Yeti DNS: Current Status

Yeti Virtual Meeting

Davey Song (宋林健) &
Shane Kerr / BII Lab

2015-09-07 / Internet

BII 天地互连

Outline

- Introduction of Yeti
- The Current Status and Challenges
- Some Technical Findings

What is Yeti?

- 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
  - Has its own DNSSEC signing and validation keys
- System is intended for Internet-scale *science*

Bii 天地互连

Things That Yeti is Not...

- ***NOT*** research into alternatives to the IANA root/namespace

- ***NOT*** interested in policy or political work
  - Although such work may eventually result from Yeti findings

Who: Roles and Participants

- **Coordinators and DMs**
  - WIDE, BII, and TISF
- **Root Server Operators**
  - 11 root servers are operating, 5 have shown their interest
- **Participants from Client Side**
  - Research labs
  - DNS software implementers
  - Developers of CPE devices, IoT devices, …
- **Traffic and Data Collector**
  - BII
- **Experiment Proposer**
  - Any interested parties or individuals

Bii 天地互连

Why:
Yeti Problem Space (1 of 2)

- **DNS Centralization vs. Network Autonomy**
  - External dependency
    - Local services rely on external root services
    - Require external management
  - Surveillance risk
    - Information leakage cause by lookup at the DNS root
- **Can IPv6-only DNS work?**
  - Some DNS servers which support both A & AAAA records (IPv4 & IPv6) still do not respond to IPv6 queries
  - IPv6 introduces larger MTU (1280 bytes), but a different fragmentation model

Why:
Yeti Problem Space (2 of 2)

- **Are we ready for KSK rollover, or not?**
  - Not all resolvers are compliant with RFC 5011
  - Larger packets may introduce risks during KSK/ZSK rollover
- **Renumbering Issues**

