The Domain Name System: Where Internet Operations, Research, Security and Policy meet

Keith Mitchell

CWRU EECS Seminar

December 2013

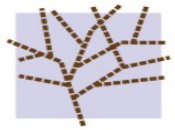

DNS-OARC

Domain Name System Operations Analysis and Research Center

30 Years of DNS

- The first documents defining the Domain Name System were published by Paul Mockapetris as RFCs 882 and 883 in November 1983
- Moved beyond ARPAnet's “hosts.txt” flat name->IP address mapping file
- Distributed, hierarchical, extensible recipe for success !
- I seem to have been messing with it since 1985...

Talk Overview

- Introduction
- The DNS and Internet Abuse
- DNS Data Gathering and Analysis
- Domain Name Public Policy
- Case Study - “Collisions”
- Conclusions
- Q&A, Discussion

DNS-OARC

Domain Name System Operations Analysis and Research Center

Introduction

DNS-OARC

Domain Name System Operations Analysis and Research Center

Speaker's Background

- Internet operations and development since 1986, co-founder of:
  - UK's first commercial ISP, *PIPEX* (CTO)
  - London Internet Exchange, *LINX* (CEO)
  - *.uk* TLD registry, *Nominet UK*
  - *RIPE NCC* Executive Board (Chair)
  - *UK Network Operators' Forum* (Chair)
- Moved to US/Cleveland 2006:
  - Internet Systems Consortium (VP Engineering until 2012)
  - DNS-OARC (President)
  - UKNOF (MD)
  - Open-IX (Board)
  - SMOTI Enterprises (Principal)

DNS-OARC
Domain Name System Operations Analysis and Research Center

Disclaimer

- My background is in network operations and start-ups, my practice is in running critical infrastructure Internet Engineering nonprofits
- I am none of a researcher, security expert, nor programmer – this talk draws extensively on the hard work of others in our community

DNS-OARC
Domain Name System Operations Analysis and Research Center

DNS 101

- *Clients:* your app, desktop, mobile...
- *Resolver servers:*
  - answer queries directly from clients
  - cache answers
  - send queries onto:
- *Authoritative servers*:
  - answer queries for a particular branch of the DNS tree hierarchy (“zone”)
  - answer with referrals to other authoritative servers for queries outside their zone
  - root servers are ultimate authority at apex of namespace
- RFC 1034, 1035 *et al*

DNS-OARC

Domain Name System Operations Analysis and Research Center

The Importance of the DNS

- Modern web sessions typically involve dozens of DNS lookups
- If your providers' DNS resolver fails, you will notice..
- If a top-level authoritative provider fails, **everyone** will notice !

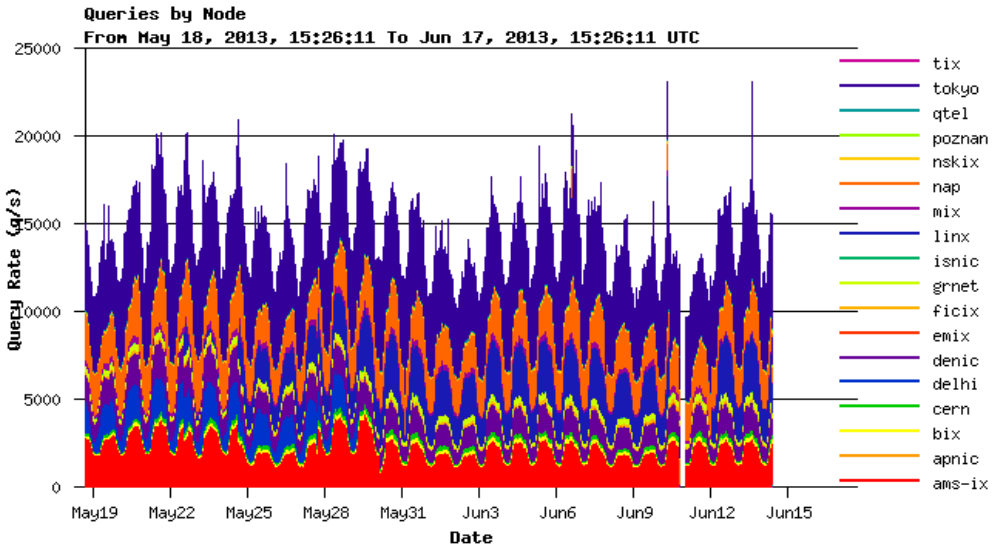

DNS Root Scalability

DNS-OARC

Domain Name System Operations Analysis and Research Center

What is DNS-OARC ?

The Domain Name System Operations Analysis and Research Center (DNS-OARC) is a non-profit, membership organization that seeks to improve the security, stability, and understanding of the Internet's DNS infrastructure.

DNS-OARC's mission is:

- *to build relationships among its community of members and facilitate an environment where information can be shared confidentially*
- *to enable knowledge transfer by organizing workshops*
- *to promote research with operational relevance through data collection and analysis*
- *to increase awareness of the DNS's significance*
- *to offer useful, publicly available tools and services*

DNS-OARC

Domain Name System Operations Analysis and Research Center

OARC Members

Afilias *(.org, .info)*
Google
ICANN
Nominet *(.uk)*
RIPE NCC

AFNIC
Akamai
ARIN
Cisco
DENIC *(.de)*
EurID *(.eu)*
Microsoft
Neustar *(.biz)*
SIDN *(.nl)*

