



T-DNS: Connection Oriented DNS to Improve Privacy and Security


John Heidemann¹
 joint work with Liang Zhu¹, Zi Hu¹,
 Duane Wessels², Allison Mankin², Nikita Somaiya¹
¹ USC/ISI, ² Verisign Labs

10 May 2014

Copyright © 2014 by John Heidemann
 Release terms: CC-BY-NC 4.0 International

don't fear connections for DNS




DNS Basics

since 1987 (RFC-1034)
 DNS is simple request-response:

client: A www.example.com ?

server: 192.0.2.1

perfect for UDP
 (TCP supported too, but as fallback and zone transfers)




T-DNS / DNS-OARC

Fear of DNS over TCP

- TCP is horribly slow: *bad client latency*
- TCP => server state : *server memory explodes*

community consensus: ~~orthodoxy~~ *dogma*
 don't use TCP*, UDP's constraints are OK

* except for fallback and zone transfers




T-DNS / DNS-OARC

Our Contributions

- analysis: **don't fear connections for DNS**
 - client latency: only modestly more
 - server memory: well within current hardware
- implementation choices to get here
- small protocol addition: TLS upgrade


=> *T-DNS: DNS over TCP+TLS*



T-DNS / DNS-OARC

T-DNS: TCP and TLS Connections

- introduction
- **why**
- how
- at minimal cost
- better than alternatives
- next steps



T-DNS / DNS-OARC


Why T-DNS

- protecting privacy
 - connections -> encryption -> privacy
- denying DoS (Denial of Service)
 - connections -> spoof-proof -> no amplification attacks
- leaving limits
 - connections -> UDP limits don't drive policies

T-DNS / DNS-OARC

Protecting Privacy

- principle: *all* traffic should be private (=> encrypted)
 - Google Public DNS, OpenDNS, others
 - traffic over WAN should be private!
- individuals avoiding transparent proxies
 - multiple ISPs intercept DNS to add ads
- DNS is more than addresses
 - anti-spam (DNSBL), embedded user IDs (facebook, etc.)
 - ex: DNSBL's spam check sends IP address of every incoming mail server over the WAN
 - even on LAN (where destinations are visible), should protect other content



advocacy of Google public DNS to avoid Turkish censorship of Twitter, 2014-03-21

T-DNS / DNS-OARC

Denying DoS

- problem: DNS attacks others
 - DNS amplification attacks
 - a growth industry in 2013: >100Gb/s attacks
- problem: DoS on DNS servers
 - work-around: massive over-capacity
- solution: TCP
 - well understood anti-DoS methods:
 - 3-way handshake precludes spoofing
 - TCP cookies shift state to client for non-spoofed

an amplification attack:
 attacker, **forging IP** of victim
 Q: ANY for **example.com**?
 (~60 bytes)

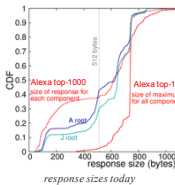
*server: let me help you,
here's 4000 bytes
of what I know about
example.com*

result: 60x more bits on victim

T-DNS / DNS-OARC

Leaving Limits

- for >25 years, *policy* decisions forced by UDP packet sizes
 - number of root servers: all fit in 512B
 - DNSsec: required EDNS for >512B
 - crypto algs and key sizes: pkt size limited
 - key rollover: temporary 2x size
- partial fix: EDNS0 deployment (10+ years, since 1999)
- what uses already discarded as too big?
 => **enough already!**



T-DNS / DNS-OARC

Doesn't DNSsec already "Secure DNS"?

A: yes, but...