“One Namespace, Many Circles” Model

IANA Model

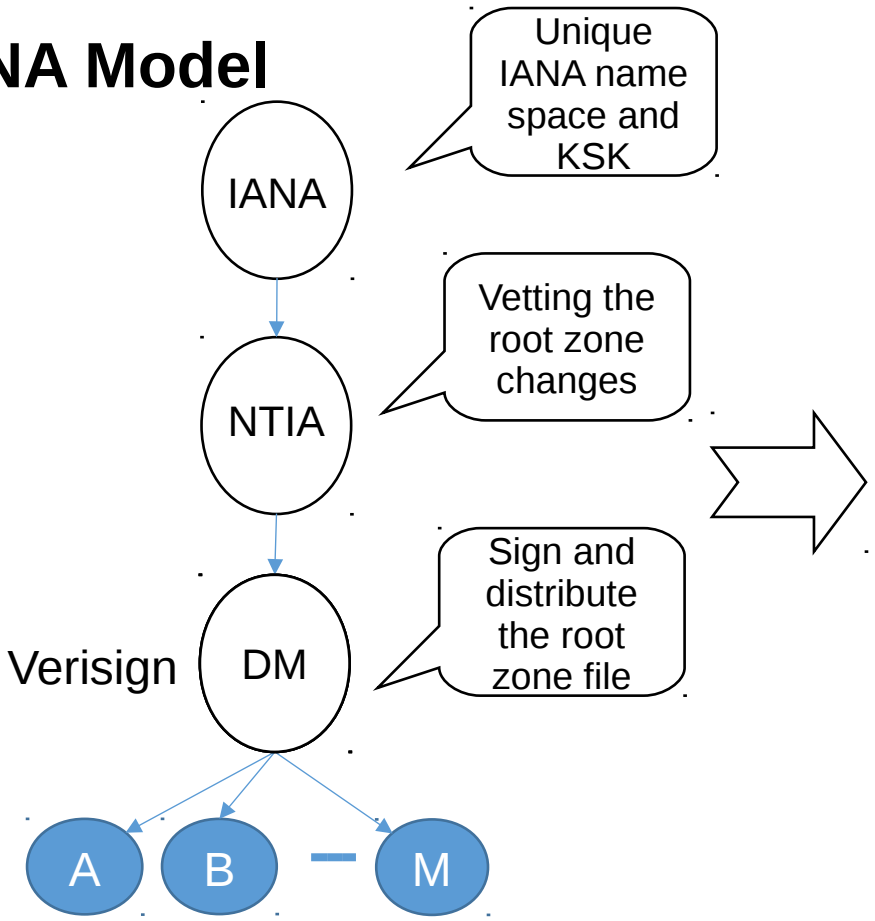

Yeti Model

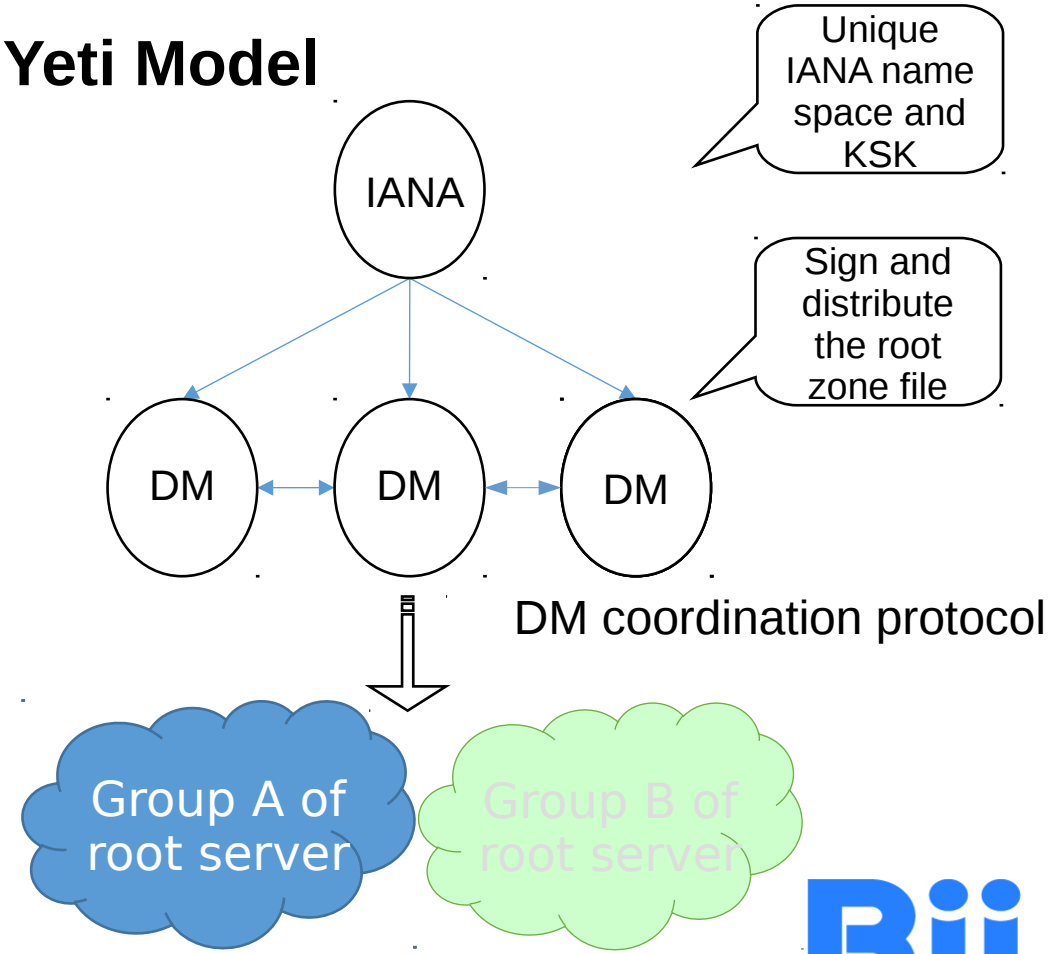

DM: distribution master

Bii 天地互连

Yeti DNS: How It Works

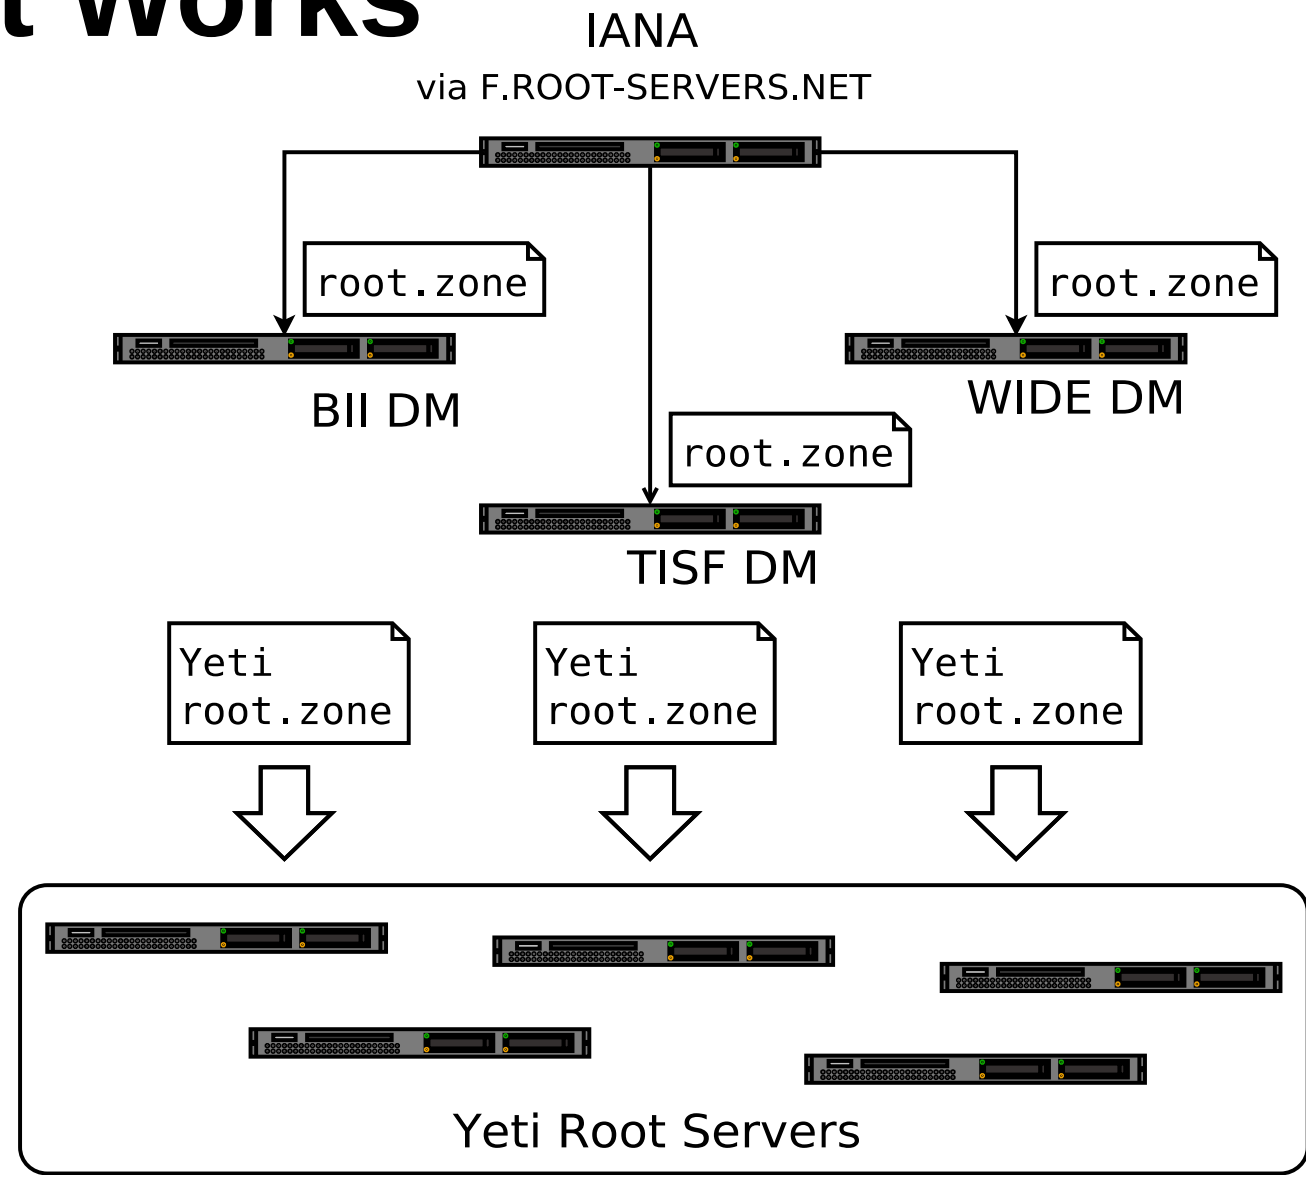

https://github.com/BII-Lab/Yeti-Project/.../doc/Yeti-DM-Setup.md

Bii 天地互连

Current Status

- System functioning
- Infrastructure up
  - Web site, http://yeti-dns.org