.CLUB
.SE
ARI Registry Services
Artemis *(.secure)*
CentralNic
CIRA *(.ca)*
CloudShield
CNNIC *(.cn)*
CORE
CZ.NIC
DK Hostmaster
Donuts
dotBERLIN
Dyn
eNom
IEDR *(.ie)*
Internet Identity

JAS Advisors
JPRS *(.jp)*
KISA/KRNIC
Mark Monitor
Minds+Machines
NIC Chile *(.cl)*
NIC-Mexico *(.mx)*
Nominum
Norid *(.no)*
NZRS
Registro.BR
RTFM
SWITCH *(.ch)*
tcinet.ru
XYZ

Comcast
ISC
Verisign *(.com)*

AFRINIC
APNIC
CAIDA
Cogent
dotua
LACNIC
McAfee
Measurement Factory
NASA Ames
Netnod
NLnet Labs
NTT
OTTIX
PowerDNS
Public Interest Registry
Secure64
Team Cymru
University of Maryland
USC/ISI
WIDE

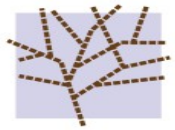

DNS-OARC
Domain Name System Operations Analysis and Research Center

OARC's Functions

- Facilitate co-ordination of DNS operations community
- Ongoing data gathering
- Run twice-yearly workshops
- Operate community info-sharing resources
  - Mailing lists, jabber, website, trust vetting
- Maintain/host DNS software tools
- Outreach via external and shared meetings

DNS-OARC
Domain Name System Operations Analysis and Research Center

The DNS and Internet Abuse

DNS-OARC

Domain Name System Operations Analysis and Research Center

Most DNS Traffic is over UDP

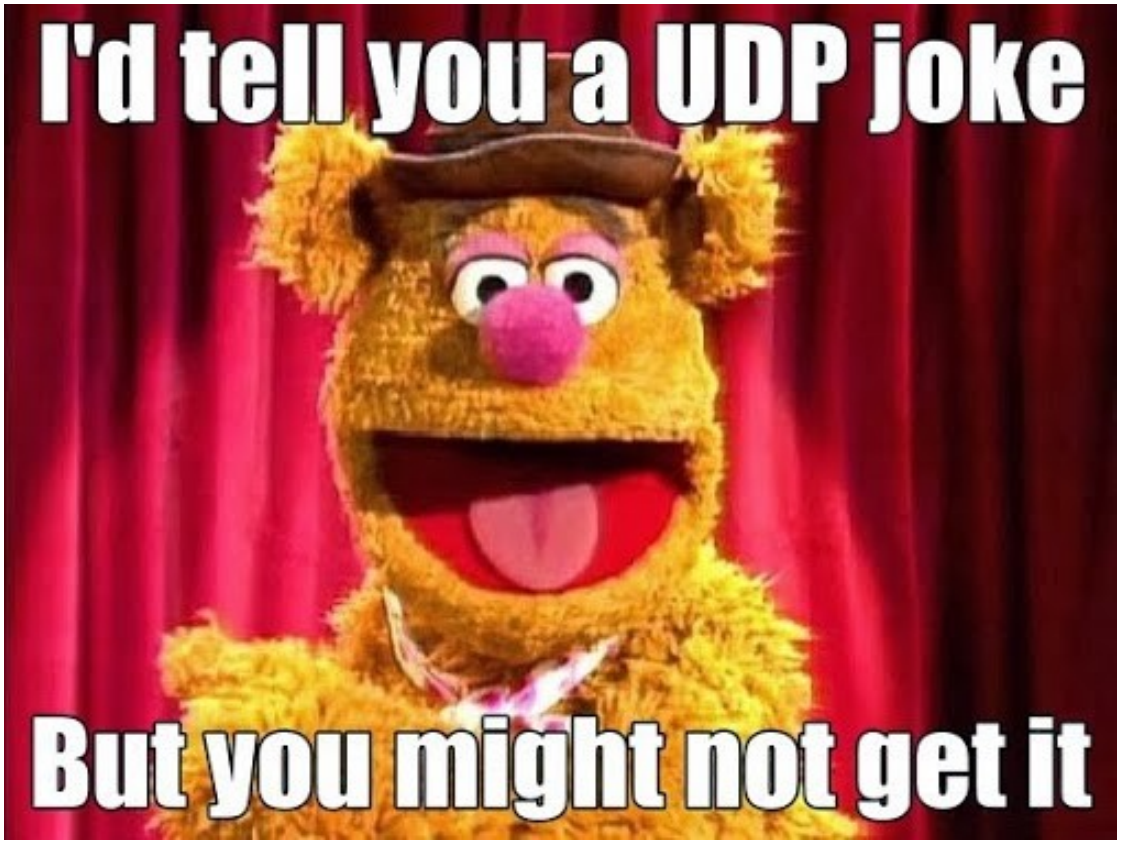

Cache Poisoning

- If a false name->IP mapping is inserted into a server you are using, your traffic can potentially be re-directed to a malicious site
- In theory, there are mechanisms to prevent this:
  - DNS transaction ID
  - application SSL certificates
  - UDP vs TCP
  - DNSSEC
- In practice, the protocol as originally designed has loopholes..

