Yeti DNS:
The First Experiments

Shane Kerr / BII Labs / shane@biigroup.cn

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

Experiments

- Yeti is for research!
- Experimental protocol
  - Lab test
  - List proposal
  - Experiment
  - Report
  - https://github.com/BII-Lab/.../Experiment-Protocol.md
- Queue of experiments
  - https://github.com/BII-Lab/.../Experiment-Schedule.md

Bii
天地互连

Experiment: MZSK (1/2)

- Yeti started with one shared ZSK
- MZSK is ”Multiple ZSK”
  - Separate ZSK for each of three DM
  - 1 KSK, 3 ZSK records in DNSKEY Rrset
  - Each DM signs the zone separately on generation
  - RRSIG per RRset is enough to validate

天地互连

Experiment: MZSK (2/2)

- MZSK does *not* use signing ceremonies
  - KSK is still shared
  - We are testing protocol, not process (mostly)
- Main concern: lots of DNSKEY records
  - In theory can have 8 DNSKEY records!
- Needed to extend DM sync protocol
  - This is a simple method for insuring consistency when generating the Yeti root zone

Experiment: MZSK Phase 1

- Have lots of DNSKEY
  - 1 KSK, 6 ZSK
  - Simulates all ZSK rolling at the same time
- Add 1 ZSK per serial, until 6 total
- Crossed 1280 bytes, saw fragmentation
- UDP failed at one root
  - But not TCP
  - Ended up being Linux kernel bug

Experiment: MZSK IXFR Issues

- Logs reporting issues with IXFR
- IXFR protocol
  - delete RR1, delete RR2, delete RR3, ...
  - add RRa, add RRb, add RRc, ...
- Problems with inconsistent masters
  - DM 1 has different RRSIG than DM 2 or DM 3
  - RRSIG delete fail if slave picks different DM
- BIND & NSD switched to AXFR
  - Knot was leaving old RRSIG (now fixed)

Experiment: MZSK Phase 2

- Actual separate ZSK, one per DM
- For each DM:
  - Add new ZSK, wait 2 days
  - Switch to new ZSK, wait 2 days
  - Remove old ZSK
- Avoid overlap (although not necessary)
  - Takes 4 x 3 = 12 days to roll in new ZSK

Experiment: MZSK Conclusion

- Multiple ZSK works, basically as expected
- https://github.com/BII-Lab/.../Report-MZSK.md
- Future work:
  - Non-shared KSK
  - Maybe zone verification by Yeti root servers?

天地互连

Experiment: BGZSK

- BGZSK is ”Big ZSK”: 2048 bit ZSK
- Moved to top of list by Verisign announcement
- Skipped lab test
  - Lots of people use 2048 bit ZSK
- Rolled new ZSK in over 12 days

Experiment: BGZSK Conclusion

- No surprises (a bit boring, but good!)
- Will be keeping 2048-bit ZSK going forward
- https://github/.../Experiment-BGZSK.md

Experiment: KROLL

- KROLL is ”KSK roll”: KSK roll
- Idea is to test a root KSK roll before ICANN
- KROLL is the first of two experiments:
  - KROLL is normal KSK roll
  - IROLL is like the proposed ICANN roll
- Takes at least 30 days, maybe 60 days ☹

Experiment: KROLL Pre-History

- Did an unplanned KSK roll early in project
  - Default BIND 9 timers, no process review
- Failed due to RFC 5011 hold-down timer
  - Actually, BIND 9 worked fine (no timer?)
  - Unbound broke (as desired?)

Experiment: KROLL Launch

- Bumpy…
- Accidentially made ZSK not KSK (fixed)
- Didn’t publish KSK in documentation
  - Meant that any new resolvers would only have the old KSK
  - Fixed, restarted RFC 5011 timer

Experiment: KROLL RFC 5011 DoS

- Wes Hardaker/Warren Kumari draft

  https://tools.ietf.org/html/draft-hardaker-rfc5011-security-considerations-01
- Published during our roll... ☹
- Which Kees Monshouwer had already pointed out and been overlooked... 😭

Experiment: KROLL RFC 5011 DoS - Explanation

- RFC 5011 has a 30-day hold down timer
- This gets re-set if new KSK not seen
- Attack is a classic replay attack
  - DNS messages can be replayed during signature validity period
  - Causes resolvers to re-start 30 day hold down timer
  - Must add the signature validity period to roll time
- ICANN proposed timings are safe

Experiment: KROLL RFC 5011 DoS – Yeti Response

- KROLL experiment continued on original timeline
  - Yeti resolvers closely monitored, low-value targets
  - Didn’t want to extend experiment again
- Leave timings for next Yeti KSK roll
  - Will perform a targetted DoS against specific Yeti resolvers

Bii 天地互连

Experiment: KROLL BIND 9 Views Problem

- BII resolver modified during KSK roll
  - New view added
- View inherited trust anchor
- View did *NOT* inherit RFC 5011 status
- Suggestions:
  1. Guidance for BIND 9 operators
  2. Modify BIND 9 behavior so views inherit global managed keys

Bii 天地互连

Experiment: KROLL Conclusion

- Formal write-up pending
- RFC 5011 basically works
  - Still some concern over BIND 9 behavior

Pending Experiments

- KSKDOS: KSK Roll with Replay DoS Attack
- RENUM: Root Server Renumbering
- 5011X: RFC 5011 Roll-Back
- FAKER: *Lots* of Root Servers
- DOT-Y: Rename Servers to .YETI-DNS
- PMTNC: Priming Truncation
- ECDSA: KSK ECDSA Roll
- FSTRL: Frequent ZSK Roll
- TCPRT: TCP-only Root

Image Credits

- Science: https://commons.wikimedia.org/...nd.svg