  - Mailing lists, DSC, RT ticketing, ...
- Docs & scripts in GitHub
  - https://github.com/BII-Lab/Yeti-Project
- Currently gathering Yeti root operators
  - 11 up now

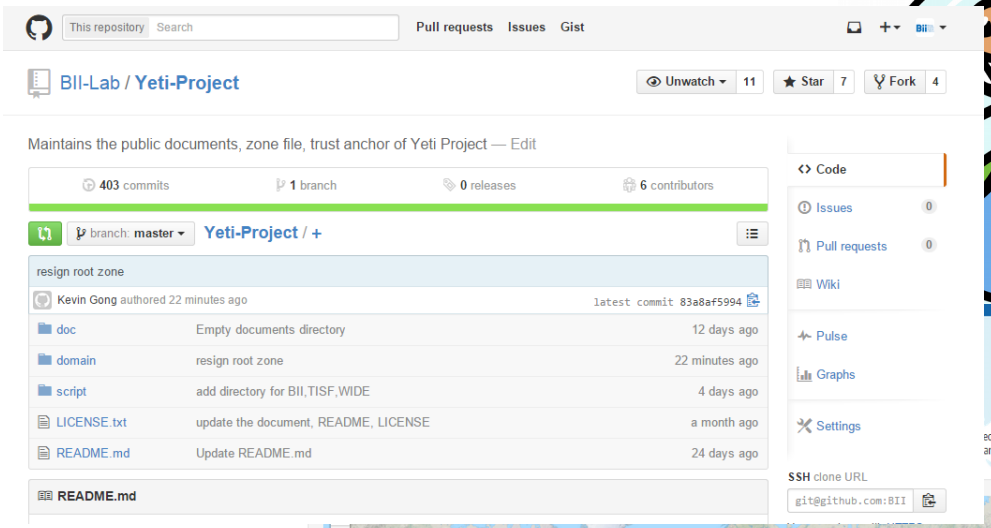

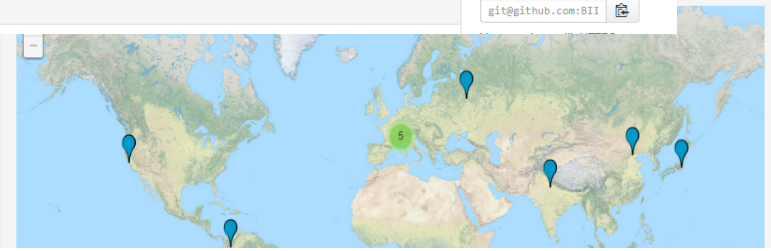

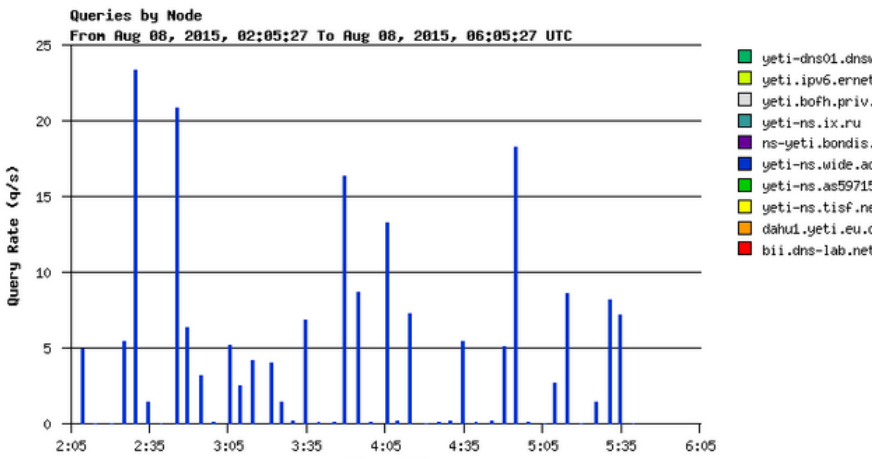

天地互连

Where and Who We Are

Bii 天地互连

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 Building-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