DNS-OARC
Domain Name System Operations Analysis and Research Center

The “Kaminsky” Attack

- In 2008, Dan Kaminsky discovered a new vector for Cache Poisoning attacks against DNS transactions
- Issue (small size of transaction ID) known for years, but new exploit via caching of additional answer records from spoofed responses
- The solution was to increase the entropy used to match up queries/responses by randomizing the UDP source port
- This was a major multi-vendor co-ordinated effort over many months
- It *appears* to have been successful, as cache poisoning attacks in the wild since then, while documented are rare

DNS-OARC

Domain Name System Operations Analysis and Research Center

OARC Web Port Tester

https://www.dns-oarc.net/oarc/services/dnsentropy

66.57.17.110 Source Port Randomness: GREAT

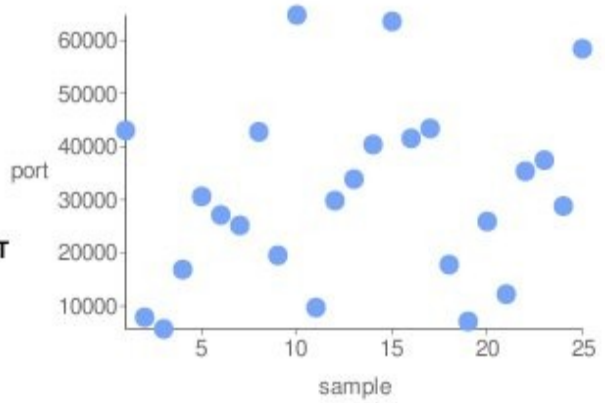

Number of samples: 25
Unique ports: 25
Range: 5691 - 64785
Modified Standard Deviation: 16827
Bits of Randomness: 16
Values Seen: 43054 7891 5691 16897 30628 27141 25182 42783 19549 64785 9724 29847 33894 40400 63576 41563 43425 17799 7114 25924 12237 35382 37464 28826 58425

DNS-OARC
Domain Name System Operations Analysis and Research Center

OARC Web Port Tester

207.217.126.41 Source Port Randomness: POOR

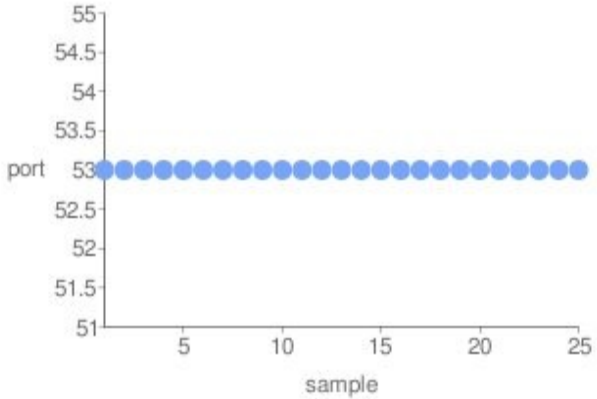

Number of samples: 25
Unique ports: 1
Range: 53 - 53
Modified Standard Deviation: 0
Bits of Randomness: 0
Values Seen: 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53 53

DNS-OARC
Domain Name System Operations Analysis and Research Center

Shulman/Herzberg Attack

- More recent work on variant cache poisoning attack:
  - **https://sites.google.com/site/hayashulman/files/fragmentation-poisoning.pdf**
  - **https://indico.dns-oarc.net/contributionDisplay.py?contribId=18&confId=1**
- DNS packets have grown in length overall since 2008, leading to greater use of EDNS0/UDP fragmentation
- The “Kaminsky” entropy is only in the first datagram fragment
- It thus becomes possible (though tricky) to insert poison records in subsequent fragments

