Recent and Future Developments in DNS Security

Duane Wessels

NANOG On The Road, Herndon VA

June 23, 2015

Outline

- Introduction to DNSSEC
- Root zone KSK rollover
- Root zone ZSK length increase
- DNS privacy
- DNS over TCP
- DANE

Introduction to DNSSEC

DNSSEC Overview

- “DNS Security Extensions”
- Extends the traditional DNS protocol so that consumers of DNS data can verify its authenticity
  - Sometimes called “data origin authentication”
- Based on public key cryptography
- Designed to detect response spoofing, cache poisoning, etc.
- May be leveraged to provide new protocol security features (DANE, TLSA, SSHFP, DKIM, etc).

DNSSEC High Level

- A public/private key pair is associated with each DNS zone
- Zone owner signs the zone with the private key
- Validating recursive name servers use public key to verify data authenticity
- Public keys are published in the zone itself
- A chain-of-trust must exist from the root zone to a leaf zone.

DNSSEC Record Types

DNSKEY	A zone's public key material
RRSIG	A signature over data (records)
DS	Delegation Signer; chain-of-trust between zones
NSEC	Next Secure; authenticated denial of existence
NSEC3	More sophisticated version of NSEC
NSEC3PARAM	Parameters for NSEC3

DNSSEC KSK / ZSK Split

- Best current practice is to have two types of DNSKEYs
- Key Signing Key (KSK)
  - Signs only the DNSKEY records
  - Changes infrequently (e.g., years)
  - Perhaps stored “offline”
- Zone Signing Key (ZSK)
  - Signs everything else
  - Changes frequently (e.g., months)
  - Probably stored “online”

DNSSEC Algorithms

5	RSA/SHA1
6	DSA-NSEC3-SHA1
7	RSA-NSEC3-SHA1
8	RSA/SHA256
10	RSA/SHA512
12	ECC-GOST
13	ECDSA Curve P-256 SHA256
14	ECDSA Curve P-384 SHA256

DNSSEC Validation

- Ensure clock is correct (NTP synchronized) on validator machine
- Add root zone trust anchor to your recursive name server configuration
- Enable validation in your recursive name server software

DNSSEC Testing with ‘dig’

```
$ dig +dnssec www.example.com

; <<>> DiG 9.9.5-3ubuntu0.2-Ubuntu <<>> +dnssec www.example.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 34139
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 2, AUTHORITY: 3, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags: do; udp: 4096
;; QUESTION SECTION:
;www.example.com.               IN      A

;; ANSWER SECTION:
www.example.com.        5       IN      A       93.184.216.34
www.example.com.        5       IN      RRSIG   A 8 3 5 20150623004549 20150615192902 6495
example.com. hE36fc8TO9SJyzTXwTH5zg44u6JLIZpNHvwcx26rkdGUhMNdlNPmUwuZ
Va54Bj575vCERzZGqpYd07q3/5ZWPgSdxZXXg4PiIy/oL2TDLgdjV4a/ hORzDa0Rj8kHABvHCu5b
+CHjwAs08vqBi1nNxay6rWAohQ6MAdR7Md3R HSg=

;; AUTHORITY SECTION:
example.com.            5       IN      NS      a.iana-servers.net.
example.com.            5       IN      NS      b.iana-servers.net.
example.com.            5       IN      RRSIG   NS 8 2 5 20150623044050 20150615192902 6495
example.com. MgX5rabwevSTkzVxSBhCcVbjPll7nF9Ka8f0BmYTSUXJ7KVOdH2cdXPv
YrRR9Kfdm9ZllhuzwqFpXKTpTD6Ukah3TTdCouuiGdOrZJDBkxraa/Xh
8tFa8IYyRdmu9vT3Wb6dSLOYgY8fIsMfhoG3Sl32apT3cIpQhrSNB5Y1 Vgo=
```