- DNSsec is about *query integrity*
 - that is: if you are told X, is X true?
 - it signs answers; signatures prove X is true
- DNSsec does *nothing* for problems
 - *everything* sent in the clear: *no privacy*
 - nothing about DoS
 - large signatures stress UDP size limits

=> need DNSsec's integrity *and* T-DNS' privacy

T-DNS / DNS-OARC


T-DNS: TCP and TLS Connections

- introduction
- why
- **how**
- at minimal cost
- better than alternatives
- next steps

T-DNS / DNS-OARC

(Review) Our Contributions

- analysis: **don't fear connections for DNS**
 - client latency: only modestly more
 - server memory: well within current hardware
- implementation choices to get here
- small protocol addition: TLS upgrade




T-DNS / DNS-OARC

(Review) Our Contributions

- analysis: **don't fear connections for DNS**
 - client latency: only modestly more
 - server memory: well within current hardware
- implementation choices to get here
- 1. small protocol addition: TLS upgrade**


(going in reverse order)



T-DNS / DNS-OARC

Protocol Changes: Goals


- minimize change
- reuse existing approaches *(as boring as possible)*
- follow IETF patterns



T-DNS / DNS-OARC

Protocol Changes: Goals

- minimize change
- reuse existing approaches *(as boring as possible)*
- follow IETF patterns
- implications:**
 - reuse TLS: Transport Layer Security
 - add a STARTTLS-like "upgrade"
 - innovation: careful implementation



T-DNS / DNS-OARC

SMTP before STARTTLS

C & S: open TCP connection
S: 220 mail.imc.org SMTP service ready


C: EHLO mail.example.com
S: 250-mail.imc.org hi, extensions are: -8BITMIME -STARTTLS DSN

C: MAIL FROM:<sender@mail.example.com>
S: 250 2.1.0 <sender@mail.example.com>... Sender OK

C: RCPT TO:<destination@mail.example.com>
S: 250 2.1.5 <destination@mail.example.com>

C: <send mail contents>

problem: cleartext mail is snoop-able (fix: TLS)



T-DNS / DNS-OARC

SMTP with STARTTLS

C & S: open TCP connection
S: 220 mail.imc.org SMTP service ready

C: EHLO mail.example.com
S: 250-mail.imc.org hi, extensions are: -8BITMIME -STARTTLS DSN

C: STARTTLS
S: 220 Go ahead

C & S: <negotiate a TLS session with a new session key, in binary>

C: EHLO mail.example.com
S: 250-mail.imc.org hello, extensions are: -8BITMIME DSN

C: MAIL FROM:<sender@mail.example.com>
S: 250 2.1.0 <sender@mail.example.com>... Sender OK

C: RCPT TO:<destination@mail.example.com>
S: 250 2.1.5 <destination@mail.example.com>


C: <send mail contents>

prologue: in clear (no privacy here)

transition to TLS

contents now private

this example: SMTP; idea used for IMAP, POP3, FTP, XMPP, LDAP, NNTP...



T-DNS / DNS-OARC

Our STARTTLS for DNS

(in draft-hzhwm-start-tls-for-dns-01)

C & S: open TCP connection *prologue*

transition to TLS

C: QNAME="STARTTLS", QCLASS=CH, QTYPE=TXT with the new TO bit set in EDNS options
S: RCODE=0, TXT="STARTTLS", with the TO bit set

C & S: <negotiate a TLS session, get new session key, in binary> *contents now private*

C: <send actual query>
S: <reply to actual query>

pros: no new port (from IANA, or in firewalls)
 cons: **extra RTT**; middleboxes may not like encrypted tlc

USC Viterbi School of Engineering T-DNS / DNS-OARC

(Review) Our Contributions

3. analysis: **don't fear the DNS connection**
 - client latency: only modestly more
 - server memory: well within current hardware
2. **implementation choices to get here**
1. small protocol addition: TLS upgrade

(going in reverse order)

USC Viterbi School of Engineering T-DNS / DNS-OARC

Careful Implementation Choices

- problem: no tuning of DNS TCP for queries *(until now!)*
- connection reuse (or restart)
 - persistent connections
 - TCP fast open
 - TLS resumption
- query pipelining
- out-of-order processing