Data from OARC Port Test Tools

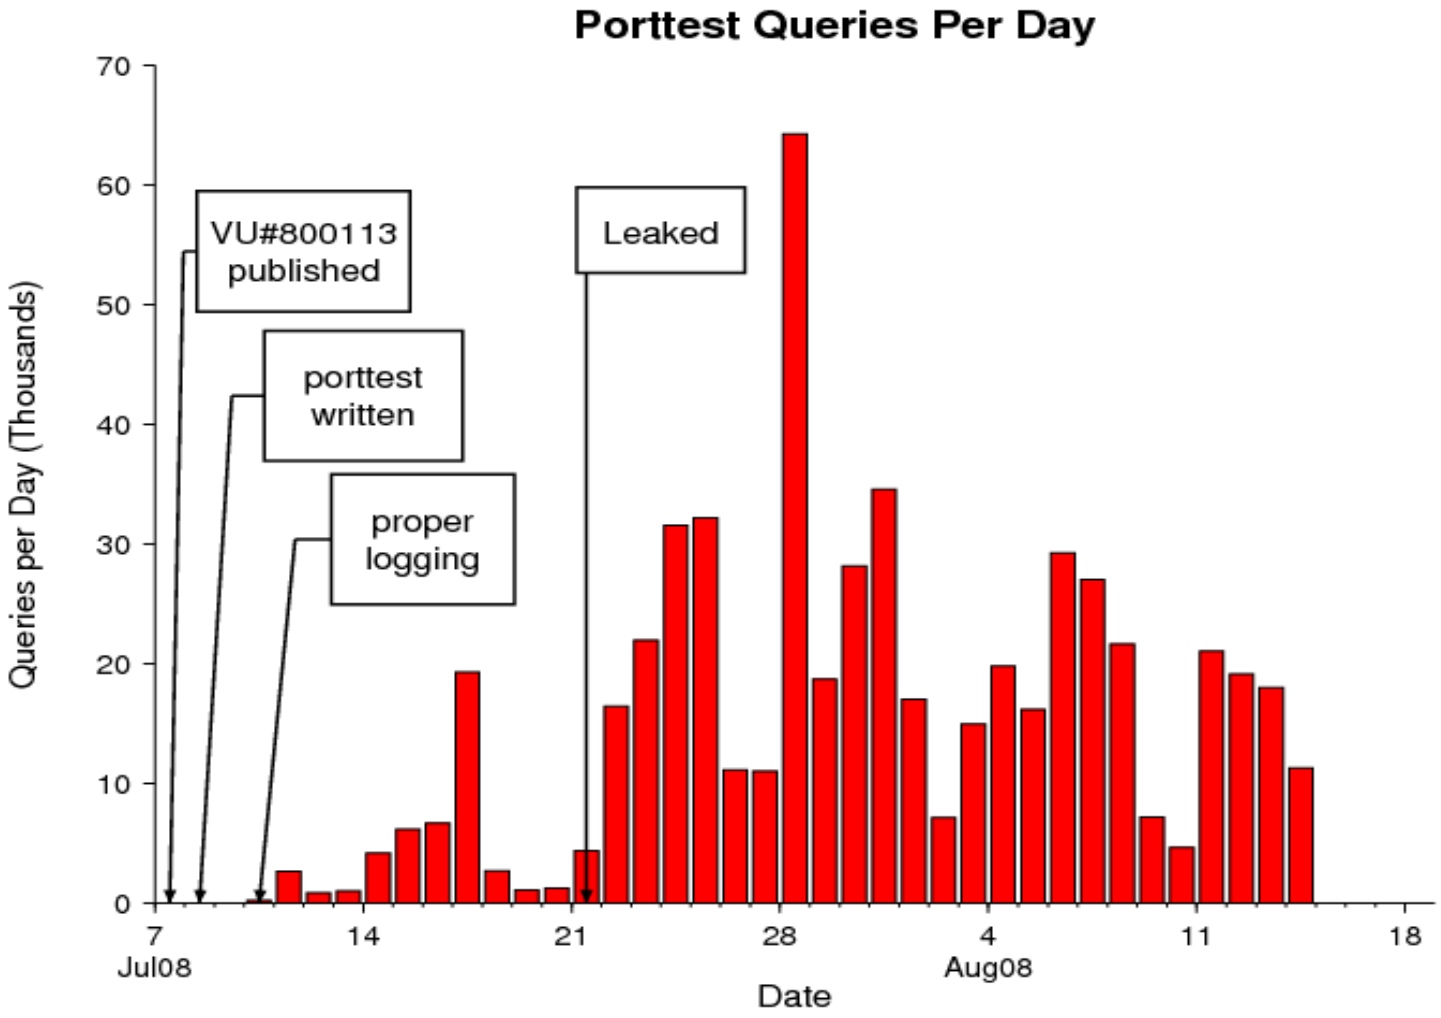

Data from OARC Port Test Tools

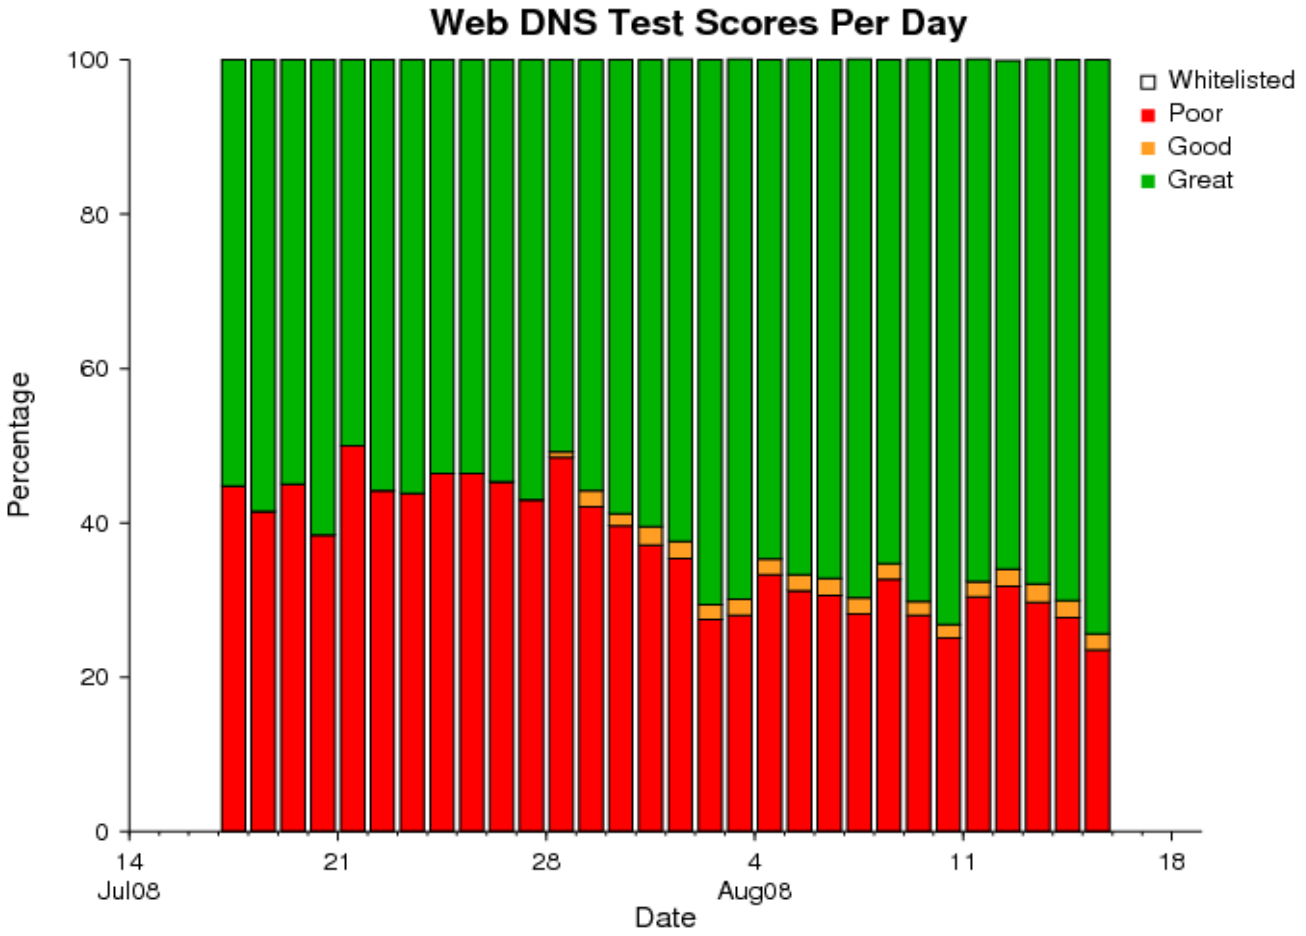

DNS-OARC

Domain Name System Operations Analysis and Research Center

Amplification Attacks

- Botnets are commonly used for Distributed Denial of Service (DDoS) attacks by bad actors
- One way attacks can have much more impact is through *amplification*
- Send a small packet to a 3rd party with a spoofed source address, which triggers a much larger packet back to the victim
- Some DNS queries (including DNSSEC, and ANY), generate a *much* larger response than query
- Not just DNS: *SNMP*, *NTP*, *Chargen/19* are all UDP-based protocols which can act as amplifying reflectors if server ports not properly restricted

DNS-OARC