DNSSEC Testing with DNSSEC Analyzer

- http://dnssec-analyzer.verisignlabs.com

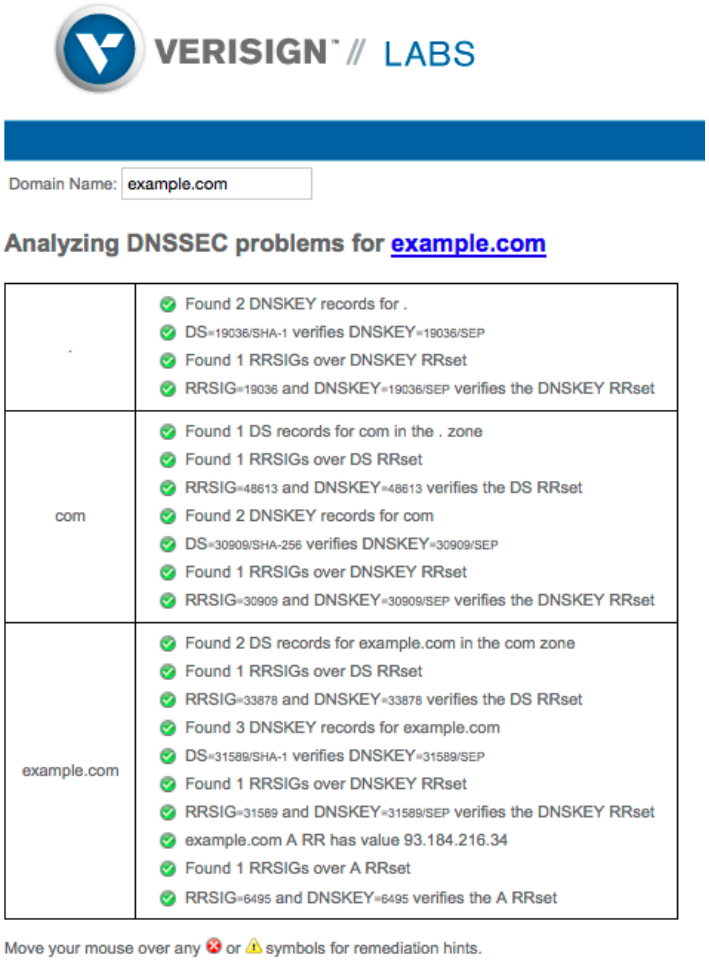

DNSSEC Testing with DNSViz

- http://dnsviz.net

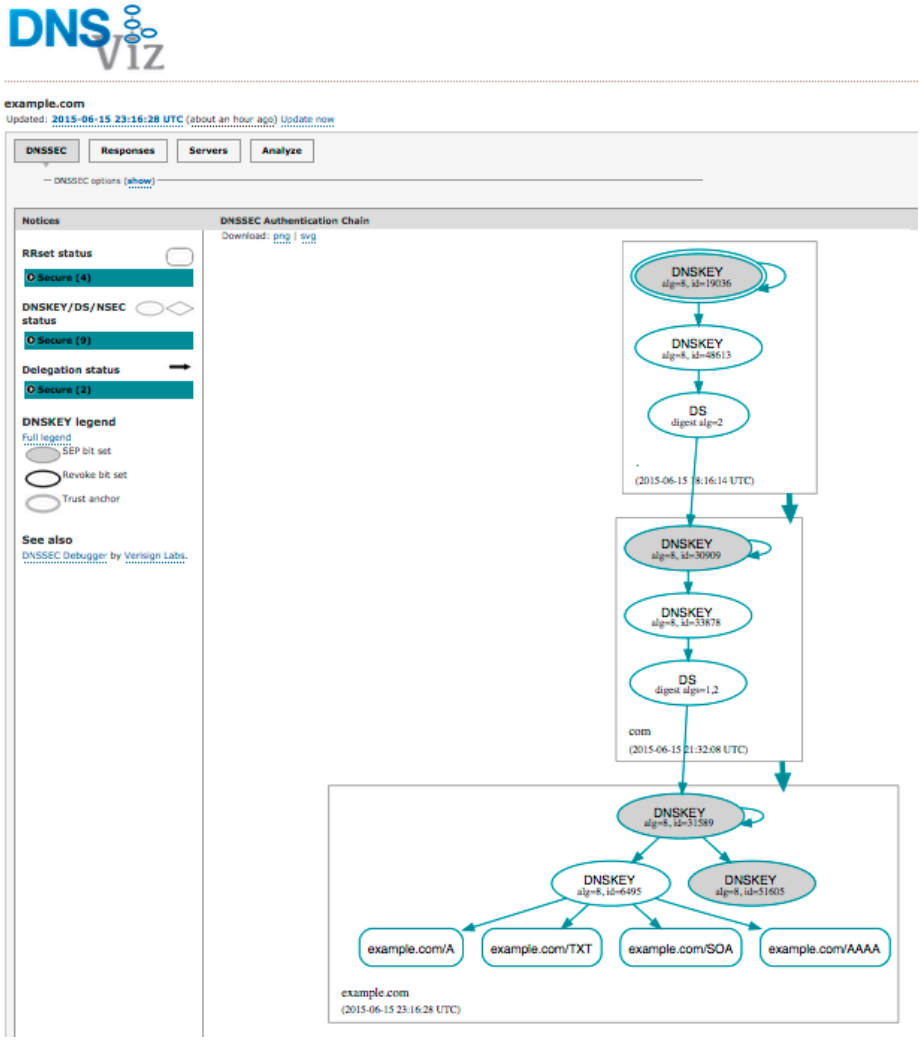

Root Zone Key Signing Key Rollover

Root Zone KSK

- The DNSSEC root “trust anchor”
- Signs DNSKEY RRset
  - Quarterly, at ICANN key signing ceremonies
- 2048-bit RSA key, algorithm 8
- Unchanged since July 2010

KSK Rollover Design Team

- Design team of seven experts
  - Joe Abley, Jaap Akkerhuis, John Dickinson, Geoff Huston, Ondrej Sury, Paul Wouters, Yoshiro Yoneya
- Plus Root Zone Management partners
  - ICANN
  - Verisign
  - U.S. Department of Commerce, National Telecommunications and Information Administration (NTIA)

KSK Rollover Concerns

- Rollovers lead to larger responses
  - Only for ‘DNSKEY’ responses
  - Still below typical Ethernet/IPv4 MTU sizes (1500)
  - But above minimum IPv6 MTU (1280)
  - May lead to more fragmentation and/or truncation
- Automatic Trust Anchor Updates
  - RFC 5011
  - Not yet tested at this scale

- Check that networks allow IP fragments and DNS-over-TCP
- Check validators for RFC 5011 automatic updates

KSK Rollover Next Steps

- Design Team to complete its work by end of June
- 40 day ICANN comment period
- Additional month to prepare final report

- Root Zone Management partners then develop plan for execution

Root Zone Zone Signing Key Length

Root Zone ZSK

- Signs most records in the root zone
  - everything except DNSKEY records
- 1024-bit RSA, algorithm 8
- Rolled every 90 days

- Recent concerns that 1024-bit RSA keys are weak

Why 1024-bit?

- ZSK key length is defined in the requirements document from NTIA
- The concerns regarding the key length of the ZSK were discussed among the Root Zone Management Partners back in 2009
- Root Zone Management Partners agreed to make an exception due to the packet size concerns
- The ZSK key length was clearly communicated to the Internet community at-large at multiple venues to solicit input
- The specification of the ZSK was intended to be reconsidered and planned when the KSK change/rollover happens
- The KSK change/rollover was delayed

Increasing ZSK size

- Verisign is investigating the requirements and consequences of increasing the size of the root zone Zone Signing Key.