USC Viterbi School of Engineering T-DNS / DNS-OARC

Latency in DNS/TLS

C & S: open TCP connection *TCP 3wh: +1 RTT*

STARTTLS: +1 RTT

C: QNAME="STARTTLS", QCLASS=CH, QTYPE=TXT with the new TO bit set in EDNS options
S: RCODE=0, TXT="STARTTLS" with the TO bit set

C & S: <negotiate a TLS session with a new session key, in binary> *TLS handshake: +2 RTTs*

C: <send actual query> *query: 1 RTT*
S: <reply to actual query>

USC Viterbi School of Engineering T-DNS / DNS-OARC

Connection Reuse

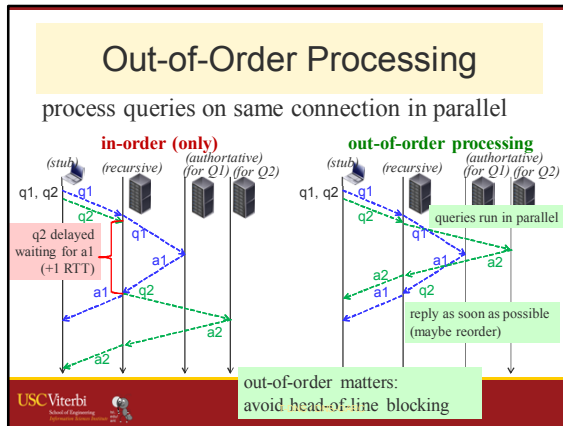
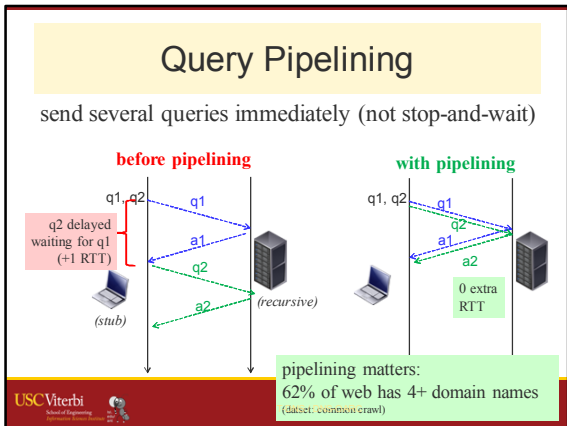
- basic idea:
reuse connection -> no setup cost
- secondary idea:
if must close, client keeps state to restart quickly

USC Viterbi School of Engineering T-DNS / DNS-OARC

Connection Reuse

- basic idea:
reuse connection -> no setup cost
 - persistent connections (in client and server)
- secondary idea:
if must close, client keeps state to restart quickly
 - TCP fast open: client has cookie to send data in 3wh
 - draft-ietf-tcpm-fastopen-08: in Linux-3.6, default 3.13
 - TLS resumption (RFC-5077): client keeps
 - RFC-5077: in OpenSSL and GnuTLS

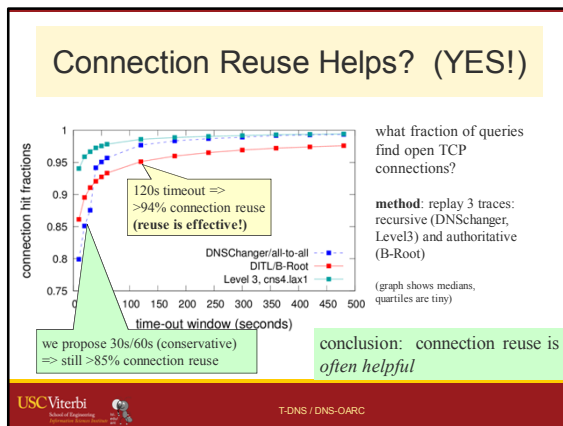
USC Viterbi School of Engineering T-DNS / DNS-OARC