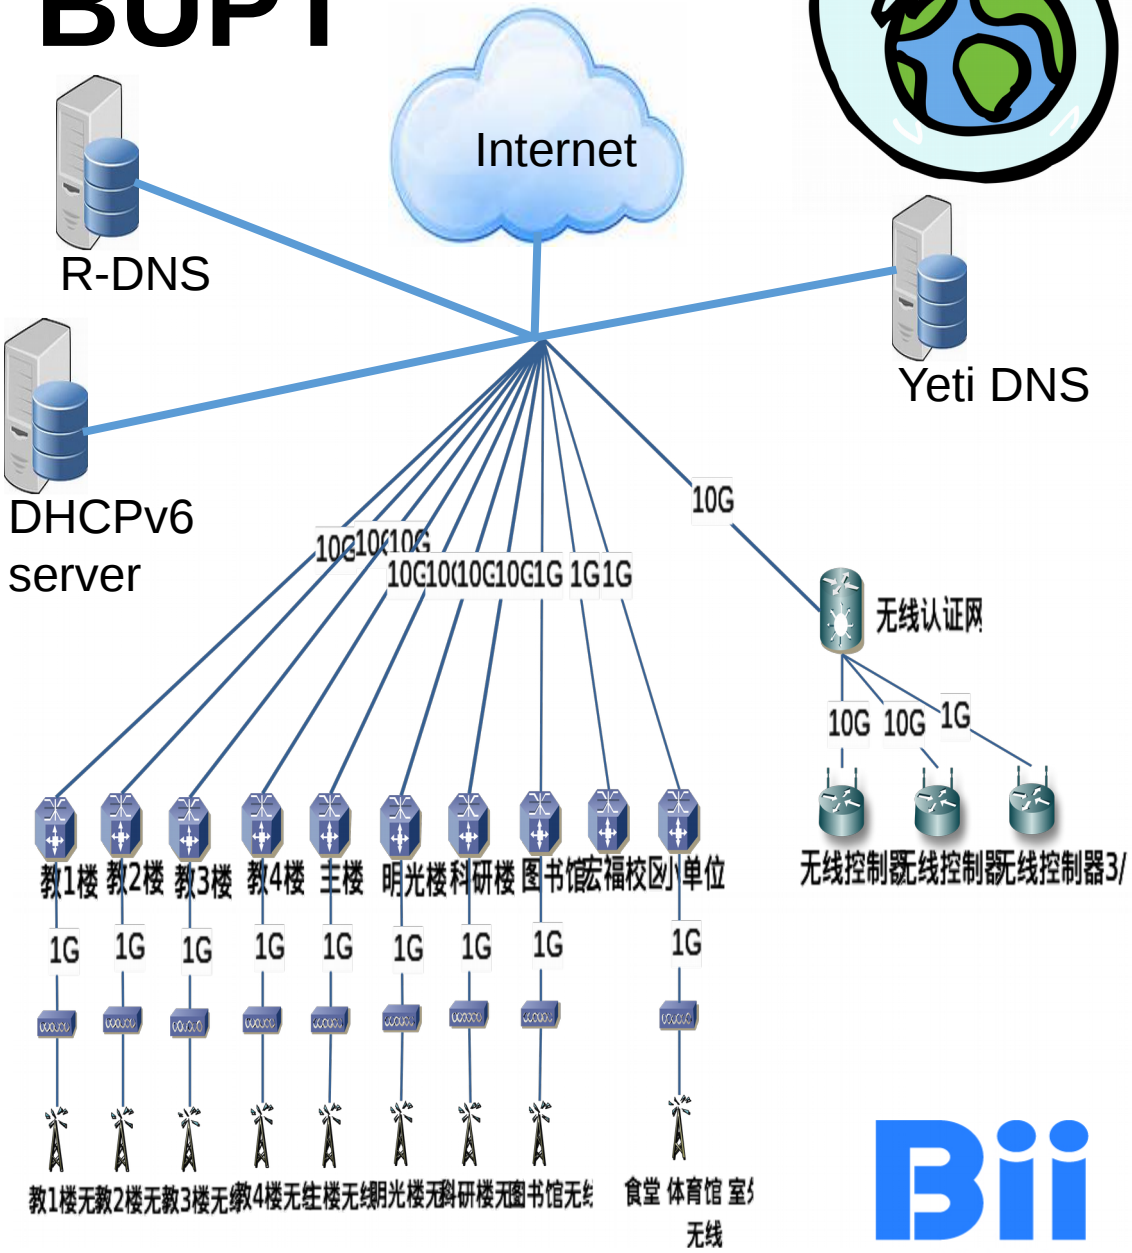

System Ready for Yeti Experiment

天地互连

Yeti R-DNS Traffic Analysis

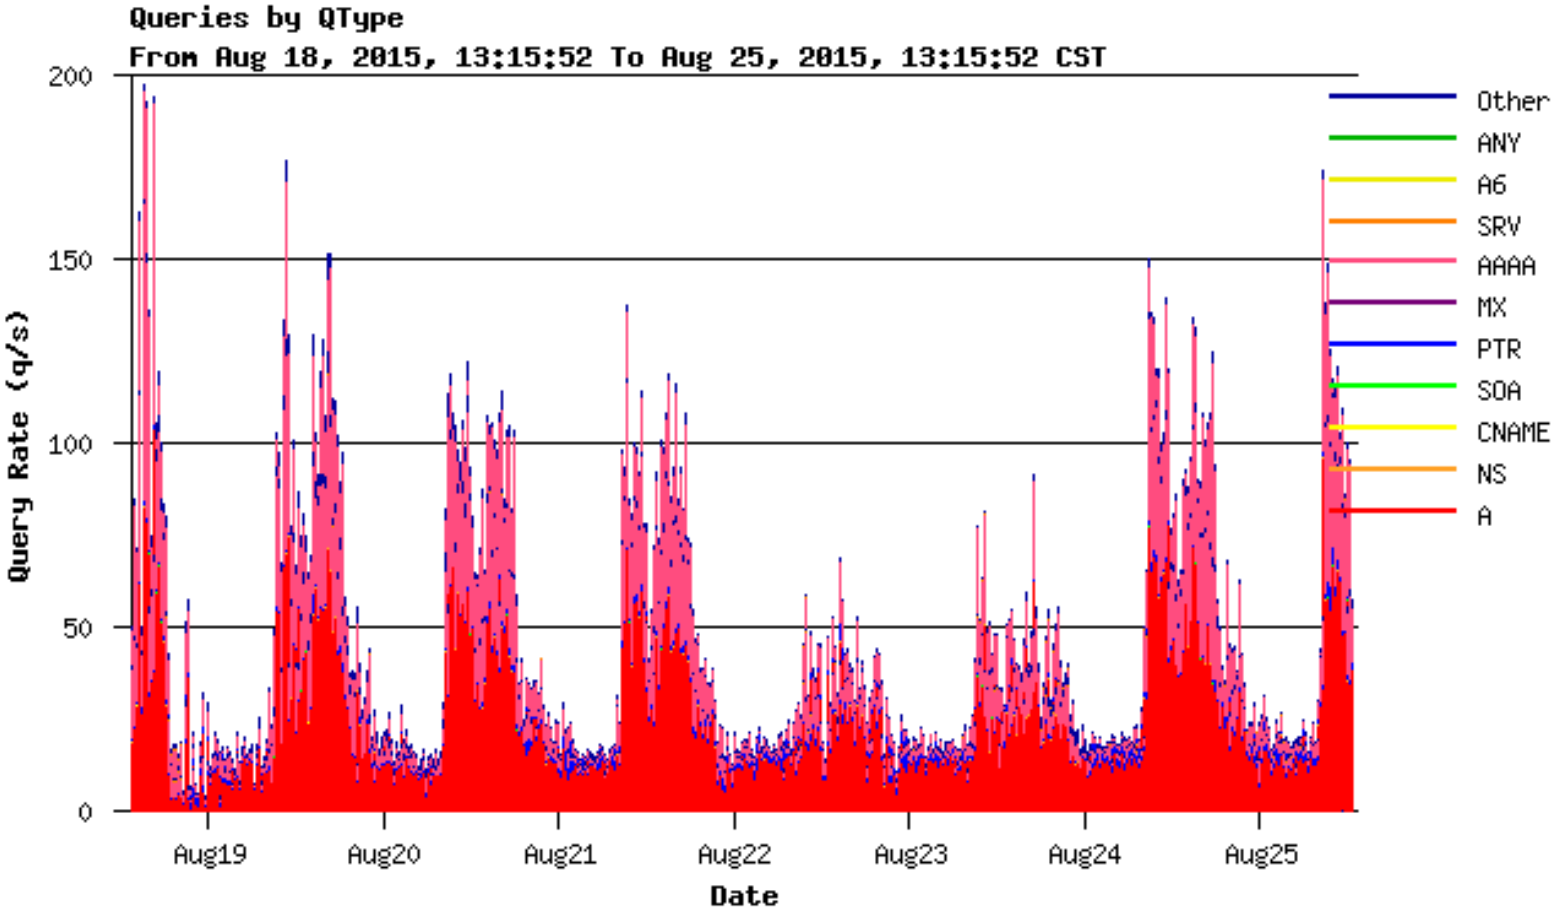

Peak: 205 queries/second

Major QTYPE: A, AAAA

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

Bii 天地互连