- How would such changes affect DNS traffic?
  - Response sizes
  - Bandwidth
  - Truncation
  - Fragmentation

Cumulative Distribution of All Response Sizes

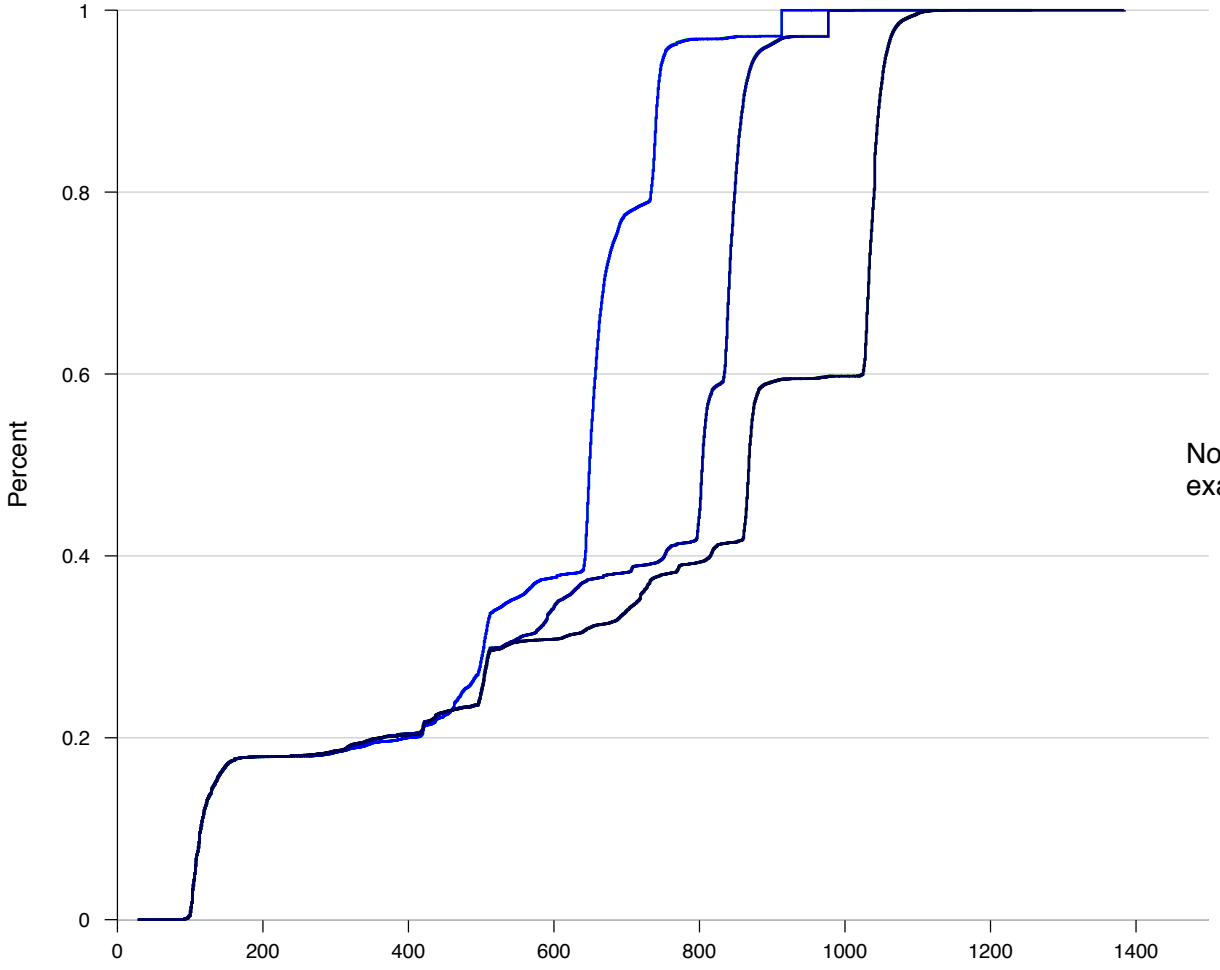

Note that ZSK Roll lines exactly match the non-Roll lines.

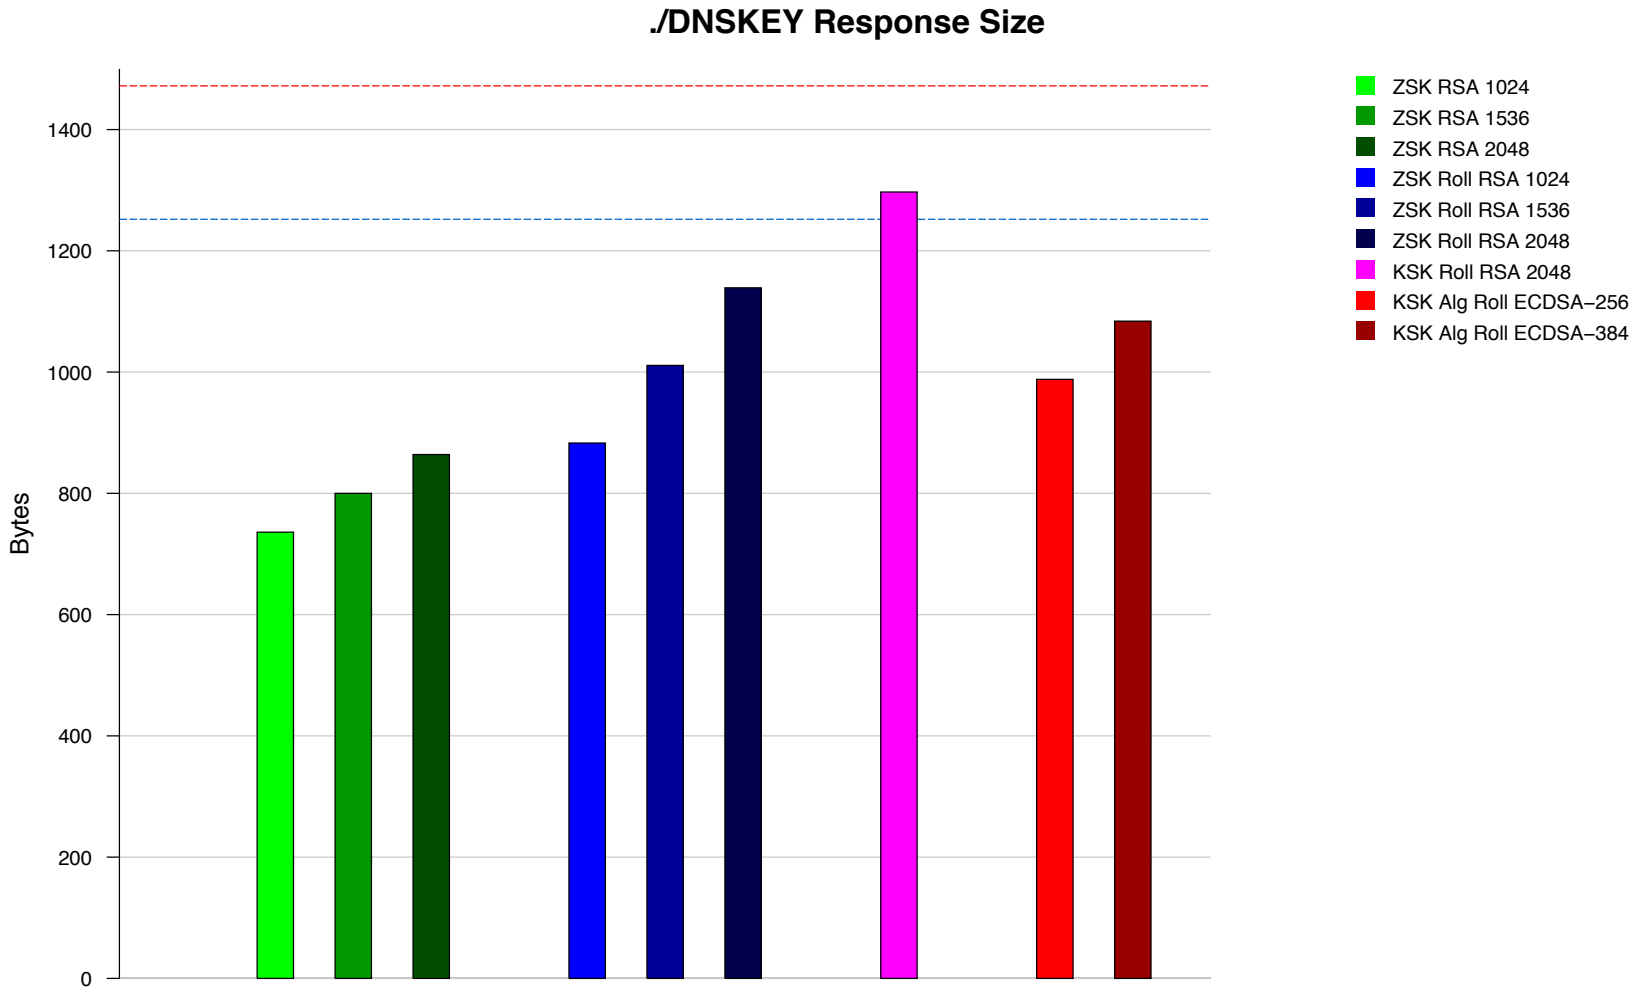
./DNSKEY Response Size
Bytes
0
200
400
600
800
1000
1200
1400
ZSK RSA 1024
ZSK RSA 1536
ZSK RSA 2048
ZSK Roll RSA 1024
ZSK Roll RSA 1536
ZSK Roll RSA 2048
KSK Roll RSA 2048
KSK Alg Roll ECDSA–256
KSK Alg Roll ECDSA–384