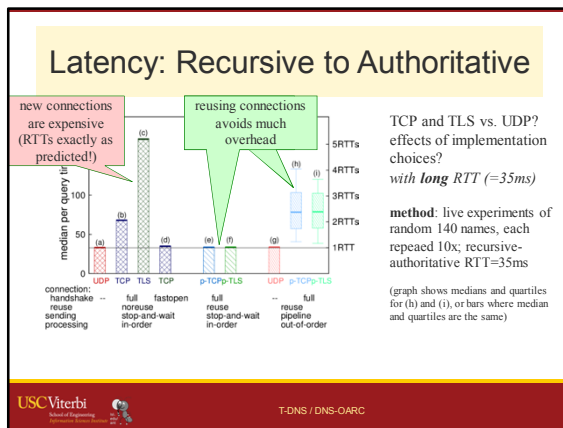
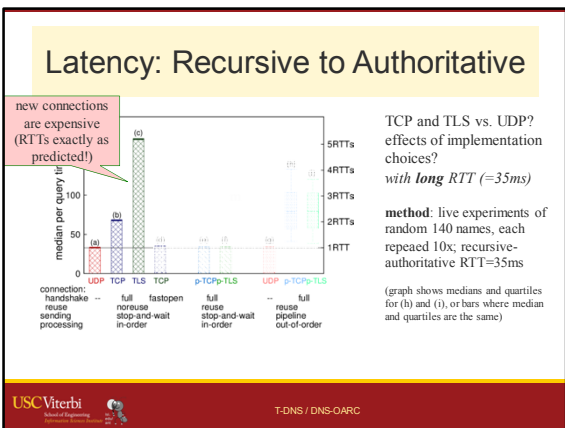
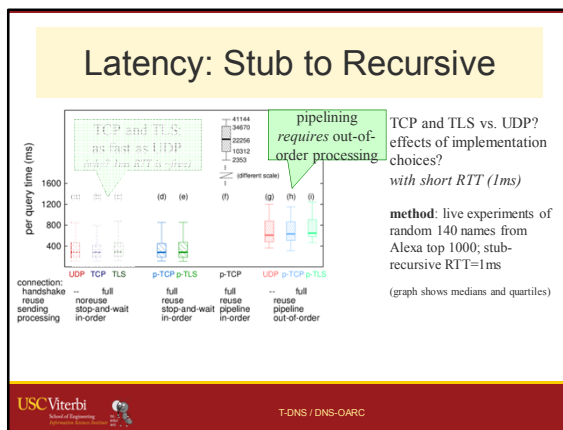
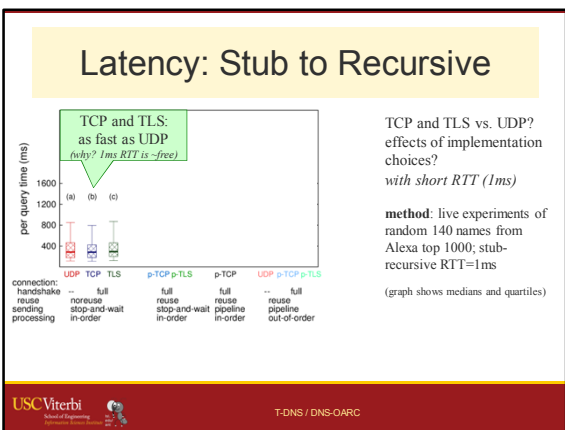
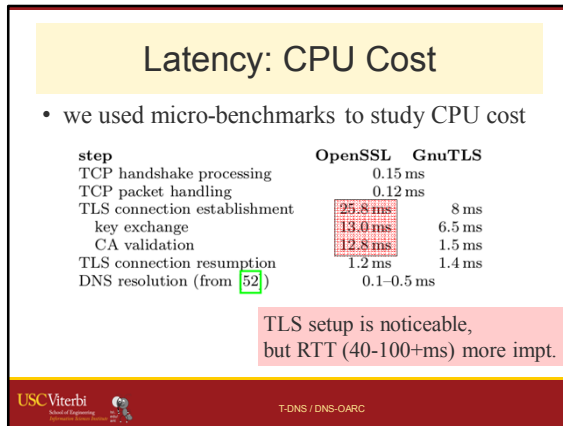
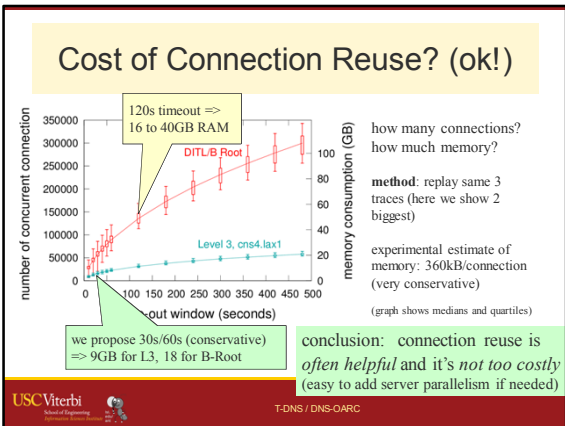