Yeti R-DNS Traffic Analysis

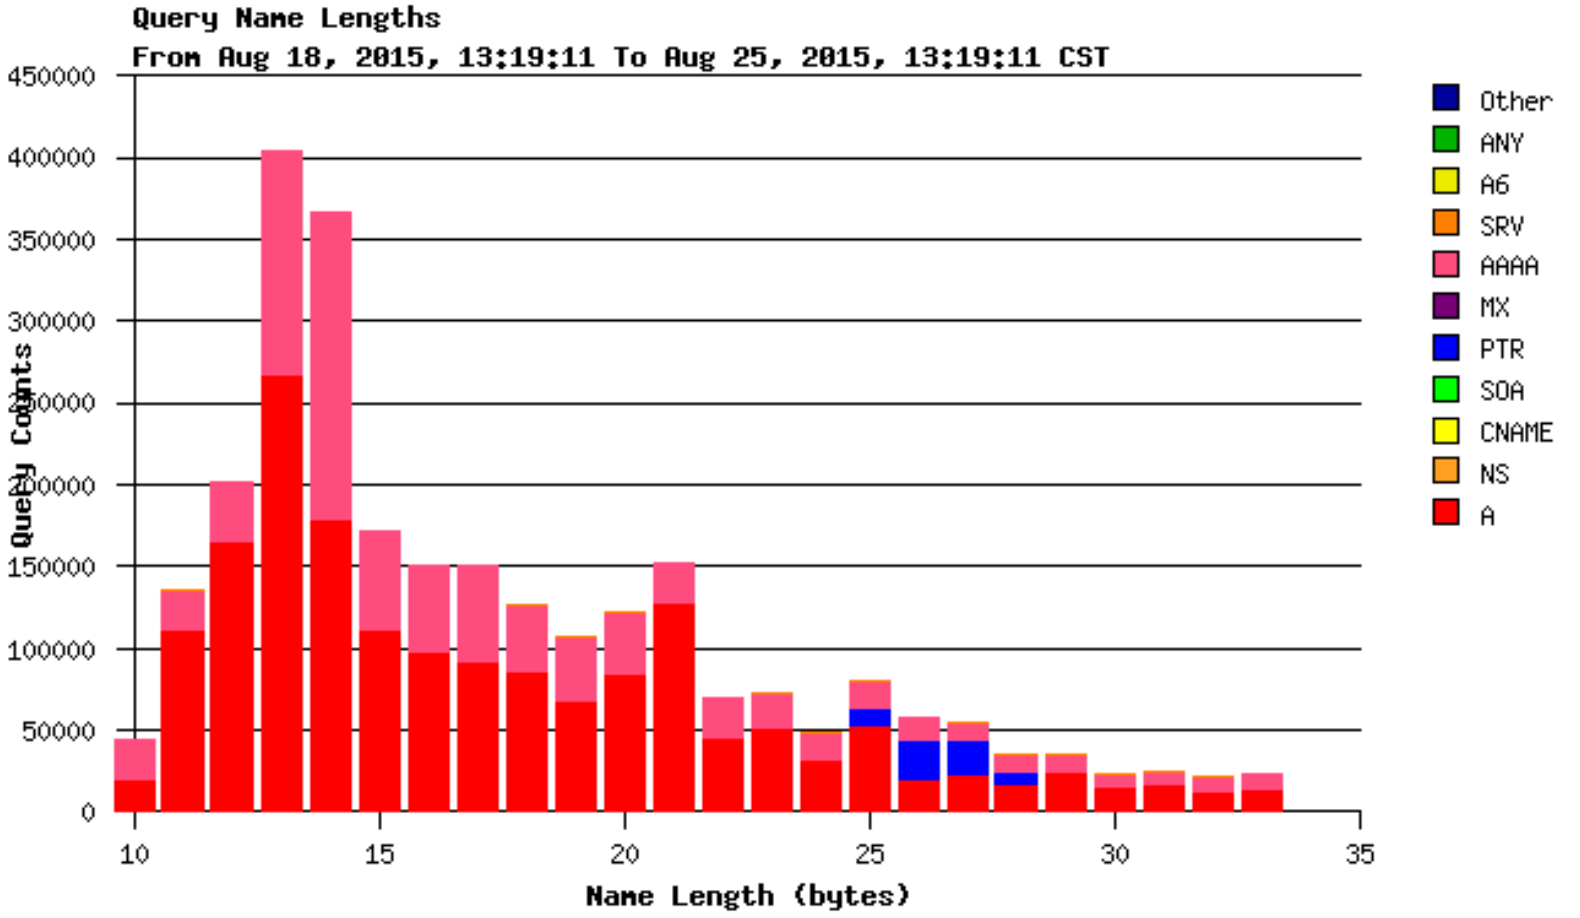

Bii 天地互连

Current Yeti traffic status

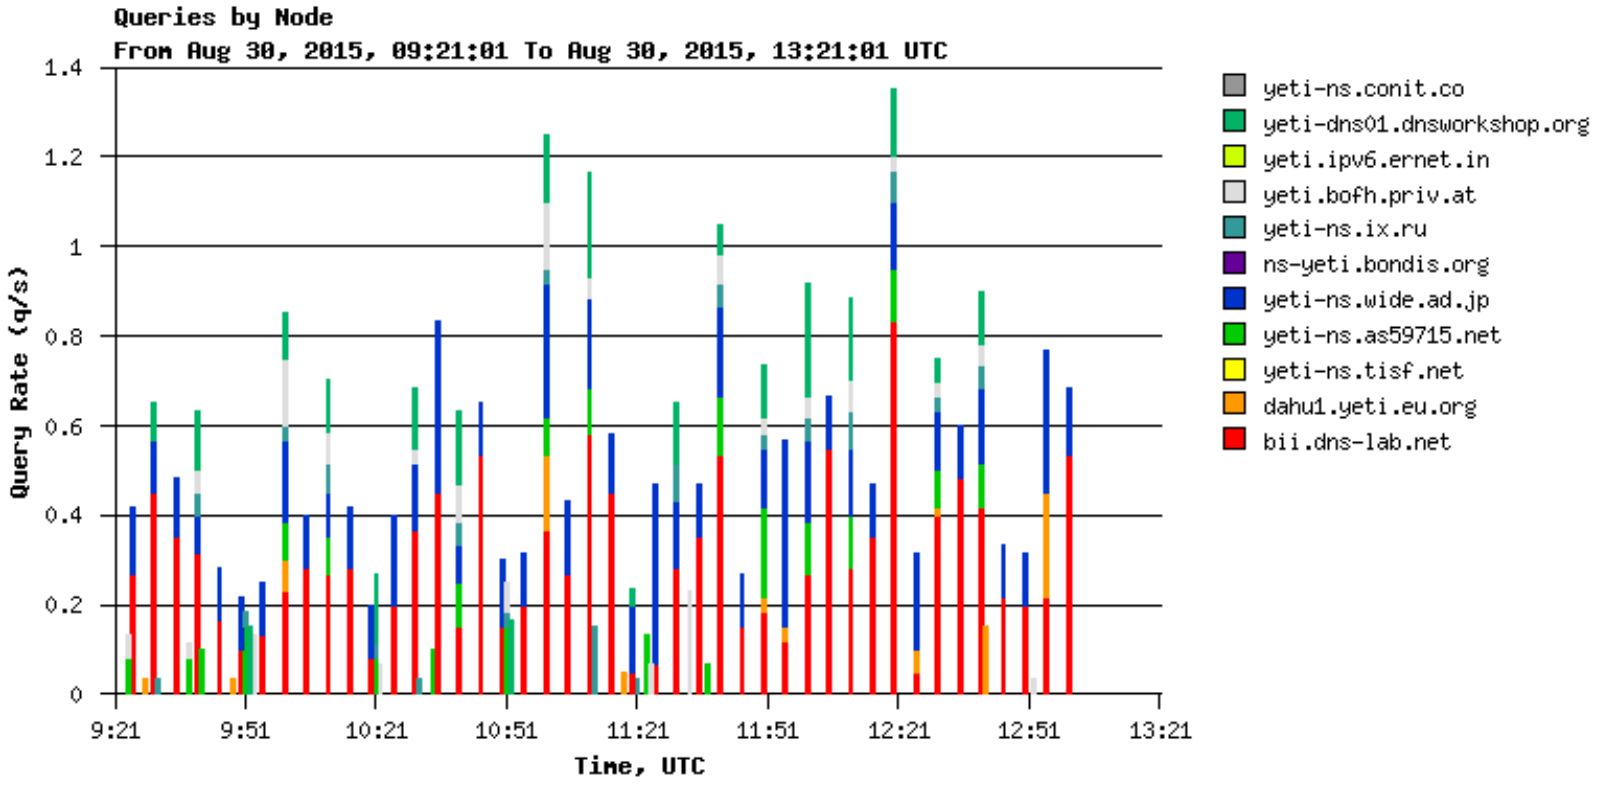

Query rate of Yeti root system (1.4 queries/second)

Bii 天地互连

Current Yeti traffic status

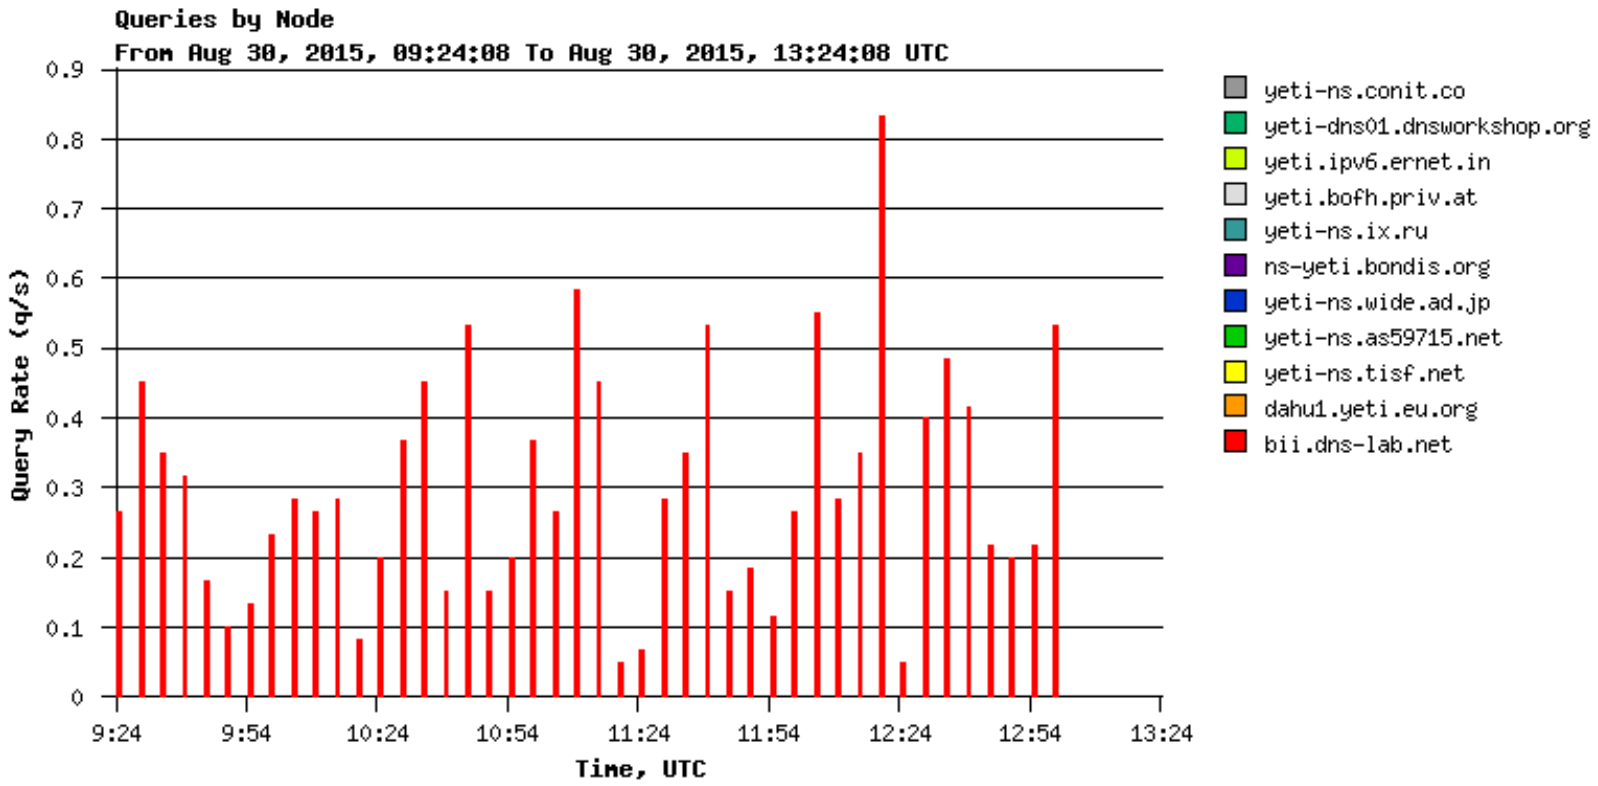

Query rate at BUPT (0.86 qps)

Bii 天地互连

Some Findings

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

Some Findings

- 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 does 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 packets

Bii 天地互连

Challenges

- We need more robust system (system/network failures)
- To resolve the inconsistency problem in Yeti with multi-DM setting
- Need more Yeti resolvers (experimental traffic)
- Need enough authoritative serve (Yeti operators)
- ...and more!

