ymmv:
Yeti Many Master Verifier

Shane Kerr / BII Labs / shane@biigroup.cn

2016-11-12 / Seoul · 서울 / Yeti Workshop

Yeti Needs Traffic

- Yeti is a production testbed
- In order to fulfill this goal, Yeti needs lots of real traffic
- DNS caching makes it very difficult to get large amounts of traffic
  - Even a very, very busy cluster of servers produces a trickle of queries
  - Google servers each send <50 queries/second to IANA root servers[1]

1 http://recs.conf.meetecho.com/Playout/wa...pter_1 (starting around 20:20)

天地互连

Sources of Traffic

- Real resolvers configured for Yeti
  - The ”gold standard” for traffic
  - May harm users, since Yeti is experimental!
- Synthetic traffic
  - Easy to generate
  - Can stress network and servers
- Mirrored production traffic
  - Real traffic, good for Yeti load & patterns
  - Problems in Yeti not detected

Bii
天地互连

ymmv

- Another way to generate Yeti traffic
- Mirrored production traffic++
  - IANA answers compared with Yeti answers

ymmv: Design Goals

- Easy
  - Simple for administrators to install & run
  - Flexible, works with existing environments
- Informative
  - For administrators
  - For Yeti operators
- Safe
  - Does not affect production service
  - Privacy-maintaining

ymmv: Design Choices

- Go language for implementation
  - Relatively fast, good concurrency support
  - Great DNS library
  - Compiles to static binaries
  - Dependencies are painless
- Input handled via separate component
  - Allows migrating to alternate format later
- Simple command-line application

