DNS-OARC

NANOG 68, Dallas TX
October 2016

The Best of OARC25

Paul Ebersman
OARC Board

What’s OARC

- The DNS Operations, Analysis, and Research Center
- Five key functions
  - Information Sharing
  - Operational Characterization
  - Workshop
  - Analysis
  - Tools and Services

OARC's Value in Action

- Gathered real-time data during end-2015/ early-2016 attacks
- Post-incident DITL-style upload of attack PCAP from H and K root (thank you !)
- Provided co-ordination resources to root-ops
- Forums for analysis/discussion of what happened, including closed member session at OARC24 workshop

OARC 25

- Just finished yesterday
- Two days of presentations from operators and researchers
  - TLDs, Developers, Academia
- 130 registered attendees
- 68 different organizations

Migrating .CZ to Elliptic Curves

- Going to where no ccTLD has gone before
- Why
  - Smaller responses and zone
  - Testing the algorithm rollover process
- How
  - Measurement experiment
  - Controlled migration of child zones
  - Migrate the parent

Migrating .CZ to Elliptic Curves

- Troubles
  - 6% of users can do RSA but no ECDSA validation (according to APNIC)
  - Communication, communication, communication
  - Measurement widget: IPv6, DNSSEC, Speed
- Finally
  - .CZ will migrate when IANA is ready
  - Public aware enough

Inter-operator transfer of signed TLDs

- Two signed gTLDs operator transfer
- Carefully planned, lots of steps
- RFC 6781
  - Assumes operators can produce slightly different zones
  - Not true when zone is produced from a backend
- A new IETF document will be produced with guidance to operators

Pre-deployment DNS Testing

- DNSviz now added pre-deployment tests
  - For domains not yet delegated
  - By running tests directly on specified addresses
  - Or running a limited instance of DNS server to answer

Anycast Latency: How many sites are enough

- Why Anycast
  - Latency, DDoS defense, collaboration
- Does it work?
  - Ideally divide the Internet into catchment areas
  - But routing is hard
- How it was measured
  - RIPE ATLAS probes against C, F, K, and L-root
  - Optimal possible latency and catchment areas

Anycast Latency

- Results:
  - Median latency generally good
  - Absolute latency nearly optimal
  - Routing policy adds some penalty
  - Location matters, specially to the tail
  - 12 instances provides good latency
    - More helps with the tail, resilience and collaboration

A study of privacy and anonymity in the DNS

- Pitiful DNS privacy
  - Only query content is protected by encryption
- Proposed techniques for privacy
  - Message padding
  - Message interleaving
  - Alter message timing by introducing artificial delay
- Proposed techniques for anonymity
  - Query chaffing

Exploring CVE-2015-7547

- So far attacks are directed to servers, or try to trick clients
- Exploitation of GLIB C bug
- Attack on the client, requires 3 conditions
  - Trigger buffer resize, force partial retry, deliver payload
  - If payload is delivered, smash the stack
- Particularly dangerous for IoT

On the search for resolvers

- Passively detect resolvers’ source address from authoritative DNS data using machine learning
- Motivation: determine which clients could be eligible for whitelisting, or special consideration
- Supervised/unsupervised learning
- Resolvers follow a distinctive traffic pattern and tend to be “sticky”

Rolling the Root Zone DNSSEC KSK

- If you do DNSSEC validation, need to be aware
- Key dates
  - October 27 2016. Generate the new root KSK
  - February 2017: New KSK operationally ready
  - July 11 2017: New KSK added to the root zone
  - October 11 2017, New KSK signs root zone DNSKEY
  - January 11 2018, old KSK revoked
- Rollover will follow RFC 5011
  - If operating correctly, trust anchor will rollover automatically
  - If not manual intervention will be needed

Testing SLD nameservers

- Domains from gTLD zones: 186M
  - Served by 3.5M nameservers
  - Many glue records using questionable addresses
  - Once resolved, nameservers are distributed across 1.5M addresses
  - Around 300,000 unique /24s at least host one nameserver
    - Around 240,000 /24's host at least two
  - Looking at the last octet, .2, .3, .4, .5, .10 and .11 are twice as common as the typical octet
  - ~6.3% of the addresses are open to recursion

