Current status of Yeti DNS Project

Davey Song @ BII Lab
2015.10.31 Yeti Workshop,Yokohama

Bii
天地互连

Outline

- Background & Motivation
- Yeti Testbed & Statistics
  - Distribution master
  - Authority server
  - Resolver & traffic
  - Data collection& Monitoring
- Some technical findings and bugs report
- Conclusion

Related work & discussion on Root System

- **ICANN ITI Panel & technical report**
  - https://www.icann.org/en/system/files/files/iti-report-15may14-en.pdf
- **ICANN RSSAC documents**
  - RSSAC 002: Advisory on Measurements of the Root Server System
  - RSSAC003: Report on Root Zone TTLs
  - History and Technical Analysis of the Naming Scheme Used for Individual Root Servers (working on)
- **ICANN Root Zone KSK Rollover Plan(draft)**
- **Scaling the Root by Geoff Huston, IPJ, March 2015**
- **IETF work on DNS Root system**
  - draft-ietf-dnsop-root-loopback-05
  - draft-ietf-dnsop-resolver-priming-05
  - RFC7626: DNS Privacy Considerations, by S. Bortzmeyer

Root system is “special”?

- The top infrastructure ／entrance of DNS system/
- The priming process& hint file stuff is not fully documented as part of DNS protocol
- Produce Root zone／ signed the Root zone ／Distribute the root zone by various parties
- The KSK of Root zone is the Trust anchor/No parent DS
- Rely heavily on BGP routing system (Anycast) to support Root system
- Regarding Internet governance for non-technical people
    - may view the root as “the control of Internet”

What is Yeti?

- Yeti is an IPv6 only Live Root DNS Server System Testbed
  - Precisely mirrors the IANA DNS namespace
  - Experimental project with 3 years duration and clear goal
- Like IANA, has diverse servers globally
  - Server operators are volunteers from many nations
- Like IANA, has DNSSEC, with a published signing key
  - Has its own DNSSEC signing and validation keys
- System is intended for Internet-scale *science*

Why: Problem Space of Yeti(1)

Conflict between DNS Centralization Vs. Network Autonomy

- **External Dependency**
  - Local services rely on external root services
  - Require external management and support
- **Surveillance risk**
  - Information leakage cause by the DNS Root lookup
  - RFC7626: DNS Privacy Considerations, by S. Bortzmeyer

Why: Problem Space of Yeti(2)

- **Can IPv6-only DNS survive**?
  - Some DNS servers which support both A & AAAA (IPv4 & IPv6) records still do not respond to IPv6 queries
  - IPv6 introduces larger MTU (1280 bytes) , but a different fragmentation model
- **Is it ready for KSK Rollover, or not?**
  - Not all resolver is compliant to RFC5011
  - Larger packets will introduce risks during ksk/zsk rollover
- **And, Renumbering issue**

https://github.com/BII-Lab/Yeti-Project/blob/master/doc/Yeti_PS.md

Hypothesis & Experiments expected on Yeti

- IPv6-only operation
- DNSSEC Key rollover and even algorithm rollover
- Renumbering with larger frequency
- Adding more than 13 root servers (How about 25 or more?)
- Multiple zone file signers
- Multiple zone file editors (some kind of Shared zone control)

“a good design could allow a political process of deciding how control for a particular zone should be shared to start” --- ICANN ITI technical report