ymmv: Components

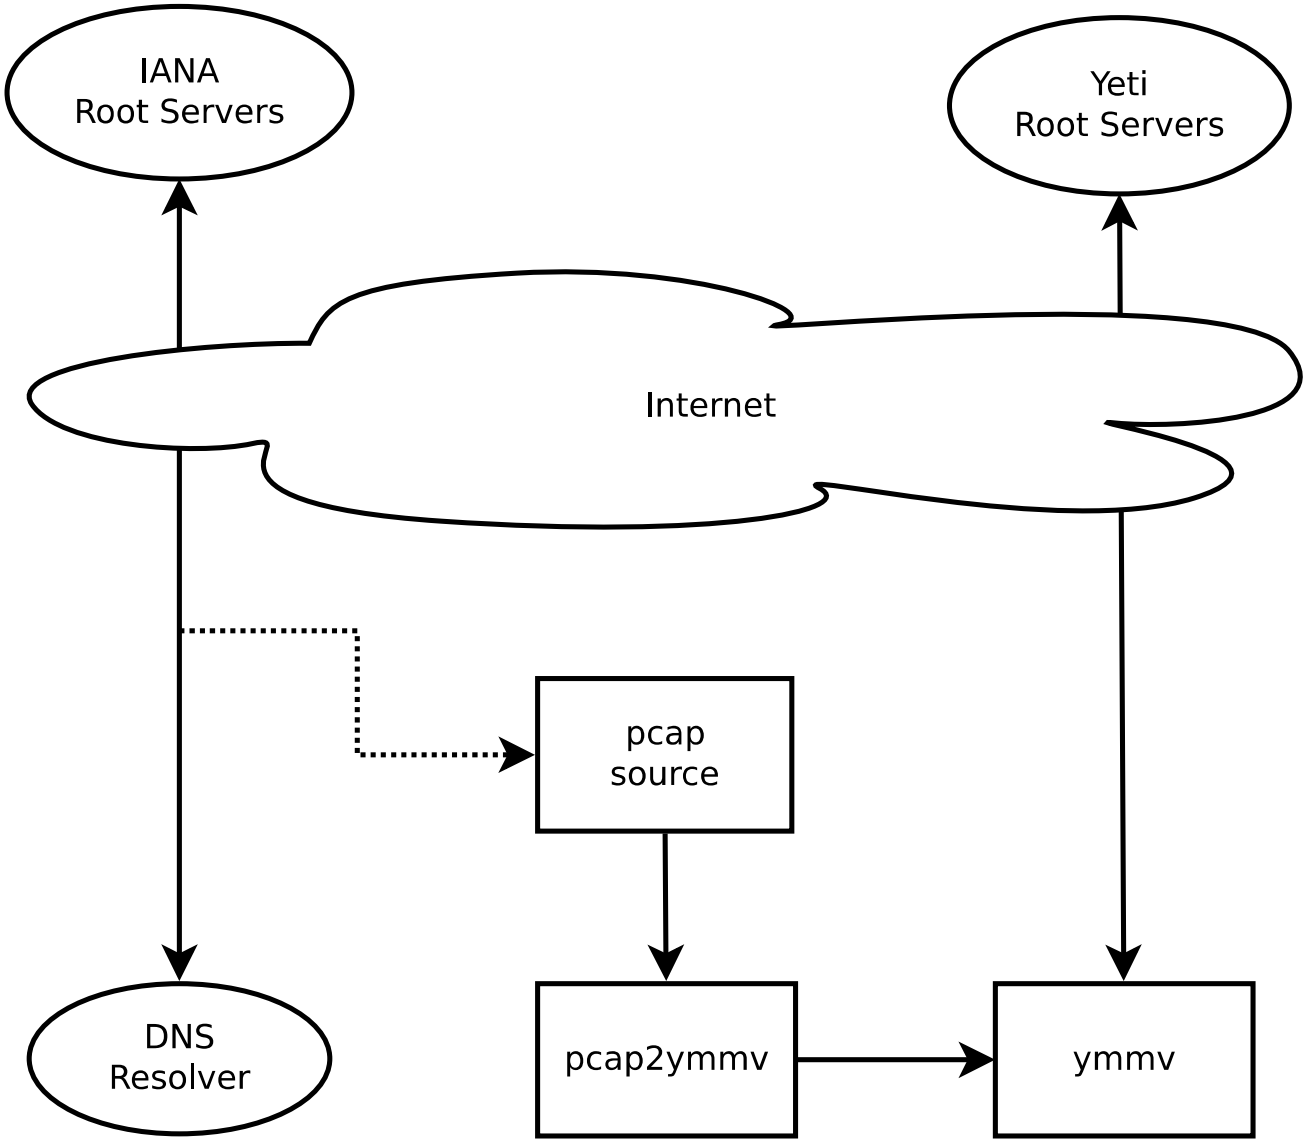

ymmv: Components

- pcap as input
  - dnstap or CBOR might be better eventually
  - tcpdump, tshark, dnscap, …
  - Can be live or replayed (live probably better)
- Pcap2ymmv
  - Convert pcap to custom format
  - Looks for traffic to IANA root servers
  - Matches queries & answers
- ymmv
  - Does actual queries & comparisons

ymmv: Features

- Detect IANA vs. Yeti differences
- Compare IANA vs. Yeti performance
- Server selection: RTT, round-robin, random, all
- Default to unusual EDNS buffer size (4093)
- Obfuscate query names (by default)
  - `example` → `example`
  - `foo.example` → `ymmv.845a838696ae1e5a.example`
- Send daily reports (opt-in) via SMTP or Sendmail

Bii 天地互连

ymmv: Usage [1 of 3]

```
$ nohup capture.sh eth0 &>/dev/null &
$ nohup capture.sh he-ipv6 &>/dev/null &
$
```

ymmv: Usage [2 of 3]

```
$ ps axo args | grep tcpdump | grep -v grep
tcpdump -i eth1 -w- -U -q udp port 53
tcpdump -i he-ipv6 -w- -U -q udp port 53
$ ps axo args | grep ymmv | grep -v grep
./../pcap2ymmv/pcap2ymmv
./../ymmv/ymmv -v 1 -p /tmp/ymmv-eth1-perf -d /tmp/ymmv-
eth1-diff -r -sendmail
./../pcap2ymmv/pcap2ymmv
./../ymmv/ymmv -v 1 -p /tmp/ymmv-he-ipv6-perf -d
/tmp/ymmv-he-ipv6-diff -r -sendmail -mail-
to=shane@blij.tk
```

ymmv: Usage [3 of 3]

Option	Description
-v 1	Increase verbosity of logging. By default, logs are placed in uniquely-named files in **`/tmp`** by the Go logging library.
-p *file_name_base*	File name to log performance to, as a CSV file. Gets date appended.
-d *file_name_base*	File name to log differences to. Gets date appended.
-r	Send daily reports of performance & differences via e-mail
-sendmail	Use local sendmail for e-mail (the other option is SMTP)
-mail-to *email_address*	The e-mail address to send to

Some Differences Discovered

- TTL problem with TISF server [fixed]
- Query with ’%’ broke PowerDNS [fixed]
- Query for DS records returns additional section from Bundy server [open]