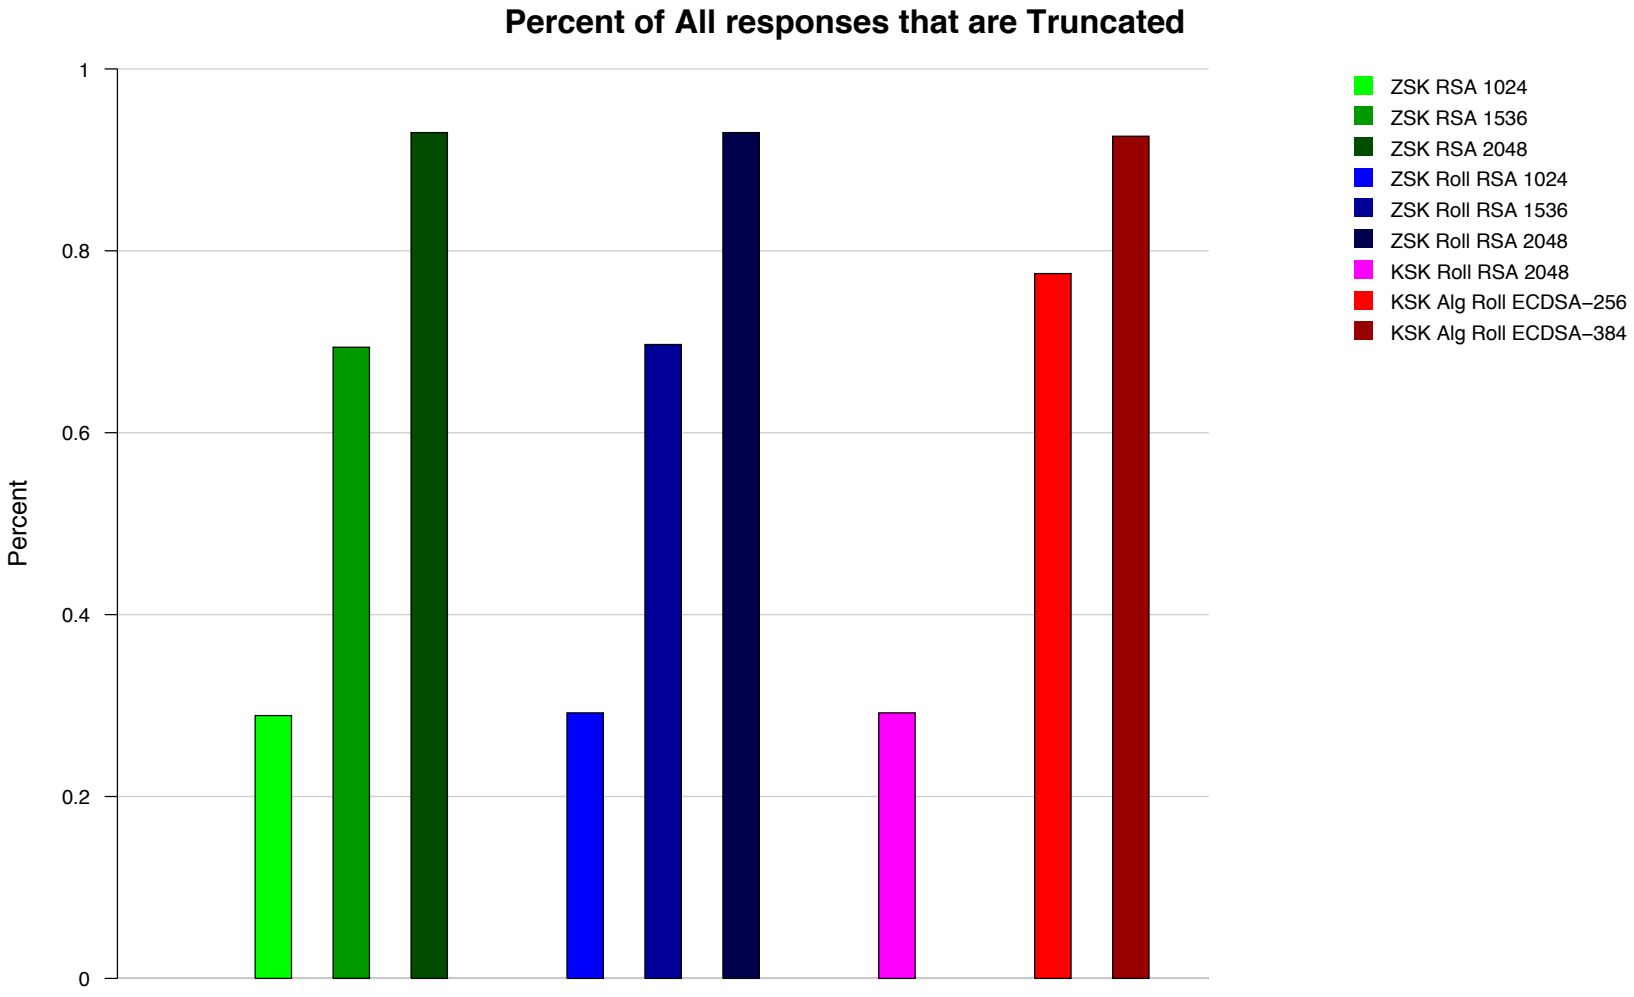
Percent of All responses that are Truncated
Percent
1
0.8
0.6
0.4
0.2
0
ZSK RSA 1024
ZSK RSA 1536
ZSK RSA 2048
ZSK Roll RSA 1024
ZSK Roll RSA 1536
ZSK Roll RSA 2048
KSK Roll RSA 2048
KSK Alg Roll ECDSA–256
KSK Alg Roll ECDSA–384

