0 and #1

Coming Soon

- In 9.11.2 and back-ported to 9.10
  - optimised name compression
  - optimised owner case preservation
- In 9.12
  - "minimal-responses" on by default
  - "glue cache" replaces "acache"
    - improves delegation performance

BIND Performance – 1M Delegations

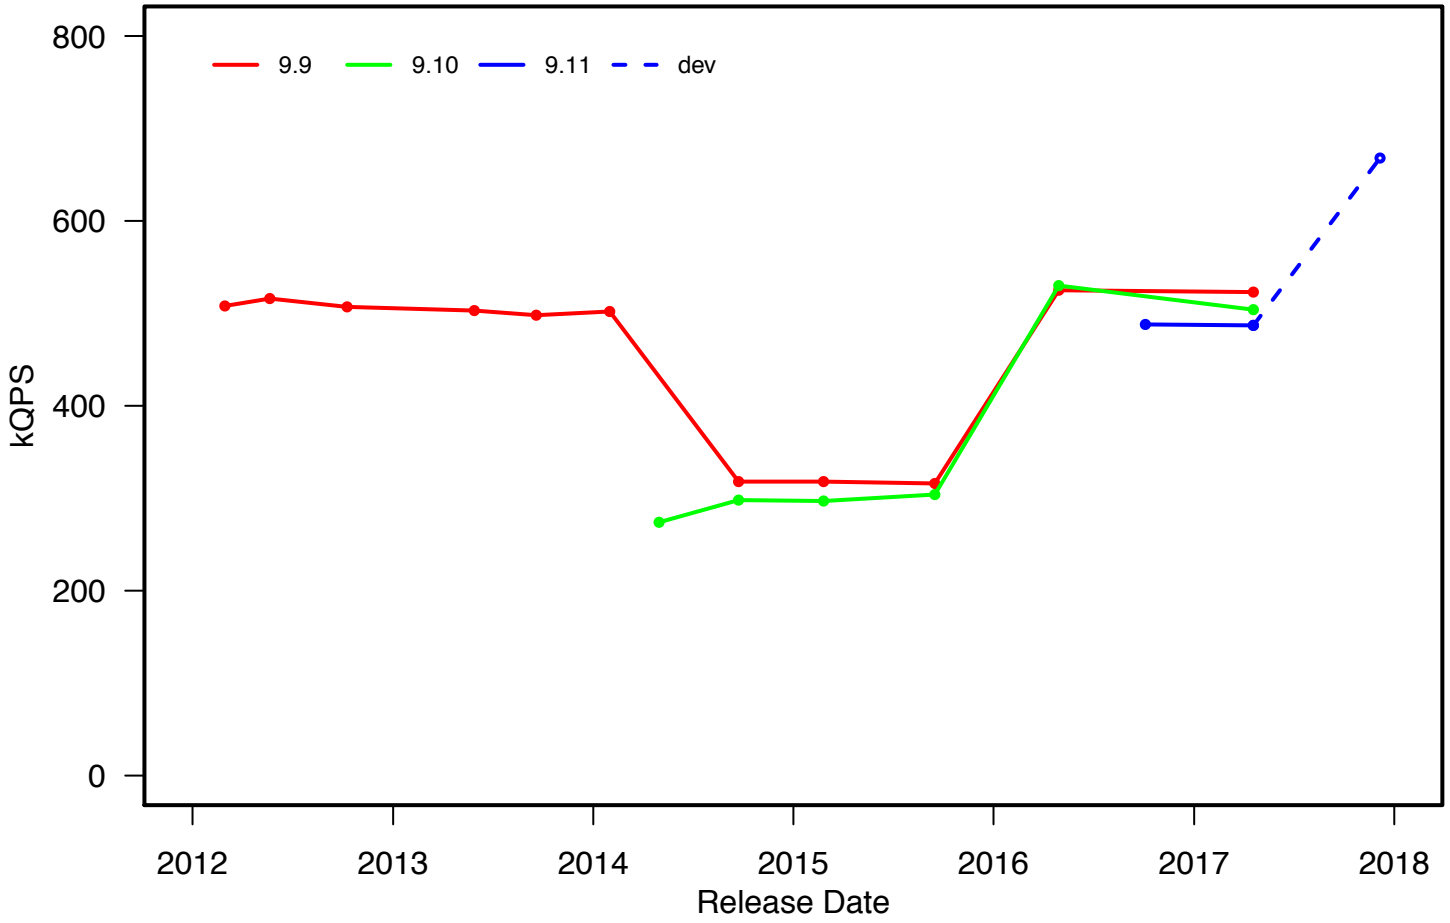

12–core Intel Xeon E5–2680 v3 @ 2.50GHz

BIND Performance – 1M Delegations

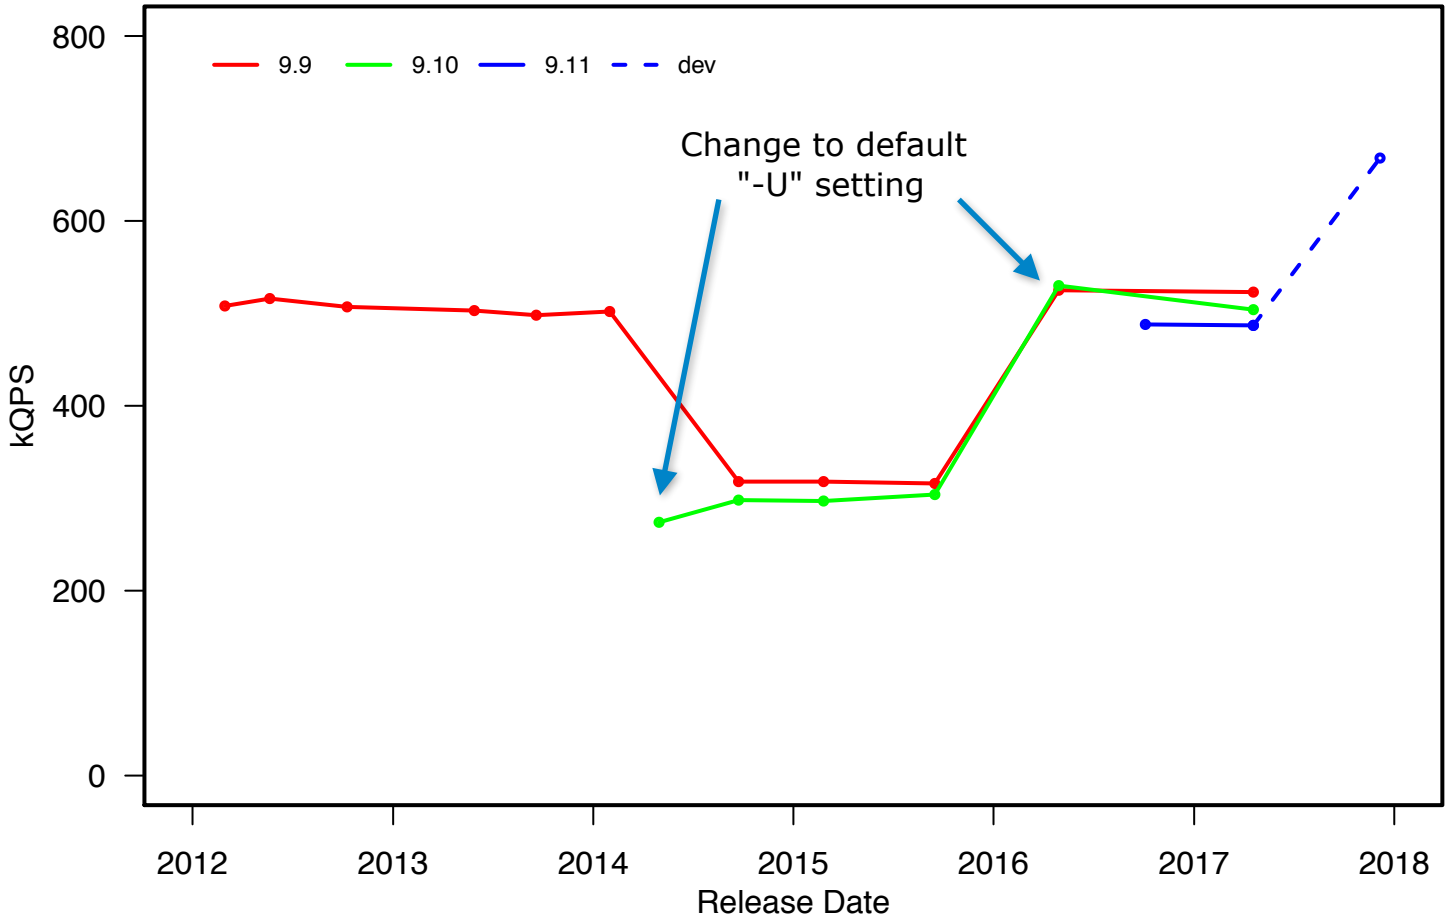