IANA RTT (IPv4) vs. Yeti RTT (IPv6)

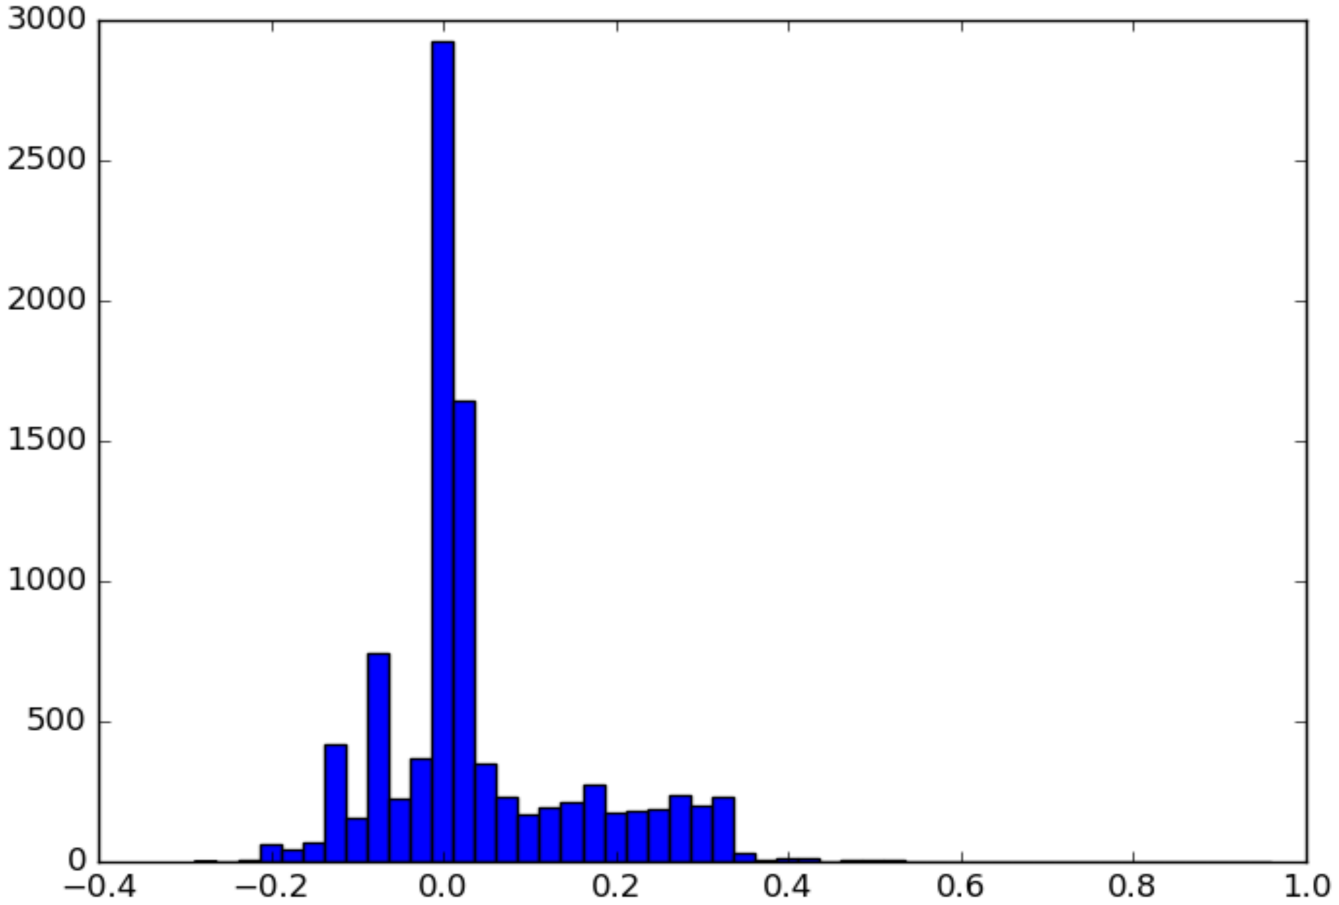

Yeti 40 millisecond slower on average,
a 105% difference.