Domain Name System Operations Analysis and Research Center

SpamHaus/StopHaus Attack

- March 20th 2013:

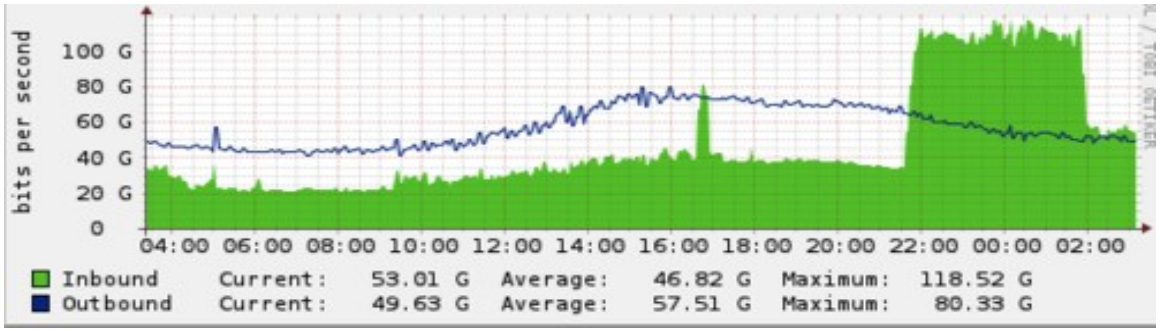

- At over 75Gb/s, this is one of the biggest ever documented DDoS attacks seen on the Internet:
- http://blog.cloudflare.com/the-ddos-that-knocked-spamhaus-offline-and-ho
- This was realized through DNS amplification..

IP Address Spoofing

- This is possible because more than 20% of Internet providers don't do source address verification (BCP38), making spoofing of source (victim) IP addresses trivial

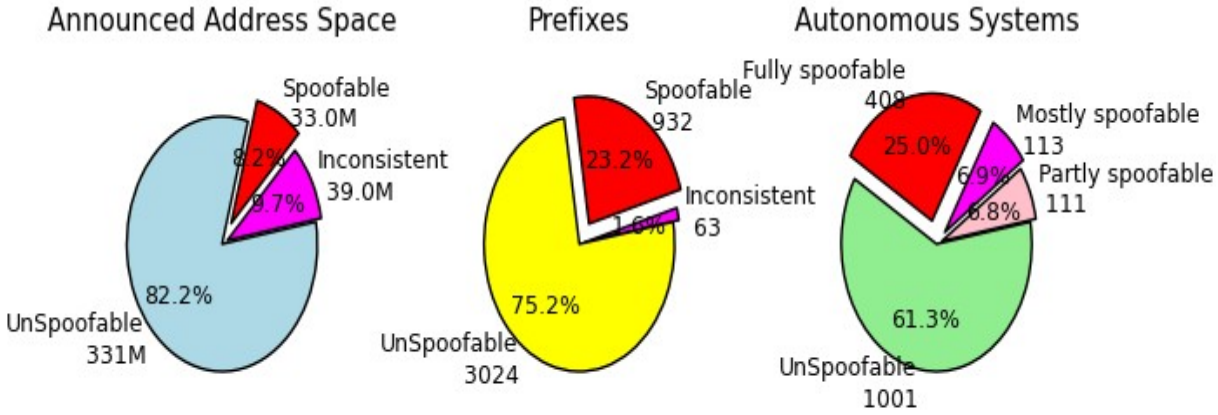

- Source: *Spoofer Project*: http://spoofer.cmand.org

Open Resolvers

- There are some 30m DNS resolvers which are mis-configured to openly respond to queries from anywhere
- Source:
  http://www.openresolverproject.org