The hunger for AAAA

- A ccTLD noted a sudden increase in AAAA for labels without associated AAAA
- It costs money when you are billed by the query
- Using a Big Data platform where able to investigate into the past, detect the main source, produce a report and have it fixed.
- This may affect other operators out there, including ccTLDs

PCAP-TO-HDFS

- CIRA built a real-time platform for DNS traffic analysis for .CA and their DNS anycast service
- Previous solutions not good enough and not scalable
- Hadoop + Flume + Impala + Actor-based concurrency
- Some experiments with Machine Learning

Domain like an Egyptian

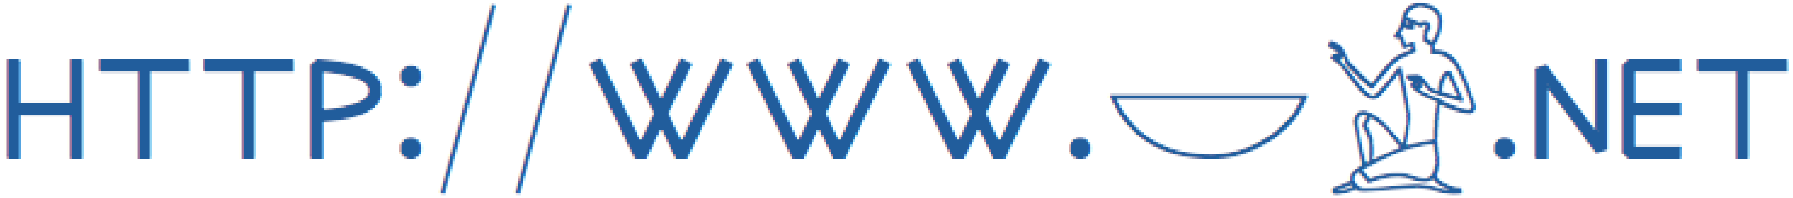

- Verisign IDN supports Egyptian Hieroglyphs
- Register a domain – not all registrars support it
- Set up a nameserver – using punycode
- Client requires font with hieroglyphs
- http://www.xn--5o7dx5d.net/

What to do with SERVFAIL

- On Aug 3, wd2go.com (cloud based storage service) disappeared.
- This causes a lot of stress to resolvers due to SERVFAIL and retries
- Using a lab with different resolver implementation to determine impact
  - Different implementations, different times for SERVFAIL cases
- Client code should have mitigation mechanism for SERVFAIL errors

Yeti DNS: The first experiments

- Yeti DNS: live root DNS server system testbed
- Three experiments so far
  - MZSK: Multiple ZSK
    - Caused troubles with IXFR, needed AXFR
  - BGZSK: Big ZSK (before Verisign announcement)
    - No surprises
  - KROLL: KSK Roll
    - Bumpy due to timers.
    - BIND 9 views problem

Anycast vs DDoS

- DDoS are bad and getting worse
- One attack hit some of the root servers on Nov 30
- During good times anycast keeps traffic contained
- But under attack, what's the best strategy?
  - Keep the site running? Switch to nodes with more capacity? Do nothing?

Anycast vs DDoS

- Summary of the attack
  - 34 GB/s aggregated
  - D, L, M not attacked. A no visible loss, E, F, I, J, K a little bit of loss, B, C, G, H a lot
- RIPE ATLAS used as vantage points to measure loss
  - Site flips from routing changes
  - Collateral damage due to shared sites in some cases

Getting DNSSEC Root TA securely

- Pure Python code
- Minimal external dependencies
- Fetch file, validate signature, generate in right format
- https://github.com/kirei/dnssec-ta-tools

ENT was here!

- In May 2016, the National Security Agency of France reported broken validation for gouv.fr
- All instances of the same anycasted nameservers caused problems: NSD
- .fr zone error free
  - The problem was a combination of ENT name subspace with no signed subzones, NSEC3+Opt-out, BIND for signing, NSD as authoritative
- Still fails with Google DNS
  - Workaround was to introduce a TXT to convert ENT into Non-ENT

When “others” measure the DNS

- When someone is looking for alternatives to host DNS, who to ask?
  - Mailing list? Search Engine? Colleague recommendation?
- Multiple services measuring the DNS
  - Sometimes with dubious methodologies
  - Same providers get wildly different rankings depending on who’s measuring them!
  - Failures accounted differently, saturated path skewing results
- Seek for Guidance

For more information
https://www.dns-oarc.net