Architecture Design for Yeti

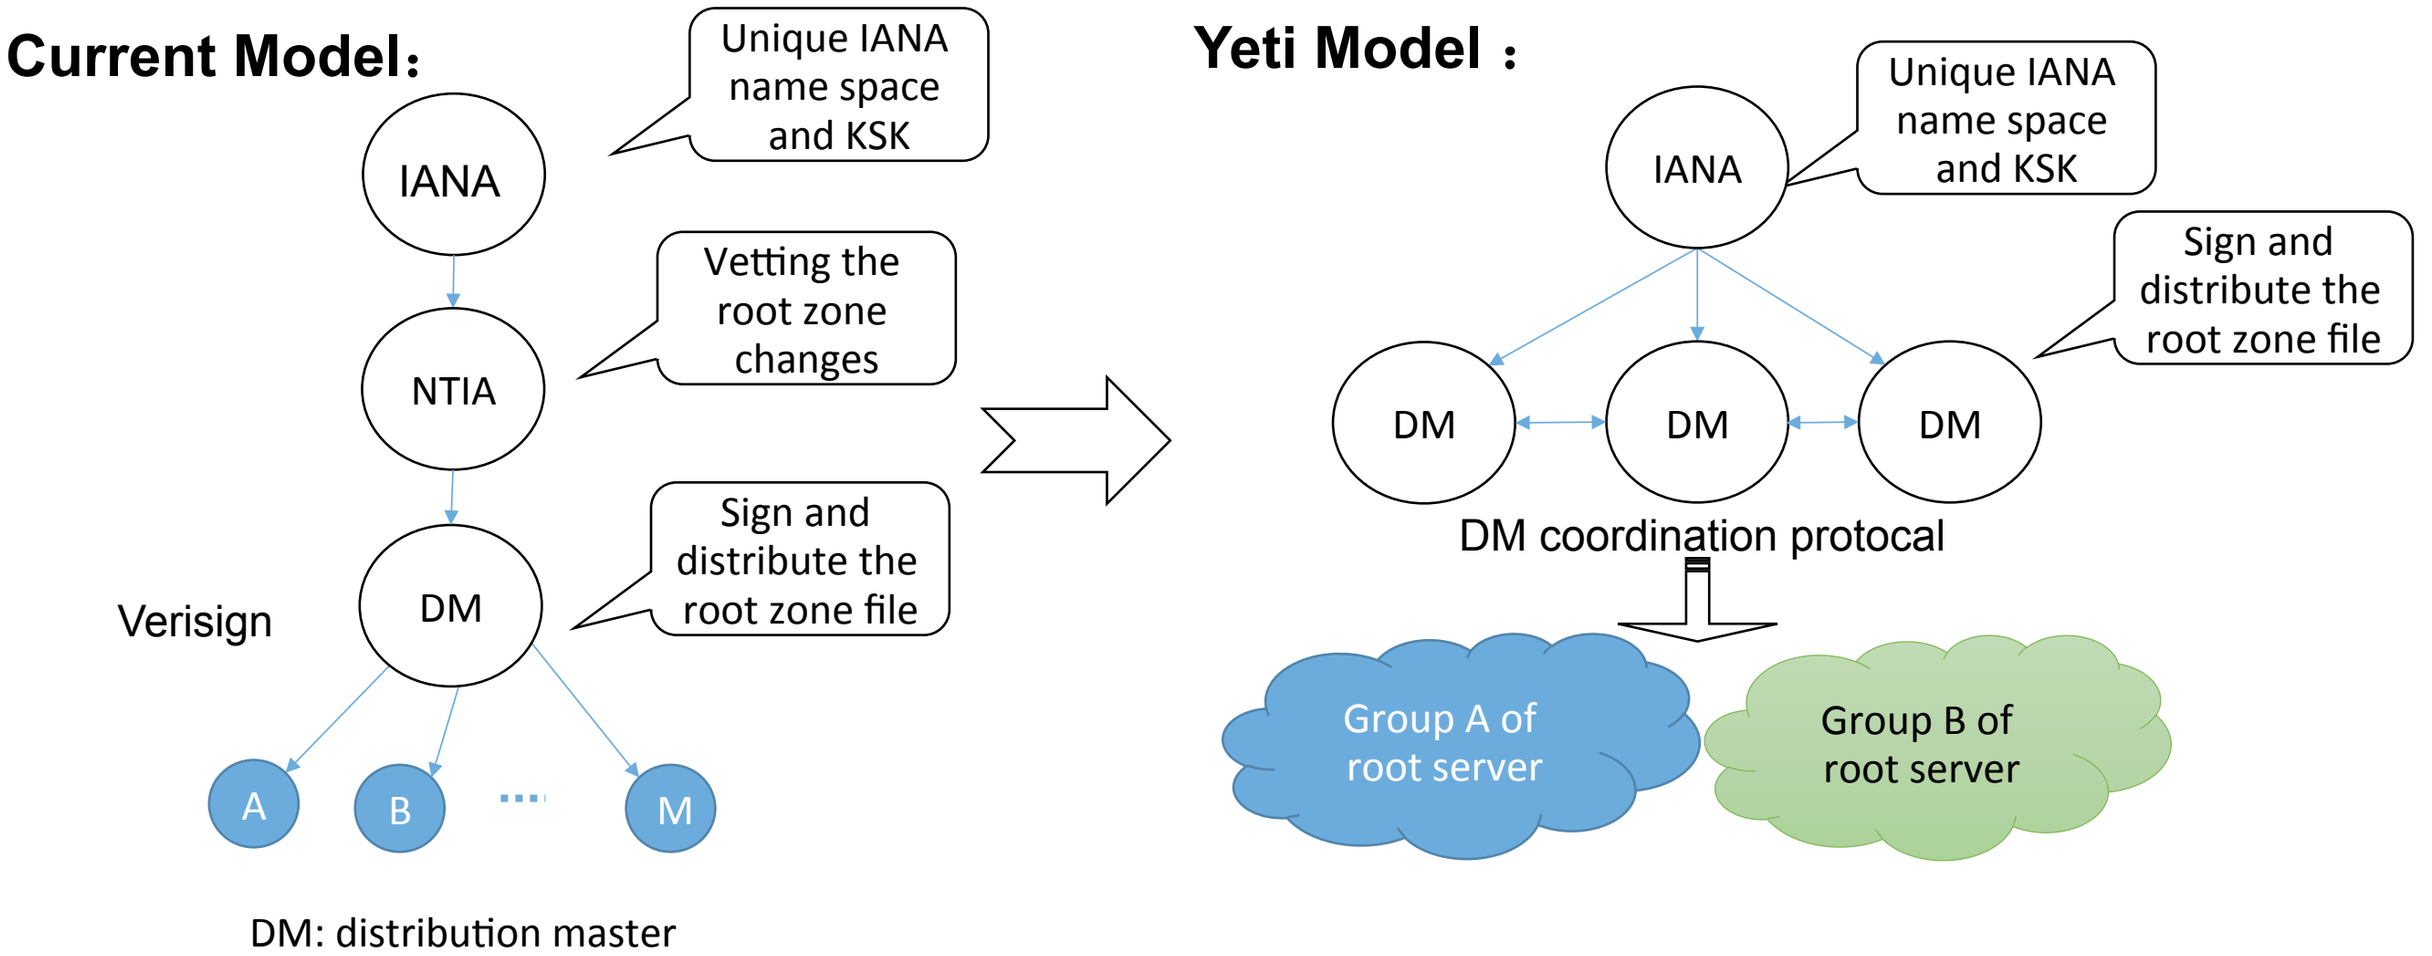

DM: distribution master

Three DMs setup and coordination

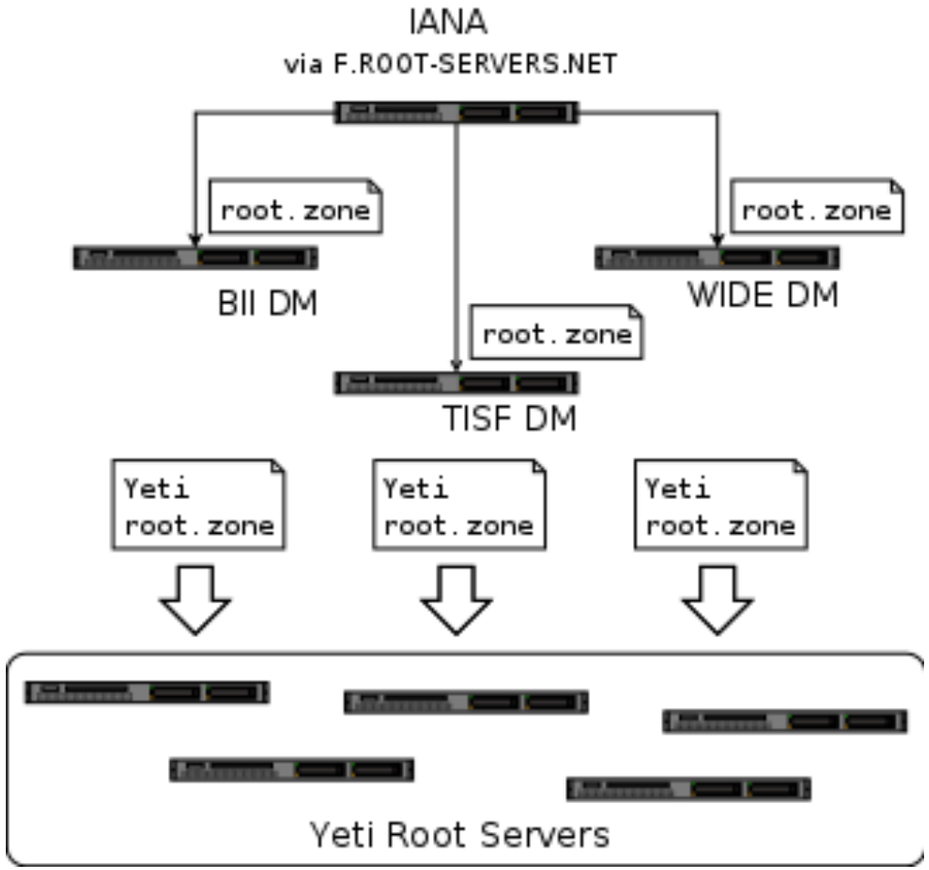

Timing setting

DM	Time
BII	*hour* + 00
WIDE	*hour* + 20
TISF	*hour* + 40

Time of Fetching the zone

Synchronizing

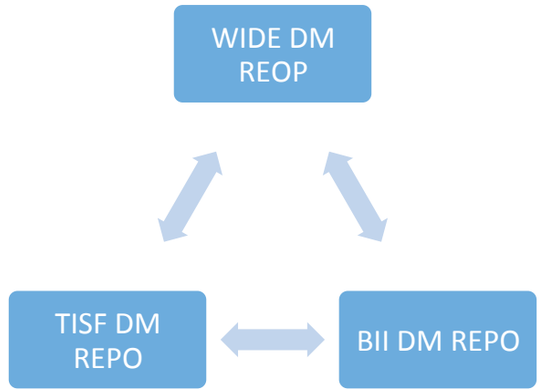

KSK, ZSK, server list,
IANA serial number

https://github.com/BII-Lab/Yeti-Project/blob/master/doc/Yeti-DM-Setup.md
https://github.com/BII-Lab/Yeti-Project/blob/master/doc/Yeti-DM-Sync.md