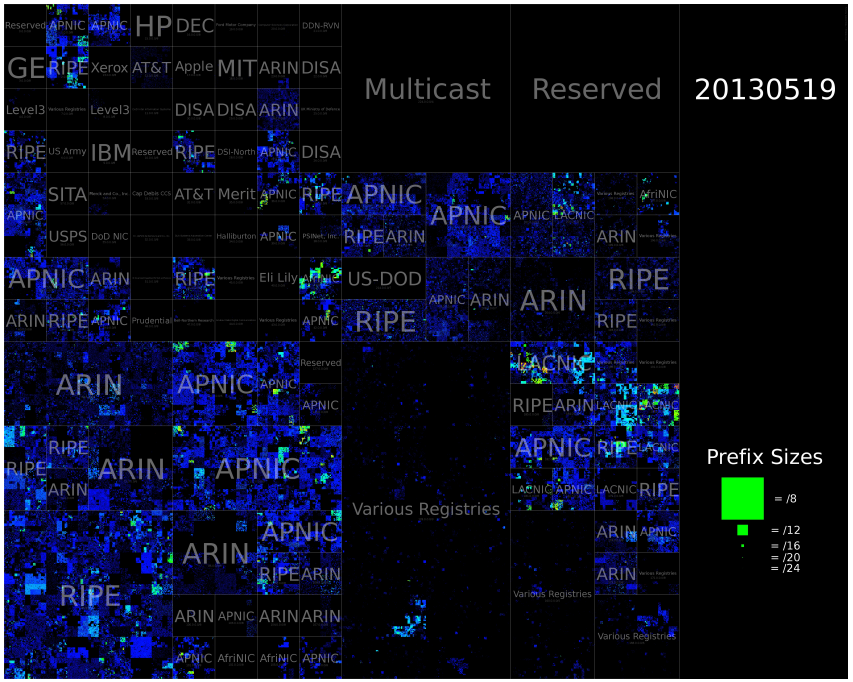

Addressing the Problem

- The work of researchers and operators doing projects like Spoofer and OpenResolver is invaluable to detecting, measuring and understanding these problems
- There is no substitute for gathering live data from the Internet
- While no panacea, the DNS is pervasive enough its use for data gathering can make it part of the solution, not just the problem..
- Solving these problems to stop the abuse is a long-haul, education based on sound data and analysis is vital to these efforts

DNS Data Gathering and Analysis

DNS-OARC

Domain Name System Operations Analysis and Research Center

DNS Data Gathering

- Generally involves sensors running on, or adjacent to servers, e.g.
  - Domain Statistics Collector (DSC) – continuous traffic analysis and summary, no payload
  - “Day in the Life of the Internet” (DITL) - full query payload for 48 hours at least once a year
  - Capturing data from user-driven test tools
  - “Passive DNS” capture of resolver->authoritative server traffic

OARC's DITL Dataset

- Since 2006, at least once per year to provide “Internet Science” baseline
- Also during key DNS events such as DNSSEC signing of root, IPv6 enabling, potentially during incidents
- Gathered from most Root and many Top-Level Domain (TLD) operators
- Full query traffic to authoritative servers
- 80Tb dataset
  - OARC has been doing “big data” for nearly a decade..
  - less challenging with modern hardware than when we first did this !
  - https://www.dns-oarc.net/oarc/data/ditl