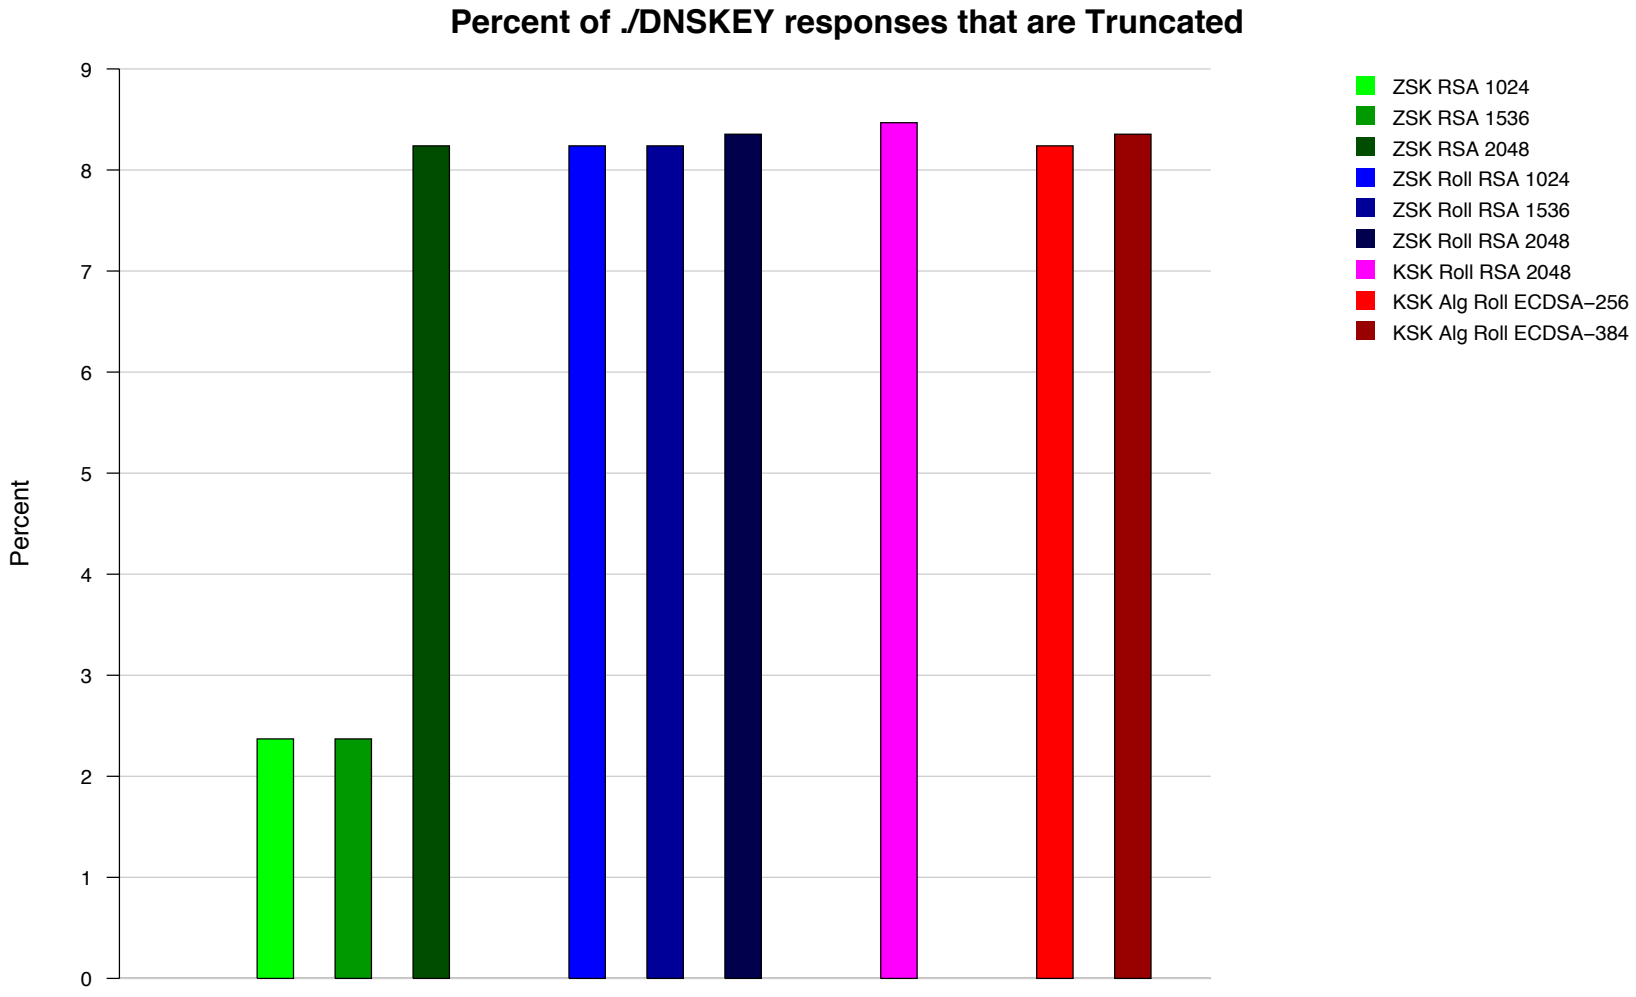
Percent of ./DNSKEY responses that are Truncated
Percent
0
1
2
3
4
5
6
7
8
9
ZSK RSA 1024
ZSK RSA 1536
ZSK RSA 2048
ZSK Roll RSA 1024
ZSK Roll RSA 1536
ZSK Roll RSA 2048
KSK Roll RSA 2048
KSK Alg Roll ECDSA–256
KSK Alg Roll ECDSA–384