12–core Intel Xeon E5–2680 v3 @ 2.50GHz

BIND Performance – Root Zone

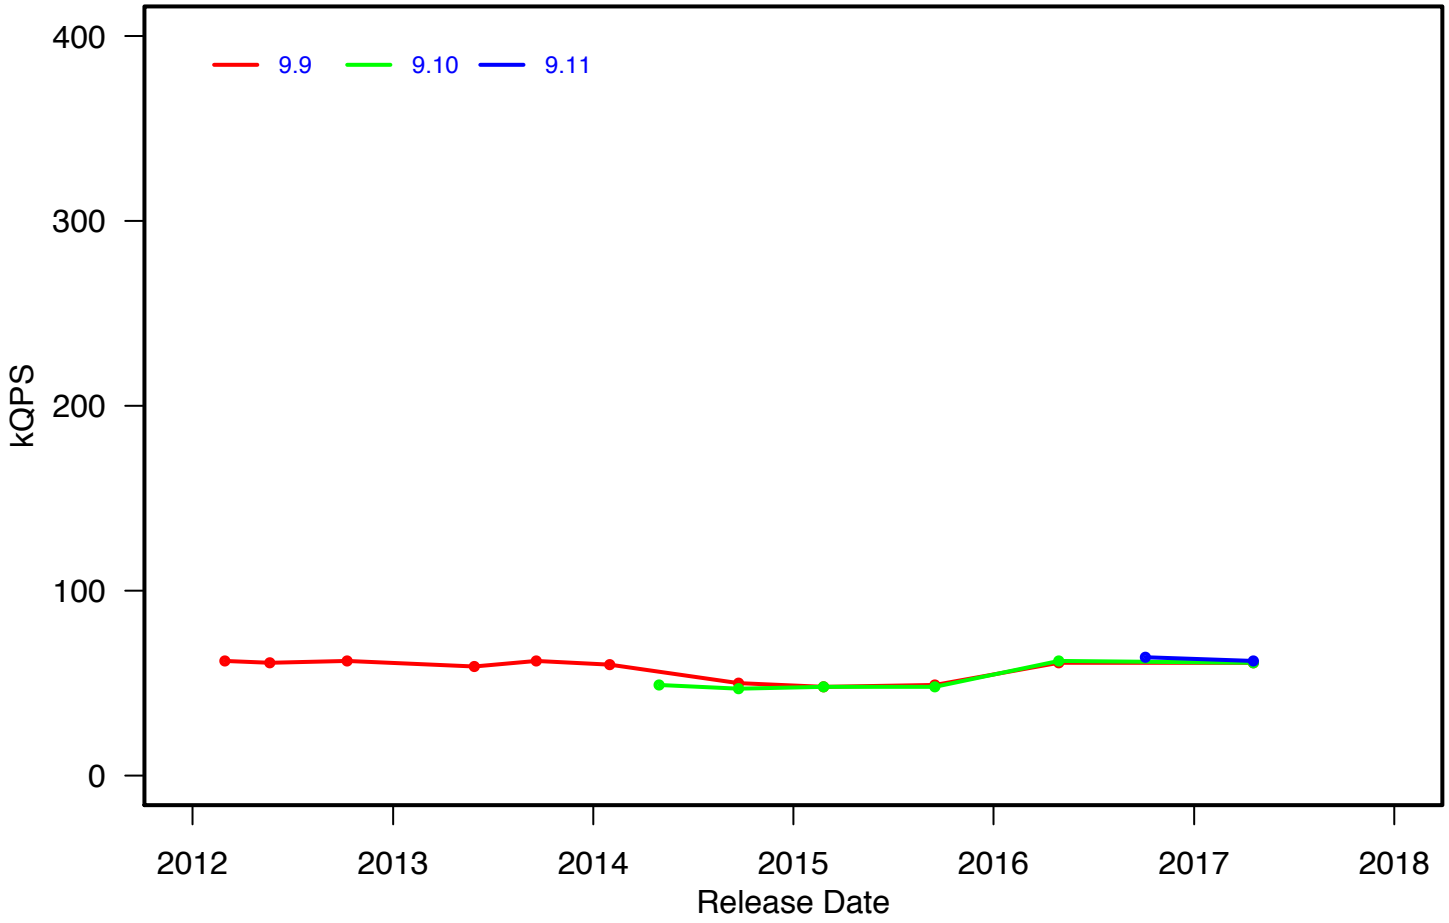

12–core Intel Xeon E5–2680 v3 @ 2.50GHz

BIND Performance – Root Zone

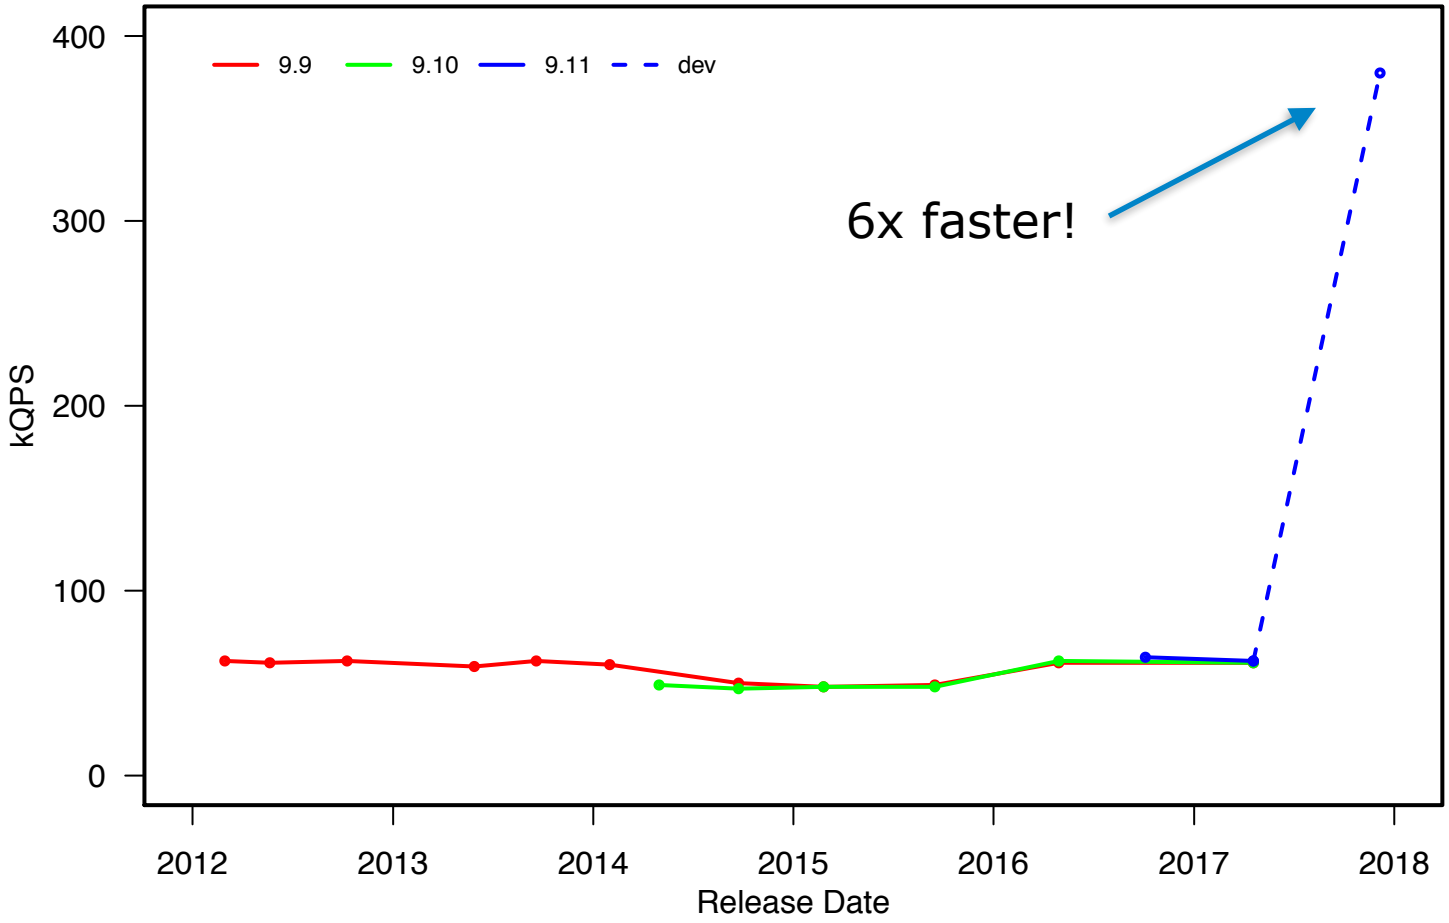

12–core Intel Xeon E5–2680 v3 @ 2.50GHz

How Do I Get It?

https://github.com/isc-projects/perflab

- This is **not** supported software
- It's not turn-key - installation is likely to require lots of fiddling
- Pull requests welcomed…

QUESTIONS?