- ### T-DNS: TCP and TLS Connections
- introduction
 - why
 - how
 - **at minimal cost**
 - better than alternatives
 - next steps

- ### (Review) Our Contributions
3. **analysis: don't fear connections for DNS**
 - **client latency: only modestly more**
 - **server memory: well within current hardware**
 2. implementation choices to get here
 1. small protocol addition: TLS upgrade
- (going in reverse order)*

- ### (Review) Our Contributions
3. **analysis: don't fear connections for DNS**
 - **client latency: only modestly more**
 - **server memory: well within current hardware**
- questions:
- a. connection reuse: hit rate? memory?
 - b. CPU cost?
 - c. latency:
 - i. stub-recursive?
 - ii. recursive-authoritative?
 - iii. end-to-end?





End-to-End Latency: Methodology

- controlled experiments are hard
 - variable stub query timing
 - caching at recursive resolver
 - different RTTs (many stubs and authoritatives)
- approach: *model expected latency*
 - i.e., just averages
 - median connection reuse from trace replay
 - other parameters from experiments

USC Viterbi School of Engineering T-DNS / DNS-OARC

End-to-End Latency: Results

protocol choices: stub-recursive and recursive-authoritative

method: modeling; vary stub-recursive RTT; assumes all optimizations (TCP, TLS resumption, pipelining, OOP)

(graph shows expected values, plus slowdown relative to case (a), UDP/UDP)

TLS (s-r, 30s t.o.) + TCP (r-a, 60s t.o.)
19 to 33% slower: modest cost -> most benefit

USC Viterbi School of Engineering T-DNS / DNS-OARC

T-DNS: TCP and TLS Connections

- introduction
- why
- how
- at minimal cost
- better than alternatives**
- next steps

USC Viterbi School of Engineering T-DNS / DNS-OARC

Alternatives

- for improving privacy
 - DNSCurve/DNSCrypt: some neat optimizations to reduce RTTs, but new and fixed stack
 - DNS over DTLS: adds back UDP limits but still stuck with most TLS RTTs
- for anti-DoS
 - on others: rate limiting
- for relaxing limits:
 - seeming alternative: live within UDP limits

USC Viterbi School of Engineering T-DNS / DNS-OARC

T-DNS: TCP and TLS Connections

- introduction
- why
- how
- at minimal cost
- better than alternatives
- next steps**

USC Viterbi School of Engineering T-DNS / DNS-OARC

T-DNS Next Steps

- more information:
 - tech report ISI-TR-2014-688 (www.isi.edu/~johnh/PAPERS/Zhu14a/)
 - internet-draft: draft-hzhwm-start-tls-for-dns-01
- code:
 - client, client & server proxies, unbound patch
 - <http://www.isi.edu/ant/software/>
- do you want DNS privacy? share feedback?
 - johnh@isi.edu

USC Viterbi School of Engineering T-DNS / DNS-OARC