DNS-OARC
Domain Name System Operations Analysis and Research Center

DITL in Action

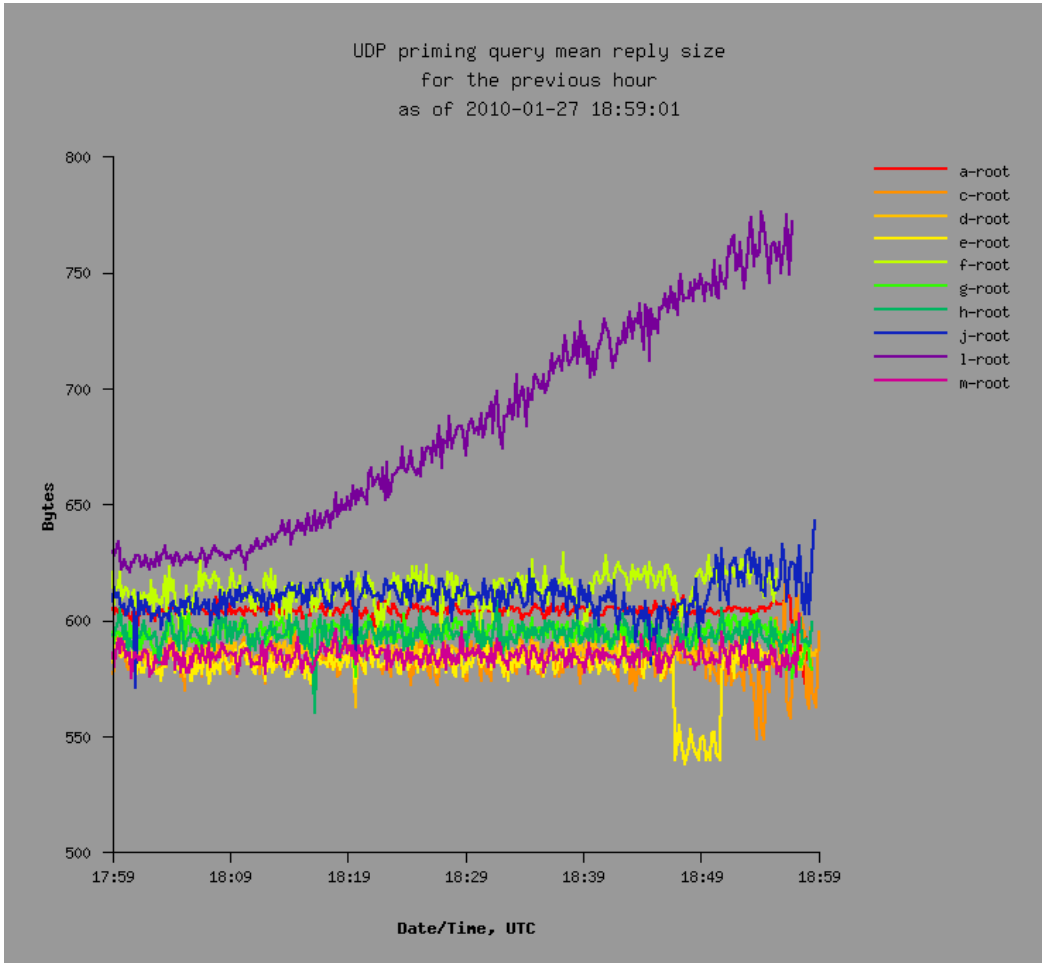

DNS-OARC

Domain Name System Operations Analysis and Research Center

The Case for DNSSEC

DNS-OARC

Domain Name System Operations Analysis and Research Center

Current DNS fixes are Interim

- Source port randomization shifts burden of protecting one application onto the operating system platform
- Increasing bandwidth and CPU power are eating away at extra entropy
- As Shulman/Herzberg have demonstrated, there's always scope for new variants on old attacks
- Switching all DNS transactions from UDP to TCP has other issues
- Nobody thought pervasive State censorship and surveillance was even a possibility when the DNS was designed ☹

DNS-OARC
Domain Name System Operations Analysis and Research Center

Reasons to do DNSSEC

- Standards and implementations are now mature
- Effective defense against cache poisoning !
- Great anti-phishing measure
- Interferes with commercial violation of Internet end-to-end principle
  - “NXDOMAIN Redirection”
  - Netalyzr will tell you if your provider is tampering
- General infrastructure integrity enhancement
- DANE could even replace SSL certs one day..

DNS-OARC
Domain Name System Operations Analysis and Research Center

Understanding DNSSEC

- Allows for cryptographic verification that DNS records are authentic
- DNSSEC enabled authoritative servers provide digital signatures in addition to “standard” DNS data
- DNSSEC validating resolvers provide authenticated responses with proven integrity
- Clients using validating resolvers get guaranteed “good” data
- Data that does not validate provides a `“SERVFAIL”` response

DNS-OARC
Domain Name System Operations Analysis and Research Center

Network Impact of DNSSEC

- Signed DNS responses are BIG - 512 byte UDP packets just don't cut it
  - Need to use EDNS0 - RFC 2671: "Extension Mechanisms for DNS"
  - Allows for bigger DNS messages via IP Fragments
- Network elements non-transparent to EDNS0 or large MTU UDP 53 may degrade DNS queries
- Testing tools:
  - https://www.dns-oarc.net/oarc/services/replysizetest
  - https://netalyzr.icsi.berkeley.edu

DNS-OARC
Domain Name System Operations Analysis and Research Center

Obstacles to DNSSEC

- Registrar support variable
- Hard to understand/configure
- Easy to break
- Difficult to use admin tools
  - getting better, e.g. BIND9.9
- Firewall and CPE equipment issues
- Education and experience-sharing can fix these

DNS-OARC
Domain Name System Operations Analysis and Research Center

Domain Name Public Policy

DNS-OARC

Domain Name System Operations Analysis and Research Center

Internet Governance Primer

- The Internet does not hold together without effort
- Balance of competition and co-operation
- Some functions are too important to be trusted to corporations or governments !
- “Bottom-up self-organizing multi-stakeholder” model
- Often embodied by mutual nonprofit organizations

DNS-OARC
Domain Name System Operations Analysis and Research Center

Internet Governance in Practice

- Standards: *IETF, W3C, IEEE, ...*
- IP addresses:
  IANA, ARIN, RIPE NCC, LACNIC, APNIC, AfriNIC
- Operations: *NANOG, RIPE, APRICOT, UKNOF, ...*
- Domain Names: *ICANN, PIR, CENTR, …*
- Policy: *ISoc, EFF, EuroISPA, ...*
- Internet Exchanges: *Euro-IX, Open-IX, ...*

DNS Governance

- ICANN is often misunderstood as “controlling the Internet”, but its remit is strictly only names and numbers
- Works with registries, registrars, ccTLDs, gTLDs, governments, root operators
- In past years, has approved 100s of new Top-Level Domains to be created (e.g. recently):
  - .bike, .guru, .xxx, .**САЙТ**, . 游戏

Evidence-Informed Policy

- Decisions to make changes at the top level of the DNS are ultimately commercial/political ones
- Many vested high-stakes commercial interests involved..
- ..but cannot be made in an operational vacuum
- Could there be adverse security/stability impacts ?
- How best to inform policy makers with hard evidence ?

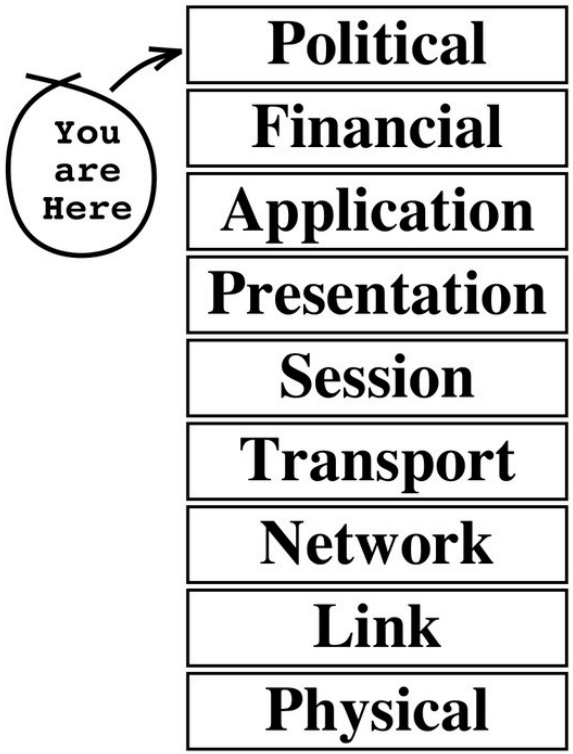

Case Study: “High-Risk Strings Collisions”