注：截至2015年10月28日

Yeti Root server

- Machine, OS system, DNS software

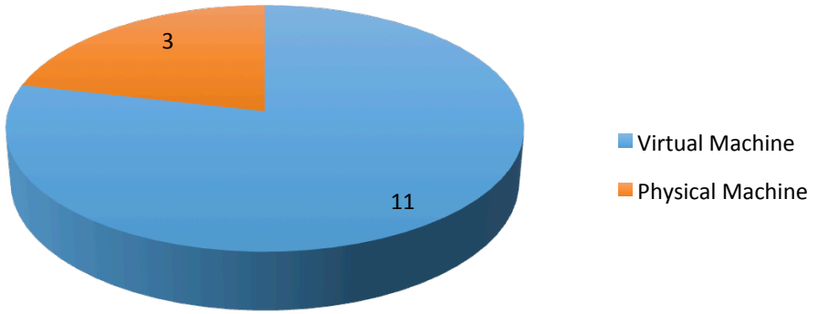

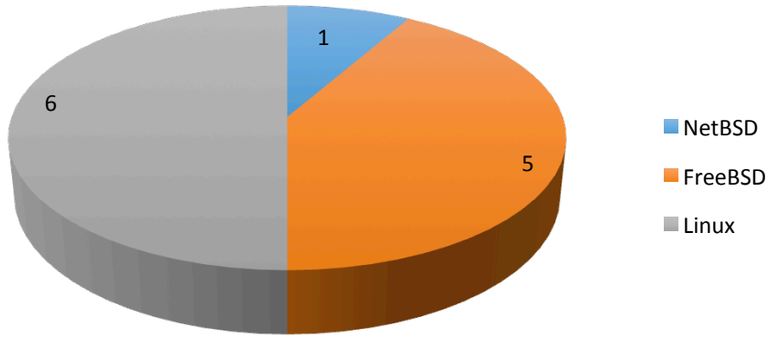

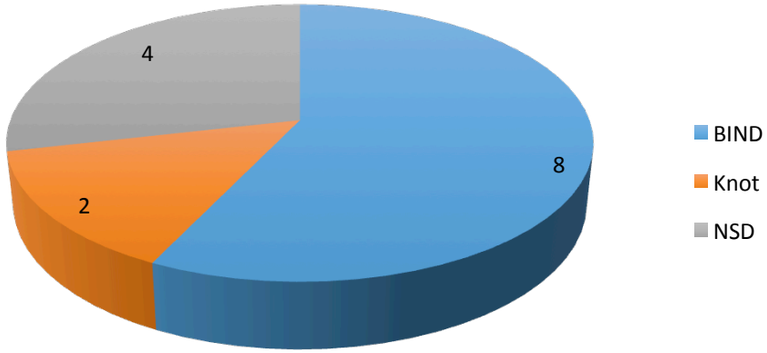

Bind9.10.3, BIND 9.10.2,
BIND 9.9.7-P2,BIND9.9.8

NSD 4.1.5, NSD 4.1.0

Knot 2.0.1,, Knot 2.1.0

Resolvers

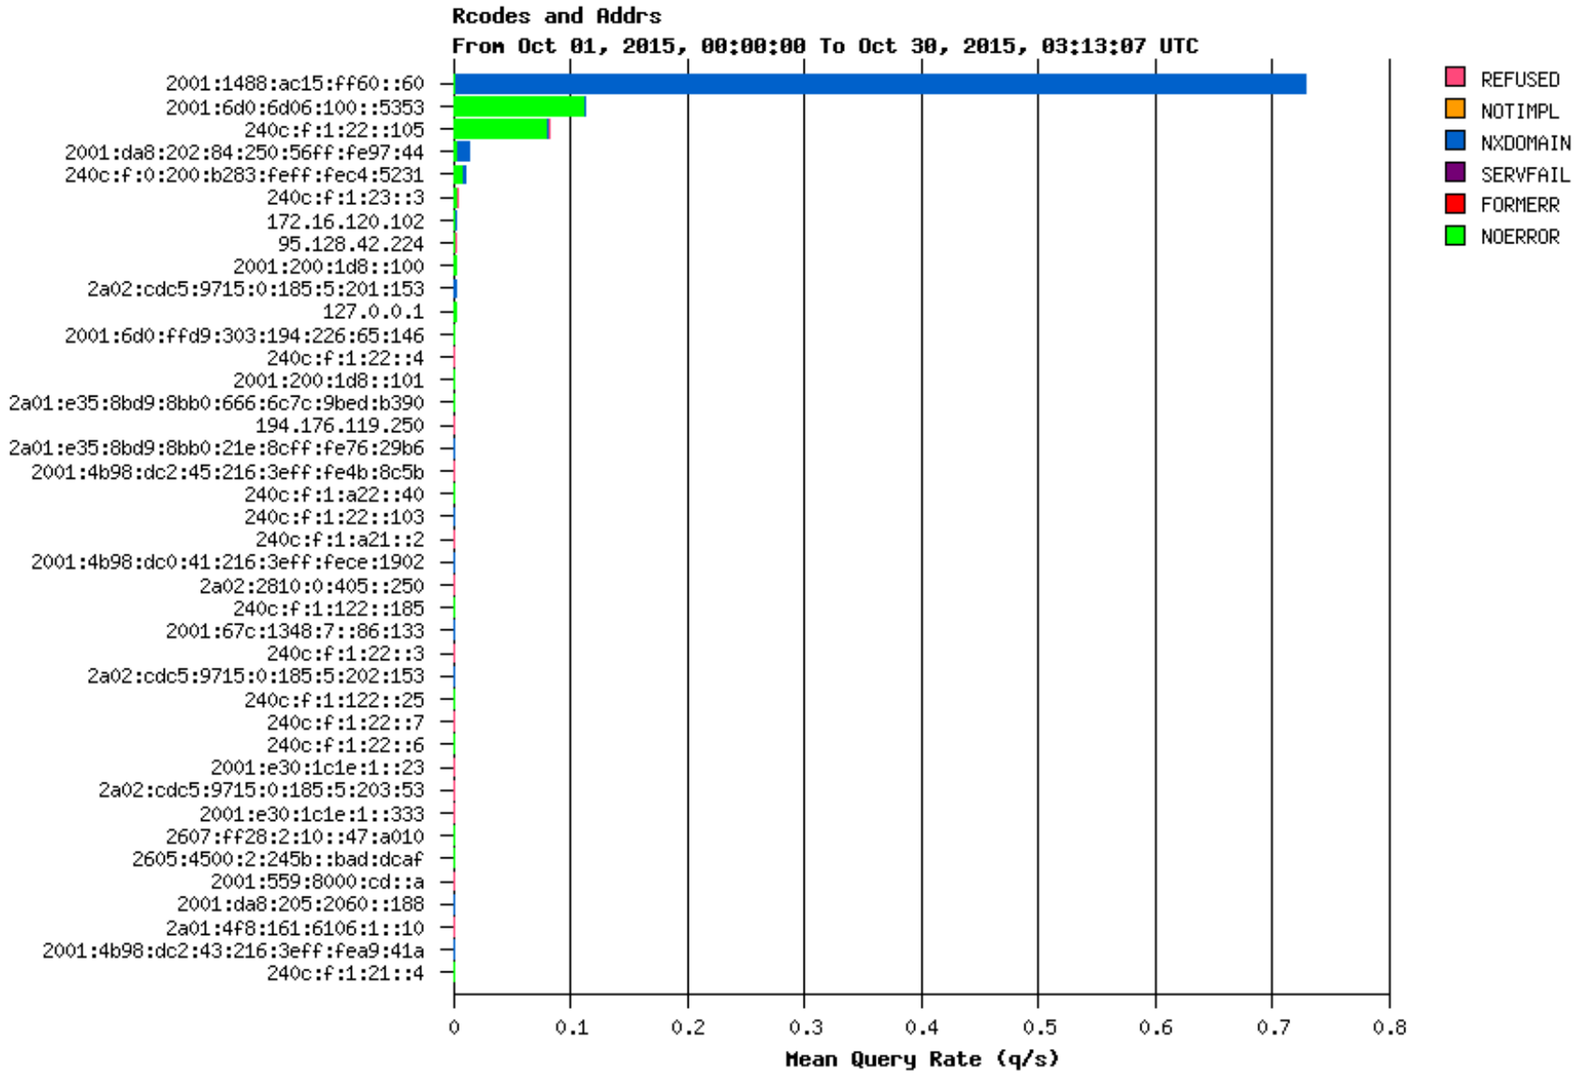
Rcodes and Addrs
From Oct 01, 2015, 00:00:00 To Oct 30, 2015, 03:13:07 UTC
2001:1488:ac15:ff60::60
2001:6d0:6d06:100::5353
240c:f:1:22::105
2001:da8:202:84:250:56ff:fe97:44
240c:f:0:200:b283:feff:fec4:5231
240c:f:1:23::3
172.16.120.102
95.128.42.224
2001:200:1d8::100
2a02:cdc5:9715:0:185:5:201:153
127.0.0.1
2001:6d0:ffd9:303:194:226:65:146
240c:f:1:22::4
2001:200:1d8::101
2a01:e35:8bd9:8bb0:666:6c7c:9bed:b390
194.176.119.250
2a01:e35:8bd9:8bb0:21e:8cff:fe76:29b6
2001:4b98:dc2:45:216:3eff:fe4b:8c5b
240c:f:1:a22::40
240c:f:1:22::103
240c:f:1:a21::2
2001:4b98:dc0:41:216:3eff:fece:1902
2a02:2810:0:405::250
240c:f:1:122::185
2001:67c:1348:7::86:133
240c:f:1:22::3
2a02:cdc5:9715:0:185:5:202:153
240c:f:1:122::25
240c:f:1:22::7
240c:f:1:22::6
2001:e30:1c1e:1::23
2a02:cdc5:9715:0:185:5:203:53
2001:e30:1c1e:1::333
2607:ff28:2:10::47:a010
2605:4500:2:245b::bad:dcaf
2001:559:8000:cd::a
2001:da8:205:2060::188
2a01:4f8:161:6106:1::10
2001:4b98:dc2:43:216:3eff:fea9:41a
240c:f:1:21::4
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
Mean Query Rate (q/s)
REFUSED
NOTIMPL
NXDOMAIN
SERVFAIL
FORMERR
NOERROR