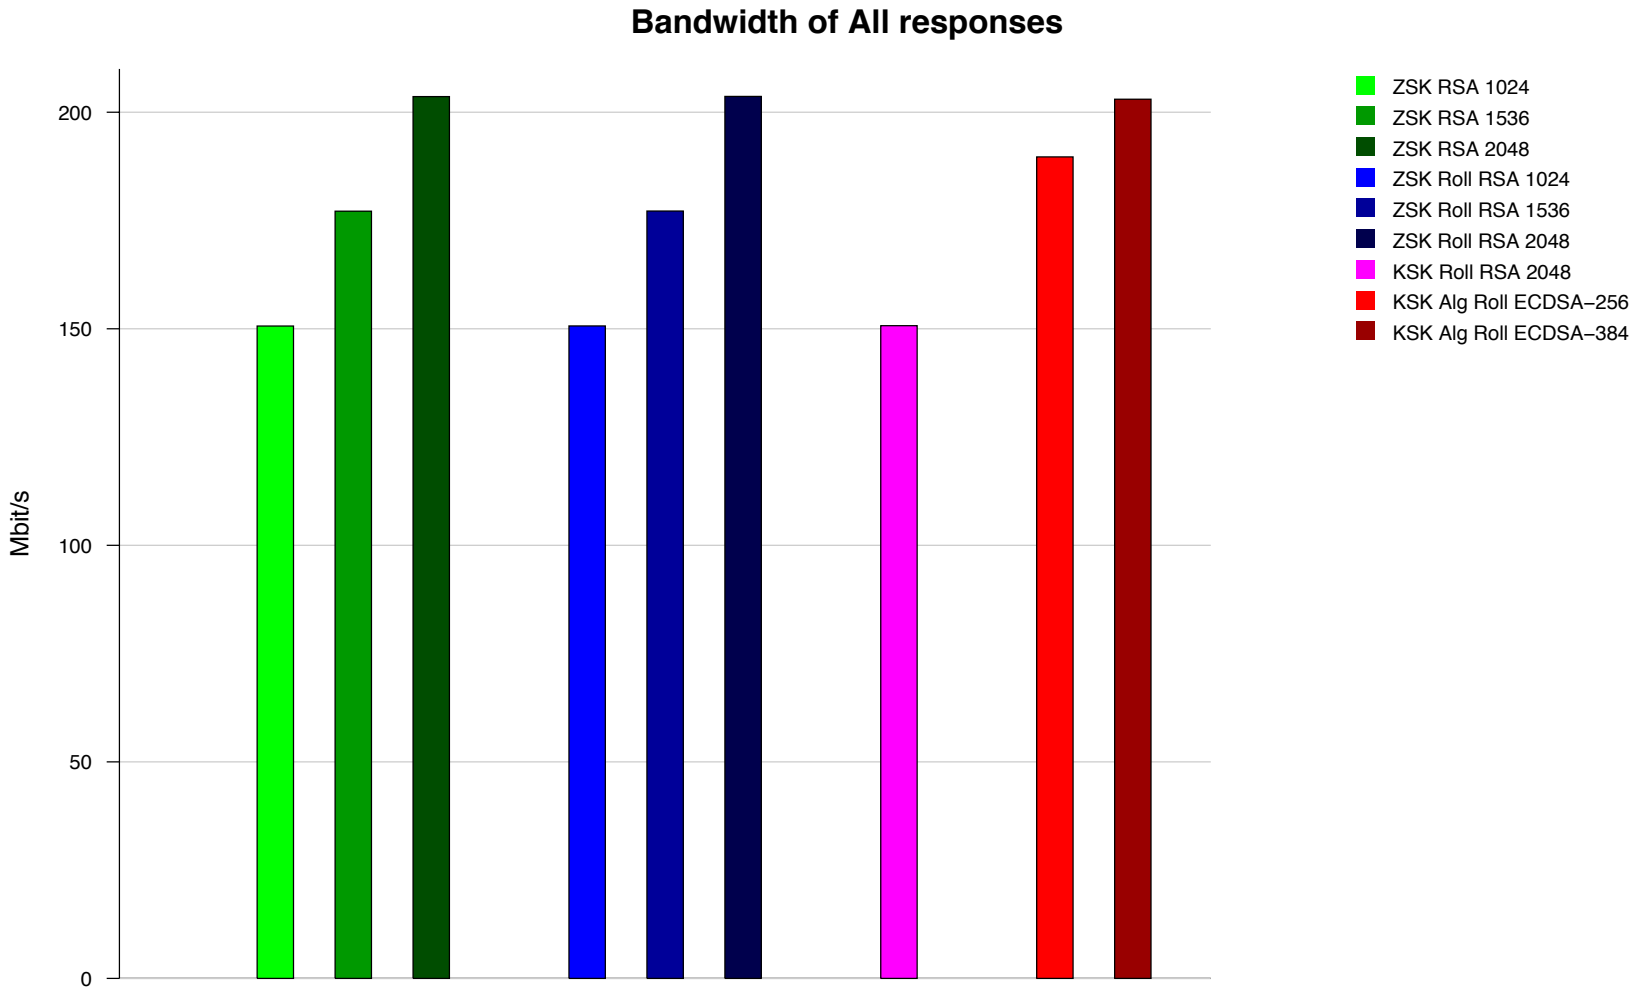
Bandwidth of All responses
Mbit/s
0
50
100
150
200
ZSK RSA 1024
ZSK RSA 1536
ZSK RSA 2048
ZSK Roll RSA 1024
ZSK Roll RSA 1536
ZSK Roll RSA 2048
KSK Roll RSA 2048
KSK Alg Roll ECDSA–256
KSK Alg Roll ECDSA–384