DNS-OARC

Domain Name System Operations Analysis and Research Center

DNS Security Collides with Policy

- ICANN approves new TLDs on a competitive bidding process
- Various domains such as “.corp”, “.home” applied for in process
- Unfortunately various entities already make non-standard use of “pseudo TLDs” in their **internal** networks
  - some of these are same as new TLDs being applied for
  - worse, some of these have “internal-use-only” SSL website-security certificates already issued for them !
- Could creating these domains on the wider Internet “collide” with their internal usage ?
- Worse, could it lead to website impersonation and hi-jacking ??

DNS-OARC
Domain Name System Operations Analysis and Research Center

OARC's Data-set to the Rescue

- Rather than debate/litigate endlessly, it's possible to analyze data already gathered to decide the extent of queries for potential new TLDs on the live Internet
- OARC's DITL dataset from 2006-2013 available for this:
  - not the perfect resource for such research, but much better than nothing at all
  - triggered donations of some extra CPU-power ☺

DNS-OARC
Domain Name System Operations Analysis and Research Center

ICANN Collisions
DITL Query Analysis

- https://www.icann.org/en/about/staff/security/ssr/name-collision-02aug13-en.pdf

Rank	Proposed TLD	As TLD	As SLD	At all other levels	Total
1	home	595,024	24,117	3,723	622,865
2	corp	122,794	31,084	39,985	193,864
3	site	13,013	212	412	13,637
4	global	10,838	8,895	13,838	33,571

- Summary:
  - **Not safe** to delegate "*.corp*" or "*.home*" new TLDs
  - Mostly safe to delegate 80% of rest
  - 20% need further study, safeguards

Conclusions

DNS-OARC

Domain Name System Operations Analysis and Research Center

Conclusions

- In a world of mobile apps and search engines, the DNS may be much less visible to end-users than it was 30 years ago
- But it still underpins the Internet in critical ways
- Yet another invisible layer in the protocol stack
- A unique place to measure and tinker

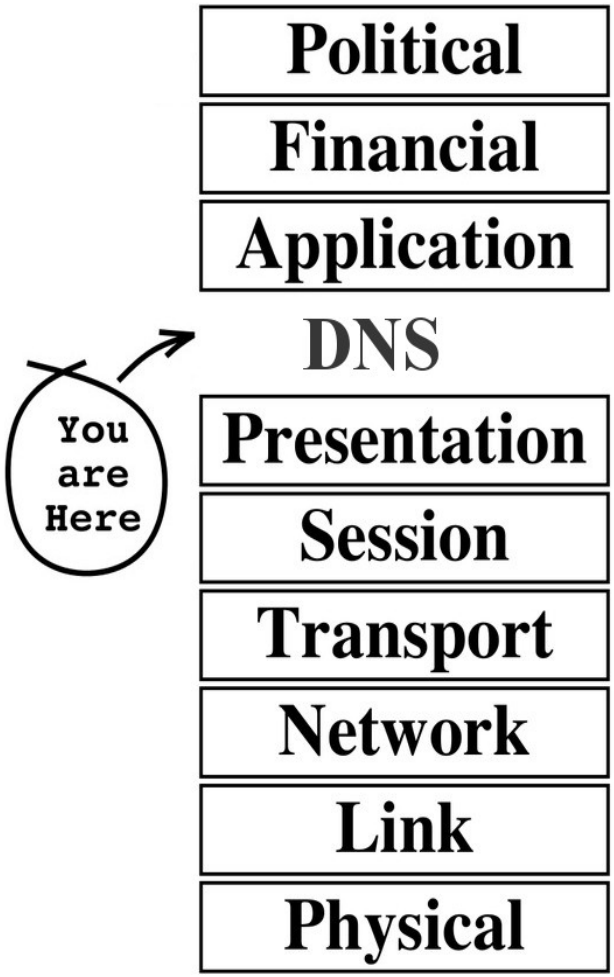

Conclusions

- There is no substitute for gathering live data from the Internet
  - this can be done whilst still respecting privacy
- The DNS is pervasive enough its use for data gathering can make it part of the solution, not just the problem
- Operators have live data network data, but don't always have the skills/insight/time to analyze it
- Researchers can greatly help understand this data, but don't always find it easy to obtain, or to interpret operational impact
- Working together we can answer important protocol, implementation, security and policy questions

DNS-OARC
Domain Name System Operations Analysis and Research Center

Further Information

- Web: https://www.dns-oarc.net
- Workshops: https://indico.dns-oarc.net
- E-mail: keith@dns-oarc.net
  dns-operations@lists.dns-oarc.net
- Social: https://www.linkedin.com/groups/DNSOARC-3193714
- IM: xmpp:keith@jabber.dns-oarc.net
- Phone: +1 650 423 1348 (EST)

DNS-OARC
Domain Name System Operations Analysis and Research Center

Questions ?

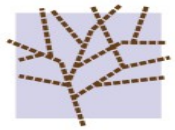

DNS-OARC

Domain Name System Operations Analysis and Research Center