天地互连

IANA RTT (IPv6) vs. Yeti RTT (IPv6)

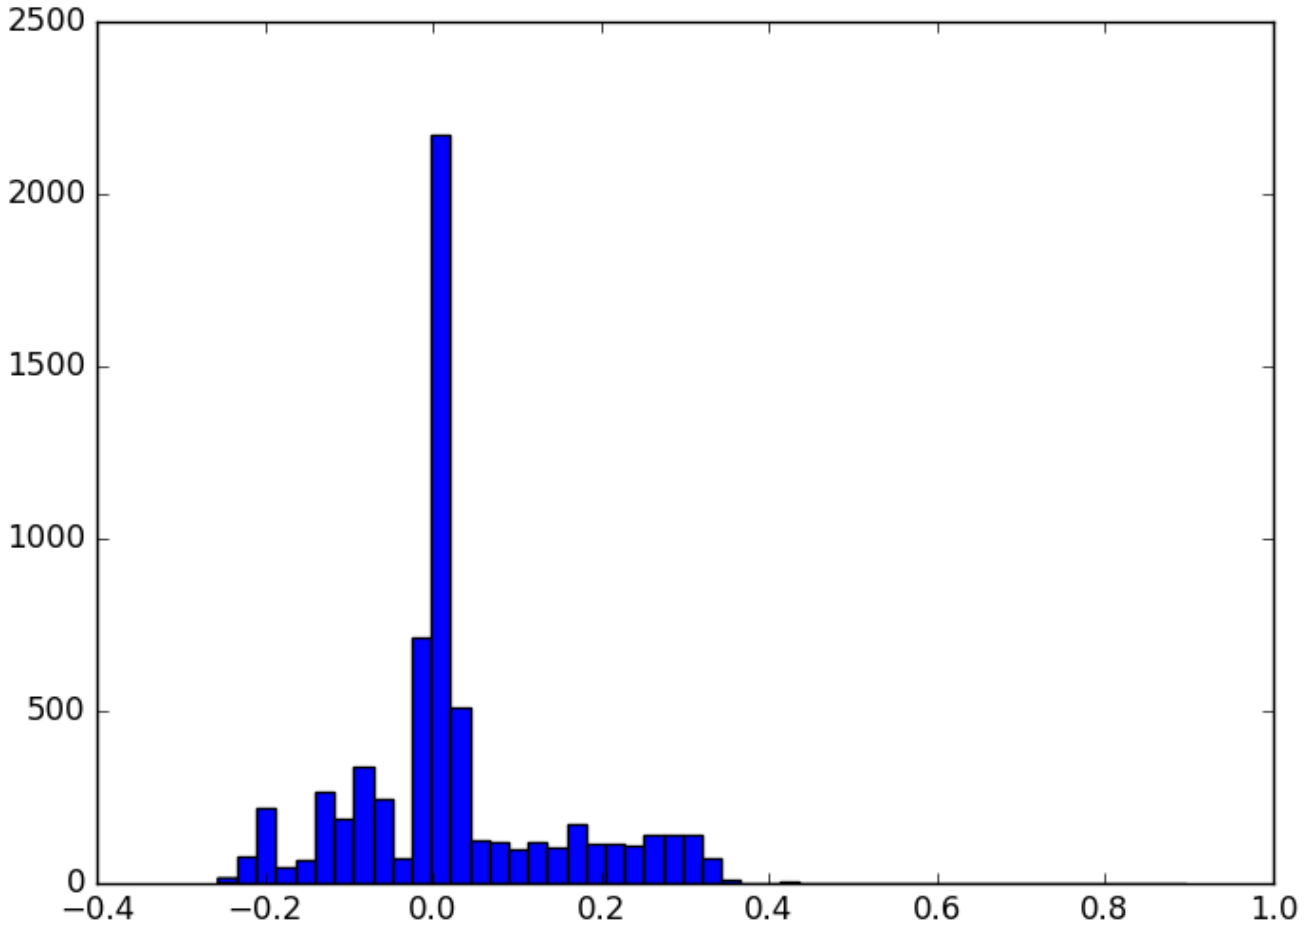

Yeti 24 millisecond slower on average,
a 50% difference.

天地互连

RTT Comparison Analysis

- Yeti than IANA is slower in Amsterdam
  - IPv6 vs. IPv6 comparison most interesting
- Not suprising… there are at least 4 IANA root servers at AMS-IX
- Server selection algorithm has huge impact
- Ability for any operator to perform similar RTT comparison may be interesting
- Getting such comparison from more locations may be interesting.

天地互连

Summary

- ymmv should be easy to run
- ymmv can provide us traffic
- ymmv can give us insight into Yeti
- Please use it! 😀

- http://dnsv6lab.net/2016/10/13/ymmv/

天地互连