Experimental traffic

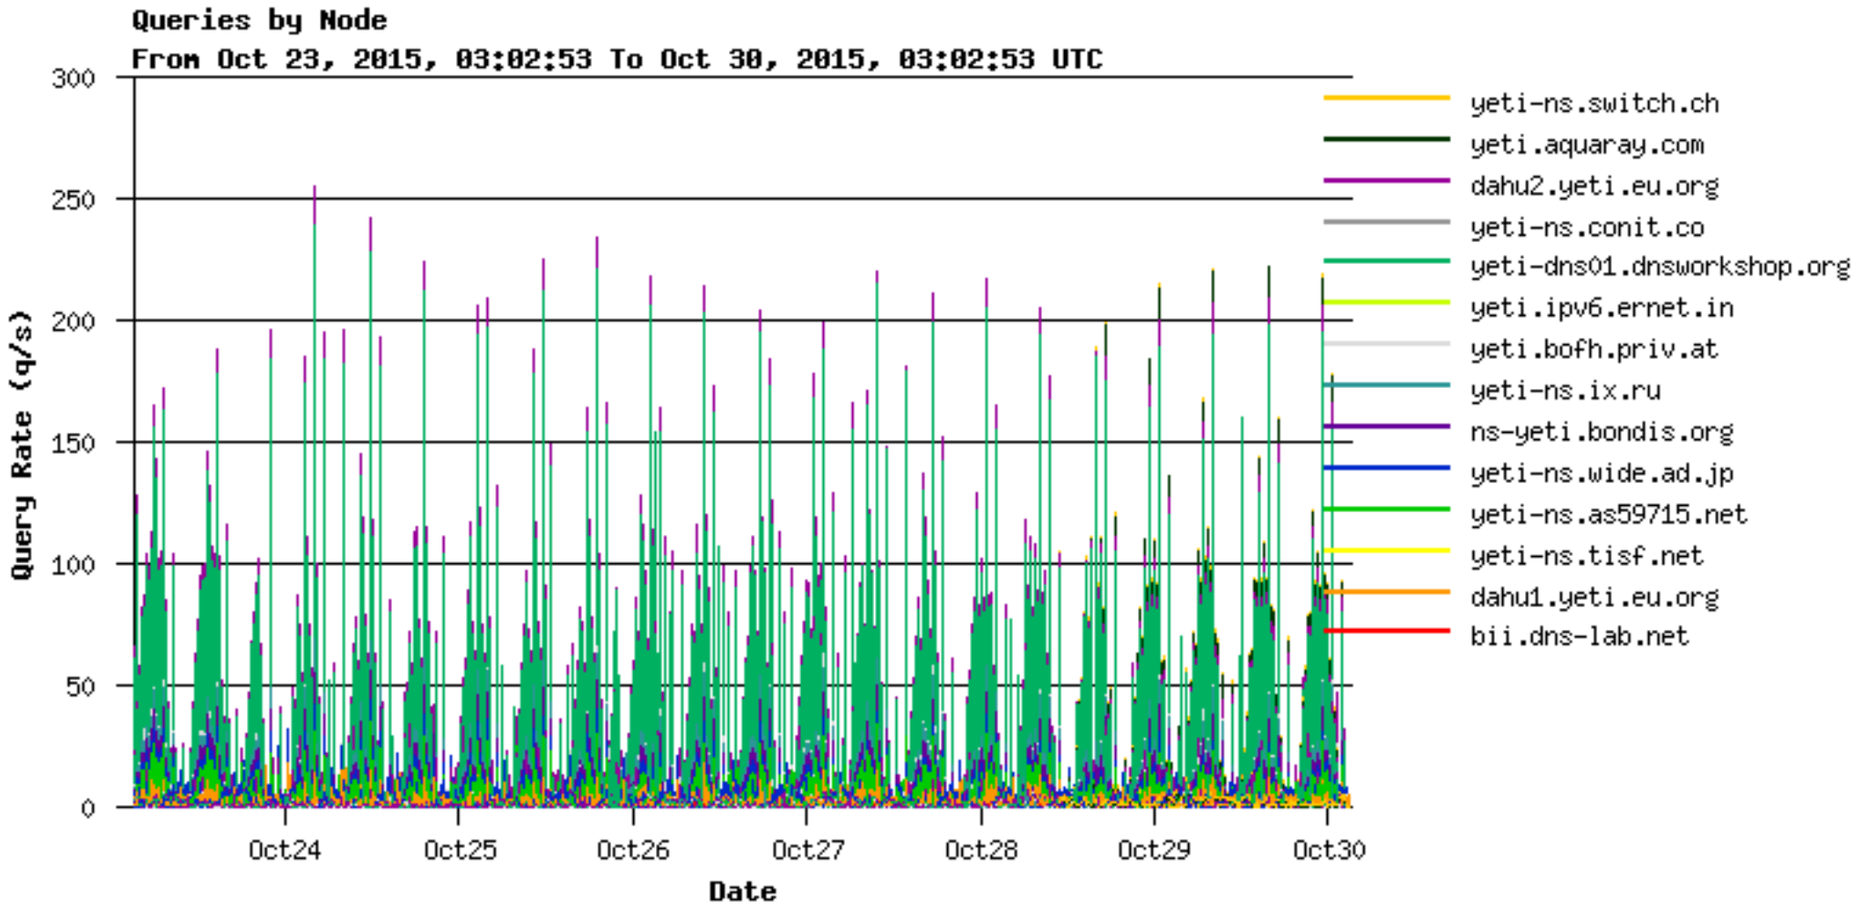

Resolvers and experimental traffic

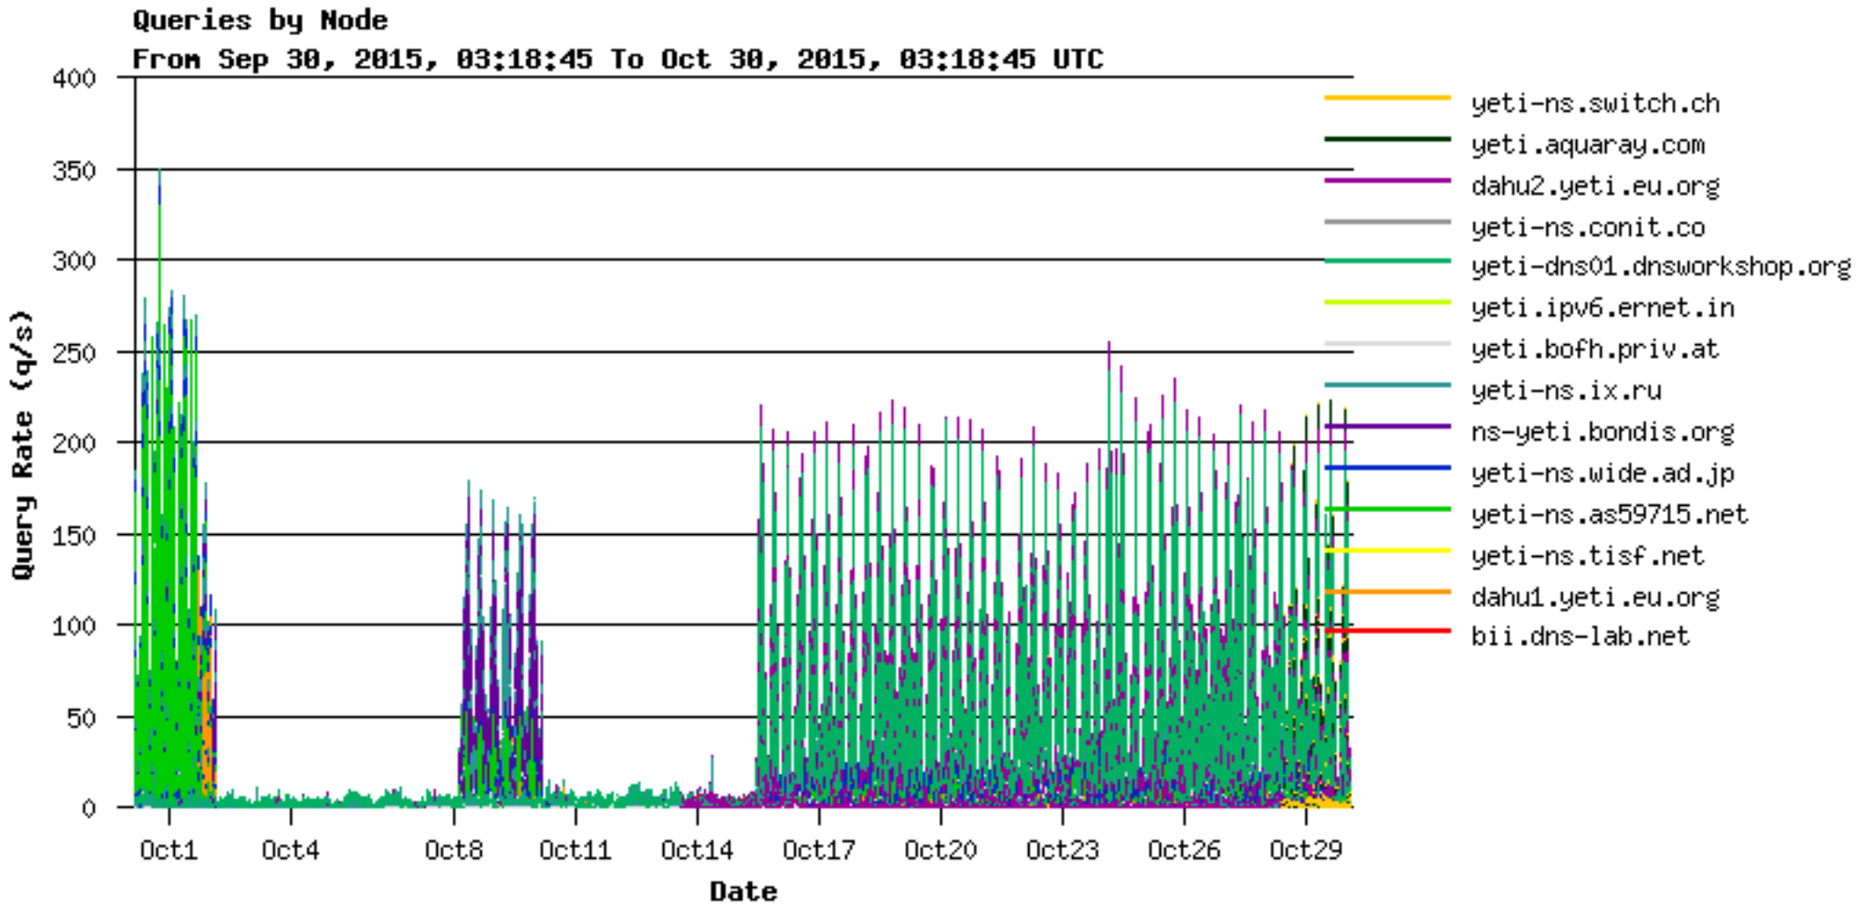

Experiment in BUPT

- Test the feasibility of Yeti concept in campus network with over 10,000 IPv6 active users
- Accessibility of one Yeti DNS root server from BUPT
- Setup a dual stack Recursive-DNS and DHCPv6 server in WiFi network of BUPT Buiding-3
- Setup IPv6-Yeti-test as one WiFi SSID
- Distribute R-DNS to IPv6 users via DHCPv6 server
- Encourage student to try
- Collect access information for further analysis