DNS Privacy

RFC 7258 - Pervasive Monitoring Is an Attack

“The IETF community's technical assessment is that PM is an attack on the privacy of Internet users and organisations.”

“The IETF community has expressed strong agreement that PM is an attack that needs to be mitigated where possible, via the design of protocols that make PM significantly more expensive or infeasible.”

DNS Lacks Privacy

- Query names are cleartext
- Full query names are sent at every step
  - e.g., root name servers see “foo.bar.example.com”
- Caching helps
- Proxying by recursive name servers helps
  - But edns-client-subnet doesn’t

Query names are cleartext

- Perhaps you've run tcpdump on a DNS server before

```
18:51:38.759042 IP x.x.x.x.56584 > 198.41.0.10.domain: 51639+ [1au] A?
clit.www.astl8.com. (47)
18:51:38.763768 IP x.x.x.x.25319 > 198.41.0.10.domain: 7834+ A?
aligntech.co.kr. (33)
18:51:38.769173 IP x.x.x.x > 198.41.0.10: ICMP host x.x.x.x unreachable -
admin prohibited filter, length 36
18:51:38.771503 IP x.x.x.x.57281 > 198.41.0.10.domain: 56647% [1au] A?
x.x.x.xM-^?^?. (43)
18:51:38.773810 IP x.x.x.x.16214 > 198.41.0.10.domain: 33833 A?
bwyeupmjjr.Home. (33)
18:51:38.773954 IP x.x.x.x.18925 > 198.41.0.10.domain: 29652% [1au] A?
ns2.securitynet.cz. (47)
18:51:38.775071 IP x.x.x.x.65302 > 198.41.0.10.domain: 46003 A?
ns11.dnsmadeeasy.com. (38)
18:51:38.775676 IP x.x.x.x.47968 > 198.41.0.10.domain: 4375% [1au] NS? .
(28)
18:51:38.775775 IP x.x.x.x.37798 > 198.41.0.10.domain: 28054% [1au] A?
gruppomgcombr.gruppomg.neen.it. (59)
18:51:38.777471 IP x.x.x.x.56052 > 198.41.0.10.domain: 48957% [1au] A?
profile-images.scdn.co. (51)
```

