

DNS: The Protocols, The Myths, The Legends

Paul Ebersman - Neustar
paul.ebersman@team.neustar
NANOG 78 – SF 10 Feb 2019

DNS Classic

BACK IN THE DARK AGES

- 1-2 Bare metal servers as auth NS
- 56k uplinks
- CPU/RAM/Disk all expensive
- We all knew each other

BACK IN THE DARK AGES 2

- You could read all the DNS RFCs in a weekend... (now over 185 RFCs, 2800 pages...)
- Everything was unicast and UDP
- Folks w/security checklists didn't know or talk to DNS folks

DNS & TCP

“CONVENTIONAL” WISDOM

- DNS was UDP port 53
- TCP was only needed for zone transfers and could be locked down to just the listed auth servers
- This Best Practices security audit checklist is flawless
- The earth is flat.

AND FOR THE SECURITY CHECKLIST FOLKS BLOCKING TCP...



Free “Best Practice” Security Checklist In Every Box!

REALITY

- TCP has always been needed for sending large packets (> 512 bytes), either in initial query/response or when TC (truncate) bit set in truncated DNS response
- There are good reasons for hosts other than those listed to do AXFR/IXFR

AND THE NEW REALITY

- EDNS0, DNSSEC, overuse of TXT records and all sorts of other things create large packets.
- IPv6 UDP PMTUD problematic (more in IPv6 section)
- TCP for DoT/DoH, pipelining

AND THE NEW REALITY

- Can load balance/shard w/TCP
- Stateful DNS, [RFC 8490](#)

IPv6 and DNS

DNS OVER IPV6 ISSUES

- PMTUD (Packet too big)
- UDP fragments dropped
<https://blog.apnic.net/2017/08/22/dealing-ipv6-fragmentation-dns/>,
<https://blog.apnic.net/2017/08/29/dealing-ipv6-fragmentation-dns-part-2/>
- Large numbers of clients don't retry on TC bit set

DNSSEC Basics

DNSSEC BASICS

- Public-key/asymmetric encryption
- Private keys kept secret/secure
- Zone data and delegations digitally signed w/private key
- Public keys published in the DNS
- DNS query results validated using public key
- Validation failure results in SERVFAIL instead of answer

DNSSEC

WHAT “EVERYONE” SAYS

- It's fragile/complicated
- The signing software is “hard” to use
- Will drive up support costs dramatically
- No benefit for extra risks

IT'S FRAGILE/COMPLICATED

- BGP isn't? Web servers aren't? 😊
- Server software vastly more mature in last 3-5 years, much easier to use (other than DS mgmt)
- Lots more large scale operational experience, both signing and validating

TOO EXPENSIVE TO SUPPORT

- Google/Comcast/Quad9 and other large resolver farms do trillions of queries a day.
- DNSSEC validation incidents are on order of dozens per month
- This percentage of errors has to be in scientific notation, it's so small

WHY DNSSEC

- Cache poisoning
- Additional protection from domain hijacking
- DANE for email/certs
- Protect CAA records
- What other scalable PKI have we done (other than kerberos/AD)

**What does
DNSSEC solve?**

BASIC SECURITY CONCEPTS

- Confidentiality
- Integrity
- Availability

WHAT DNSSEC DOES SOLVE

- Integrity
 - Cache poisoning
 - False authoritative servers

**What doesn't
DNSSEC solve?**

WHAT DNSSEC DOESN'T SOLVE

- Confidentiality
- Availability
- Correct DNS data
- Parent zone security

New Encrypted Transports (DoT/DoH)

POST-SNOWDON ERA

- RFC 7624:

- In the face of pervasive monitoring, we should encrypt anything we can encrypt.

ENCRYPTED TRANSPORT

- DoT (DNS over TLS): [RFC 7858](#)
 - For stub resolver to recursive resolver, encrypts all queries/responses using TLS (ADoT, recursive to auth DoT proposed but not yet standardized)
- DoH (DNS over HTTPS): [RFC 8484](#)
 - For application (like browser) to recursive resolver, includes all queries/responses in-band in HTTPS session

WHAT DOES THIS SOLVE

- Confidentiality

WHAT DOESN'T THIS SOLVE

- Integrity
- Availability

WHAT ARE VENDORS DOING

- **Mozilla:** <https://blog.mozilla.org/futurereleases/2019/09/06/whats-next-in-making-dns-over-https-the-default/>
 - opt-out, not opt-in...
 - canary domain for enterprises (use-application-dns.net)
 - uses cloudflare 1.1.1.1 by default as DoH server
 - bypasses OS stub resolver, enterprise/ISP resolver, sends query to US company

WHAT ARE VENDORS DOING

- **Google:** <https://blog.chromium.org/2019/09/experimenting-with-same-provider-dns.html>
 - opt-in for now, has backed off opt-out by default
 - uses currently configured resolvers of OS, checks for DoH, then DoT, then does in the clear

WHAT ARE VENDORS DOING

- **Microsoft:** <https://techcommunity.microsoft.com/t5/networking-blog/windows-will-improve-user-privacy-with-dns-over-https/ba-p/1014229>
 - opportunist use of DoH if configured resolvers support it
 - done in system stub resolver, so all apps/browsers will use DoH (or not)

WHAT SHOULD ENTERPRISE/ISP DO

- Set up canary domain if you don't want mozilla/cloudflare getting your queries
- Set up your own DoT/DoH on the same IPs you have your current resolvers on.

Q & A

Thanks!

Further Reading

RELEVANT IETF WORKING GROUPS/EMAIL LISTS

- [DNSOP](#): DNS operations
- [DPRIVE](#): DNS privacy
- [ADD](#): Applications Doing DNS proposed WG
- [ABCD](#): Application Behavior Considering DNS
- [EDDI](#): Encrypted DNS website/mailing list

FURTHER READING

- <https://datatracker.ietf.org/doc/draft-ietf-dnsop-dns-tcp-requirements/>
- [RFC 7766](#): DNS Transport over TCP - Implementation Requirements
- [RFC 8490](#): DNS Stateful Operations