System Ready for Yeti Experiment

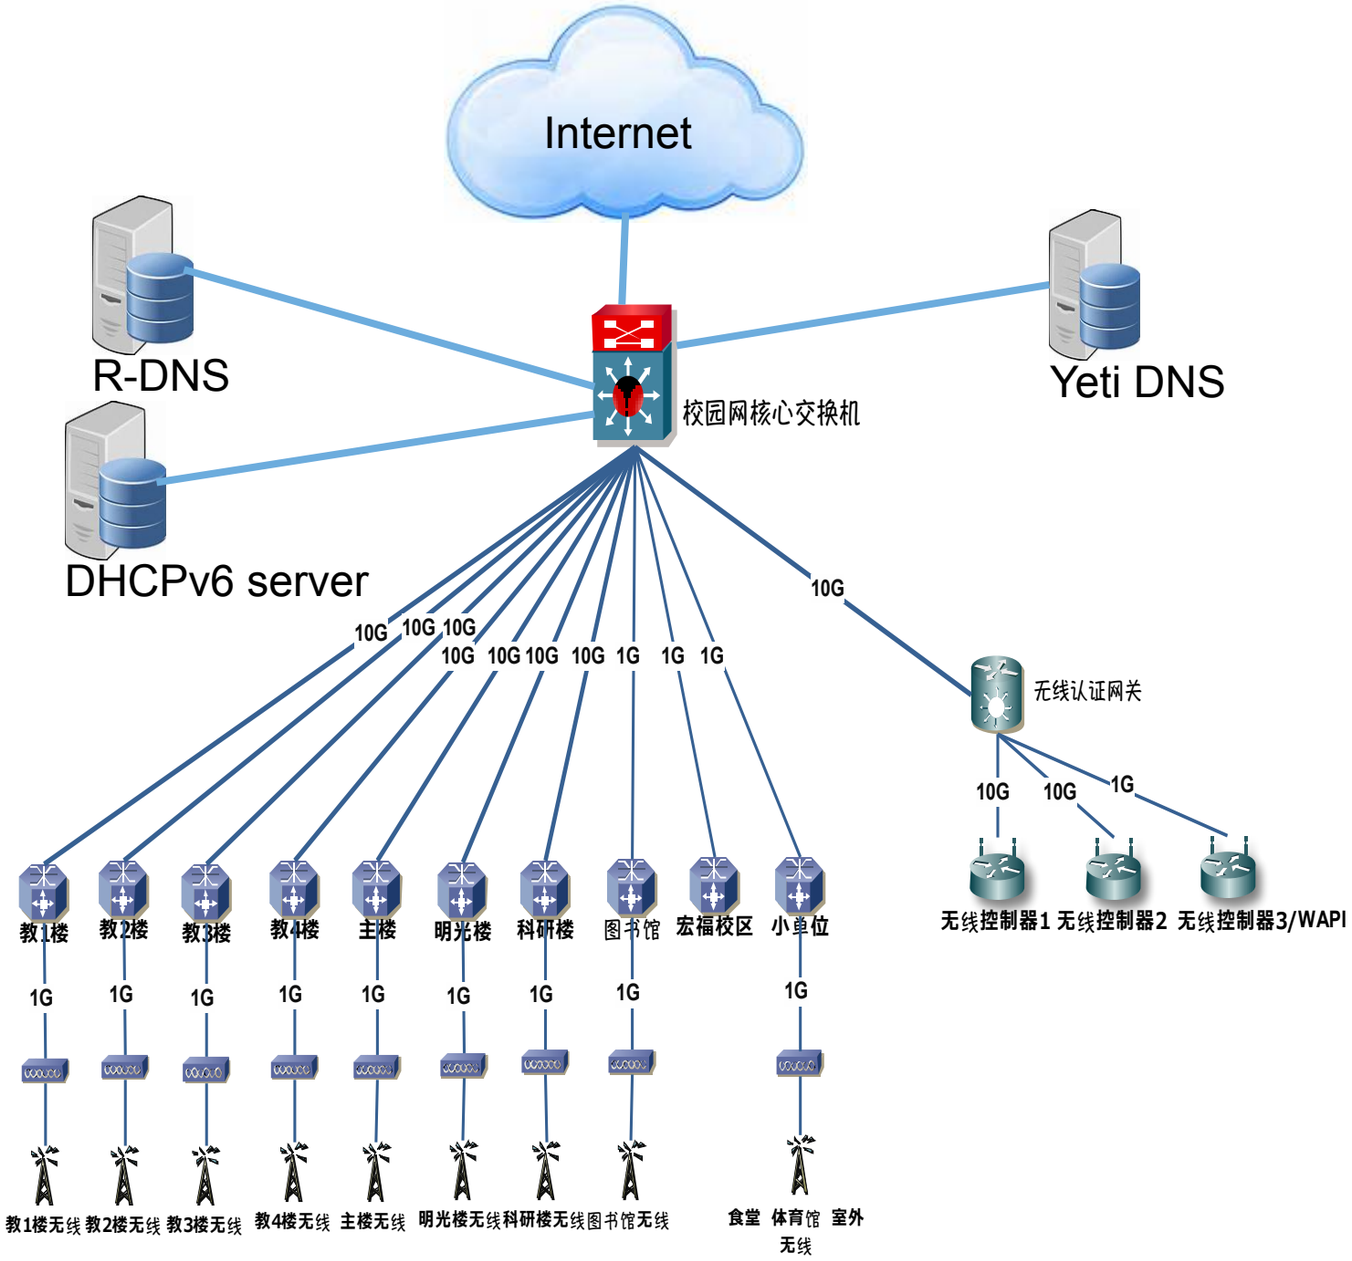

Yeti R-DNS Traffic Analysis

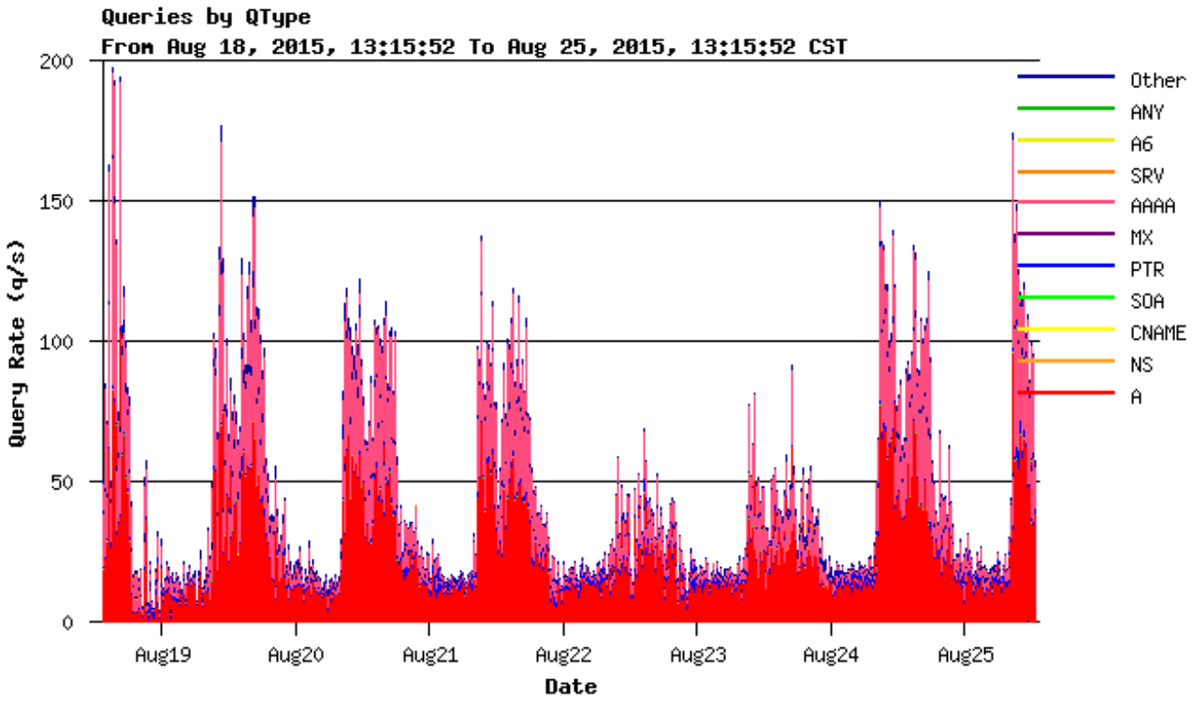

Peak: 205 qps

Major Qtype: AAAA，A

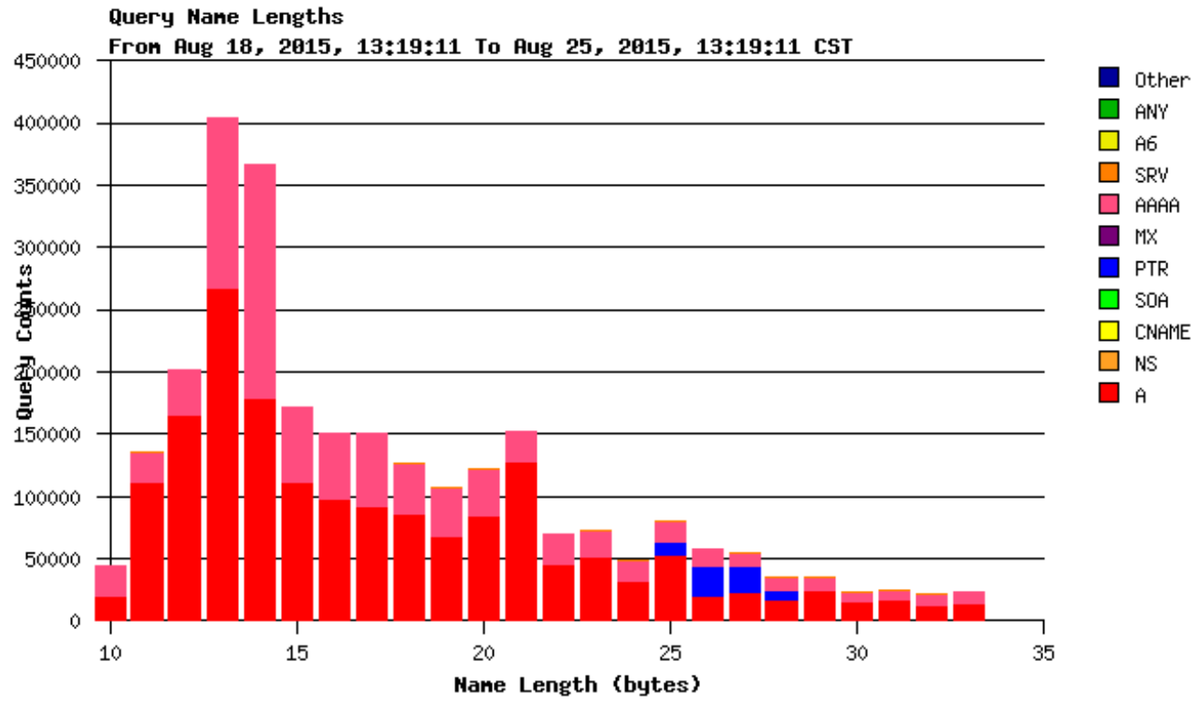

AAAA query：37%
A query: 58%
Other Qtype: 5%

Data collection and monitoring

- DSC page in Yeti website： http://yeti-dns.org/statistics.html
- Health monitoring page： http://yeti-dns.org/yeti_server_status.txt
- Yeti debug page： http://yeti-dns.org/resource/yeti-bug.txt