Where We Are {At,Going}

May/June

Kick off and start to invite participants

September 7

Yeti virtual meeting

October 31

Yeti Workshop & call for more experiment proposals

Next 3 years

Bii 天地互连

Next Steps (1 of 2)

- Get ”enough” Yeti root servers
- Introduce experiment traffic from universities and research labs
- Design and conduct some experiments in Yeti Testbed
- Deliver some experiment report and feedback to the community and/or standard bodies
  - IETF DNS-related WGs
  - ICANN KSK rollover design team
  - ICANN RSSAC Caucus

Next Steps (2 of 2)

- Outreach to CPE and middlebox vendors

Thank you!

Bii 天地互连

Yeti DNS & SCIENCE!!!

Yeti Virtual Meeting

Shane Kerr / BII Lab
2015-09-07 / Internet

BII 天地互连

Yeti is for Science

To benefit the Internet development as a whole, the proposal of Yeti Project is formed to build a parallel experimental live IPv6 DNS root system to discover the limits of DNS root name service and deliver useful technical output. Possible research agenda will be explored on this testbed covering several aspects but not limited to:

- IPv6-only operation
- DNSSEC key rollover
- Renumbering issues
- Scalability issues
- Multiple zone file signers

Proposed Experiment Protocol

1. Proposal
2. Lab Test
3. Yeti Test
4. Report of Findings

https://github.com/shane-kerr/.../Experiment-Protocol.md

Experiment: KSK roll

- Do a ”normal” KSK roll
- Based on KSK Double-DS roll in RFC 6787
  - Although no DS, since there is no parent
- Will result in large packets (> 1280 bytes)
- Takes at least 60 days to complete
  - Accursed RFC 5011 hold-down timer
- Method documented in GitHub branch

https://github.com/shane-kerr/Yeti-Project/.../KSK-rollover.md

Experiment: ICANN KSK roll

- Do a KSK roll as proposed by ICANN
- Strong desire to minimize packet size
- Timings constrained by DPS
- Re-introduces old (“incumbent”) KSK, so that it can be marked as revoked
  - This limits packet sizes, but...

http://.../root-zone-ksk-rollover-plan-draft-04aug15-en.pdf

Experiments: KSK/ZSK bits

- Change the ZSK to RSA 2048 bits.
  - ”Realistic” scenario, as RSA 1024 bits is reaching end-of-life.
- Change the KSK to RSA 4096 bits.
  - Seems unrealistic. RSA 2048 bits is enough until 2030 (by NIST recommendations).
  - Perhaps focus on algorithm roll instead?

Experiment: ZSK algorithm roll

- Change the ZSK to ECDSA.
- Since ZSK is used on every reply, this is more important
  - Packet size & CPU implications
- Slightly tricky. ;)

Experiment: KSK algorithm roll

- Change the KSK to ECDSA.
- P-256 is equivalent to 3072-bit RSA, so stronger than 2048-bit RSA, while much smaller.
- Also slightly tricky. ;)

Experiment: Frequent ZSK roll

- In principle the ZSK can be rolled much more frequently.
- Lower limit set by RRSIG length and TTL.
- Difficult to map to ICANN processes.
  - Maybe introduce ZSKSK? ;)

Experiment: Multiple ZSK

- Yeti model shares ZSK among DM.
  - Synchronization required.
- Try separate ZSK for each DM.

Experiment: RFC 5011 hold-down

- RFC 5011 has a hold-down timer
  - 30 days
- But... what happens if it is actually needed?
- Introduce a new KSK, then... revoke it?

Experiment: Priming Fragmentation

- Expand the size of the priming answer
  - Already >512 bytes (a “natural experiment”?)
- Measure (apparent) retries, TCP
- Various boundaries:
  - 1280, 1460, 1500, 4096, 9000

Experiment: Priming Truncation

- Measure worst-possible packet case
- Truncate all priming query answers
  - Return TC=1 to all UDP NS queries for root
  - Requires custom software (or possibly a proxy that changes EDNS buffer size on UDP packets)
  - Do we return any data?

Experiment:
We Love TCP

- See what happens in a TCP world
- Truncate ***all*** answers
  - Possibly check for retries and fallback to UDP?
  - Again, do we return any data?

Experiment:
We Love HTTP/TLS/...

- Configure clients to use HTTP/TLS/...

Bii 天地互连

Experiment:
Lots O' Root Servers

- While we hope to get more roots than now, it would be nice to see what happens if we have a ***lot*** of roots.
- We can create this synthetically in two stages:
  1. Make lots of names pointing to current IP addresses (as many as will fit in a DNS message).
  2. Use lots of IPv6 address on a few servers.
- Still synthetic (no route diversity).

Experiment:
Churn Root Servers

- Rapidly update the root servers
  - Remove rand[N,M] root servers each week
  - Add rand[N,M] root servers each week
- Weak form: just add & remove same servers
- Strong form: add & remove unique IP

Bii 天地互连

Experiment:
Root Servers in TLD

- Create a TLD like “.YETI-ROOT” and stick all of the root servers in there.
  - IP addresses do not need to change.
  - Sign it, using normal delegation.

Experiment:
Root Servers without delegation

- Create a *label* like “.YETI-ROOT” and stick all of the root servers in there.
  - IP addresses do not need to change.
- Do *not* delegate from the root.

Plenty Left to do...

- Test QNAME minimization impact?
- ”Minimal responses” from root
- Many that I forgot
- Outcome of research creates ideas
- ...

Bii 天地互连