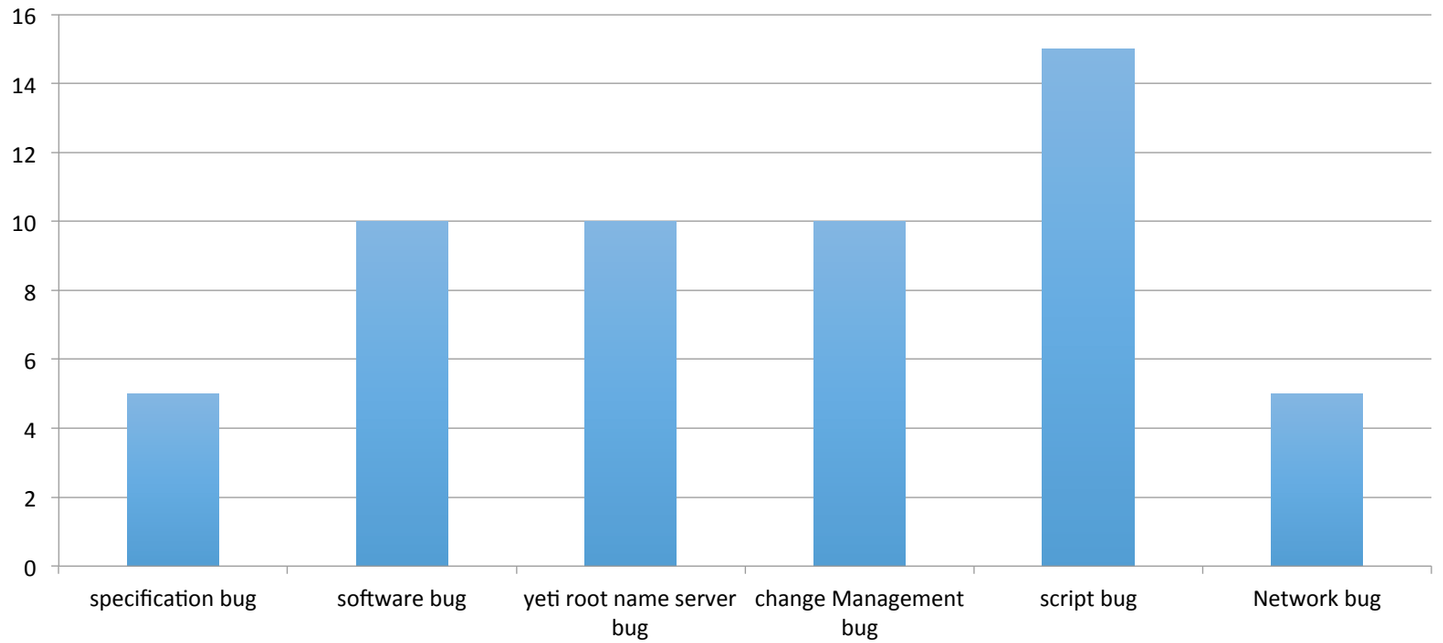

Findings & bugs

- Root Glue issues (**Resolved**!)
  - Current root servers answer for the **root-servers.net** zone, but Yeti root server dose not (independent domain), Without this setup, BIND 9 does not include glue in answers to priming queries.
  - Resolved！ With a patch for BIND9
- Related issues
  - .arpa. zone issue
  - Unused Glue issue

Findings & bugs

- A Bug in Knot 2.0 (**Resolved!**)
  - Knot 2 compress even the root. It is useless since it is a zero-length label, only one byte. Knot 1.6 used for K-root do not do that
  - Resolved! https://gitlab.labs.nic.cz/labs/knot/issues/398

- DNSCAP issues
  - Current DNSCAP(both DNS-OARC and Verisign versions) was observed losing some packet which is not ideal

Findings & bugs

- Failure on Root server zone transfer
  - Some authoritative server on some VPS failed to pull the zone from Distribution Master
  - One fact : TCP fails to respect IPV6_USE_MIN_MTU (**draft-andrews-tcp-and-ipv6-use-minmtu-04**
  - Another fact : there are bugs in Virtual machine software failing to receive IPv6 fragments (One Example: FreeBSD on VMware ESXI 5.5)
- Recommendation:
  - 1) Change the IPV6_USE_MIN_MTU setting on server side to 1500 (DM in Yeti case)
  - 2) Or set TCP MSS to 1280 on client side (Root server in Yeti case)

In conclusion

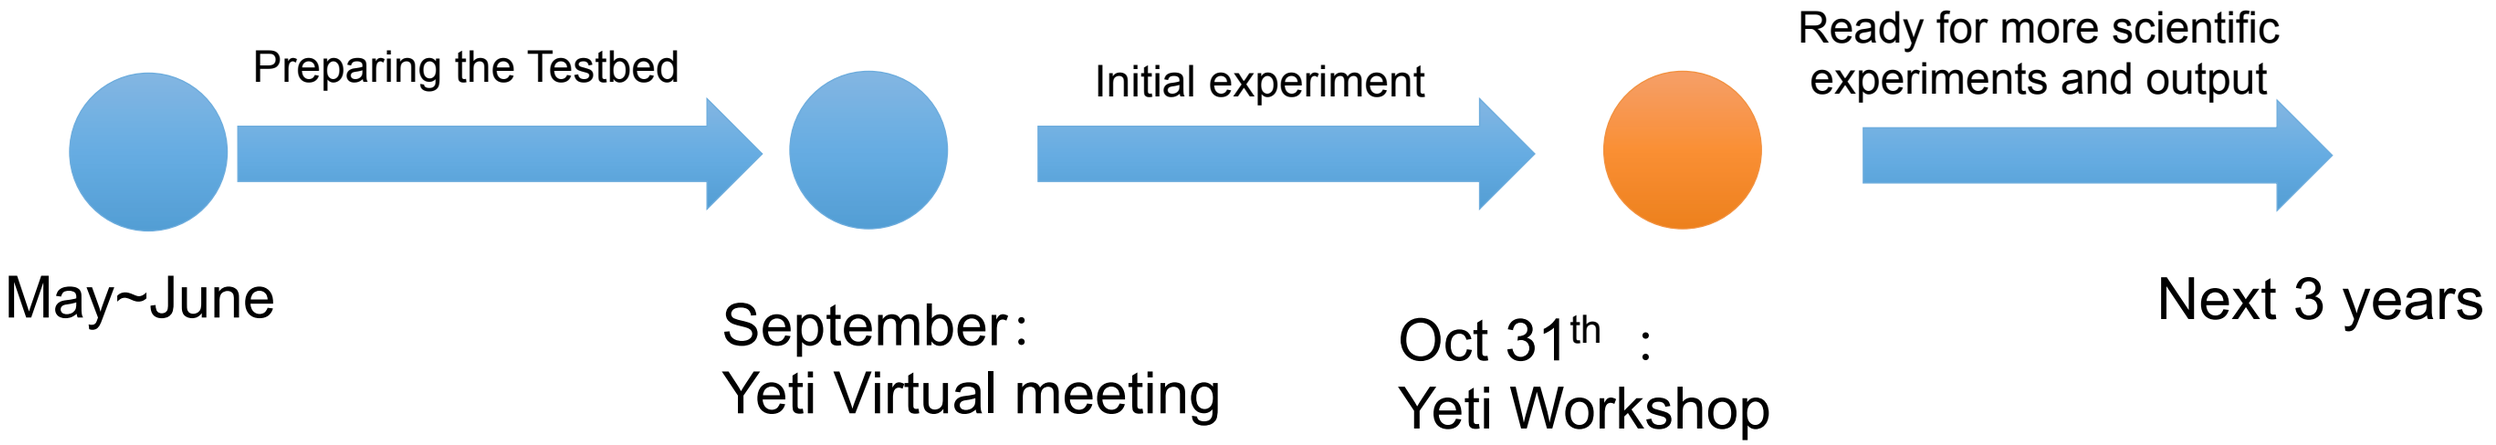

- All most finish the engineering part of Yeti testbed
- Three DMs are running, more than 13 root servers are running
- Lack of traffic , resolvers, and end-to-end measurement
- Experiments agenda expected

Thank you! Any Questions?

More information on website:

http://yeti-dns.org/

https://github.com/BII-Lab/Yeti-Project

http://lists.yeti-dns.org/mailman/listinfo/discuss