Full Query Names

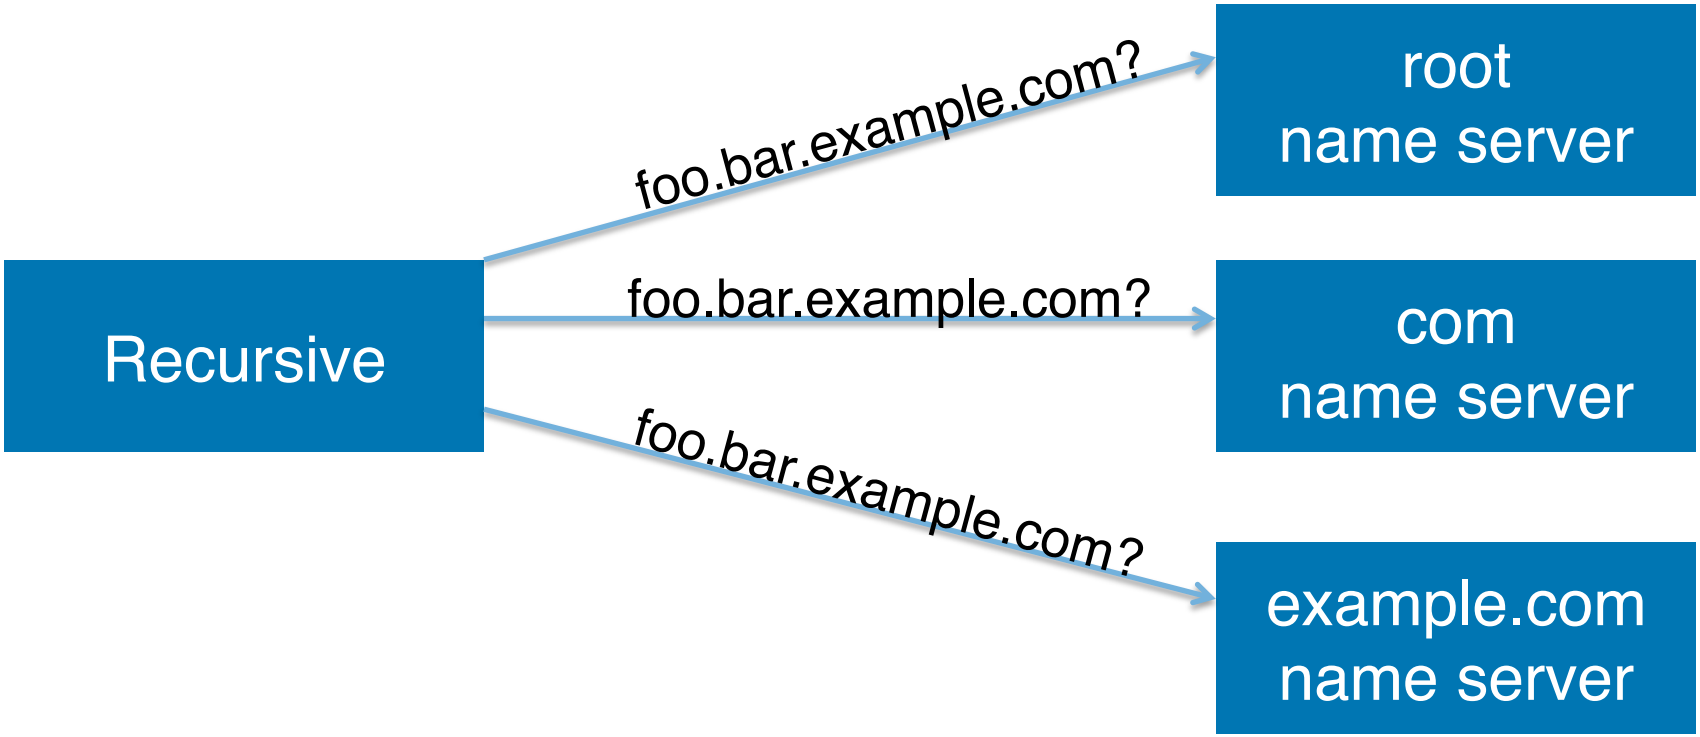

Qname Minimization

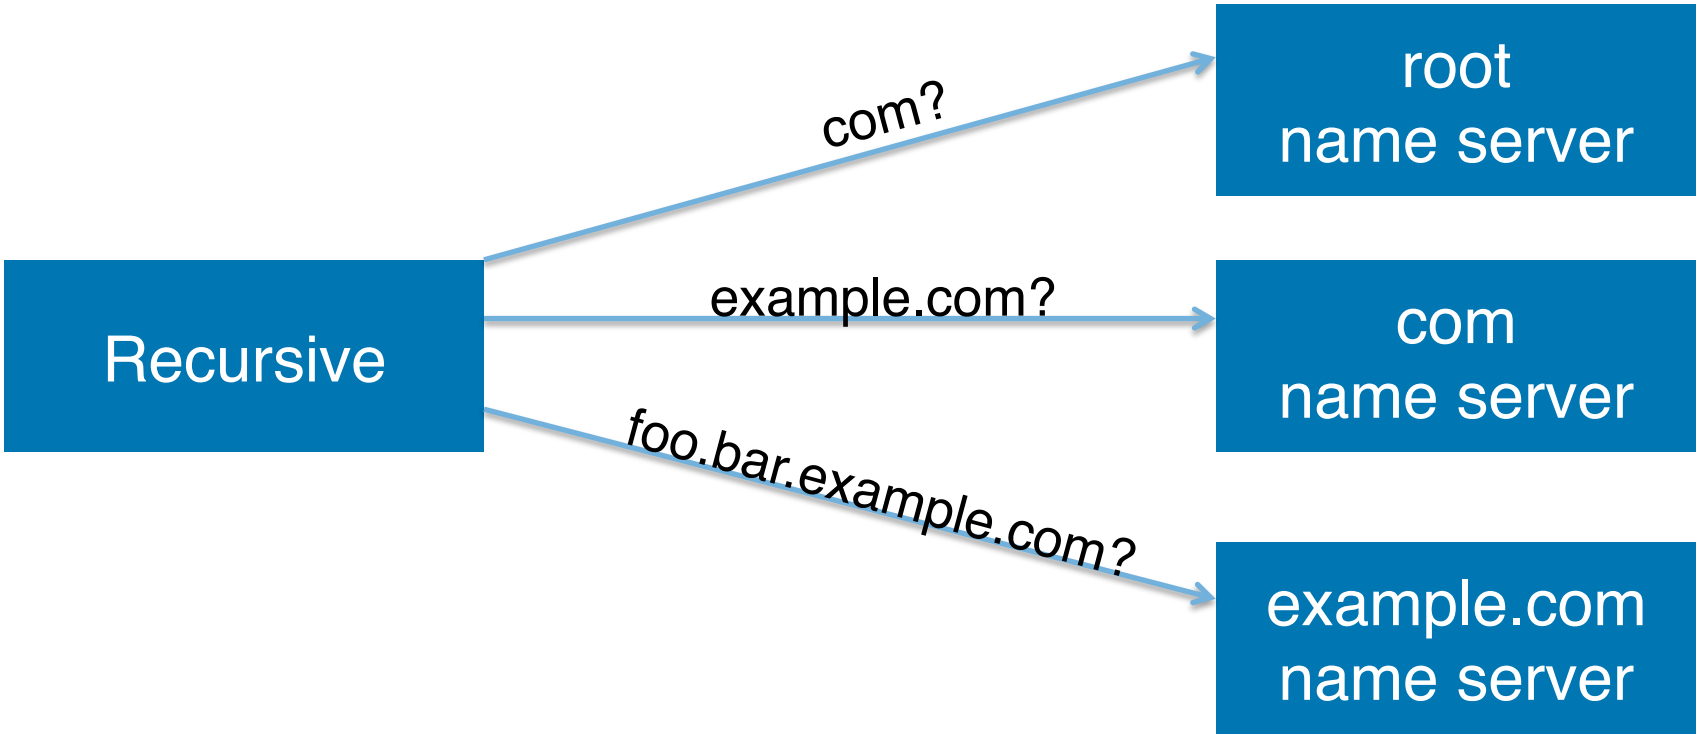

draft-ietf-dnsop-qname-minimisation-03
“DNS query name minimisation to improve privacy”

EDNS Client Subnet

- Problem:
  - Content Distribution Networks prefer to give different DNS answers depending on “where” you are
  - Recursive name servers obscure end-user IP addresses
  - Large (public) recursive name servers have users from everywhere.
- Solution:
  - Include end-user IP address data in DNS requests
  - Masked by some size netmask
  - Goes in EDNS0 OPT record
  - draft-ietf-dnsop-edns-client-subnet-01

EDNS Client Subnet Drawbacks

- End-user IP addresses (networks) no longer obscured by recursive
  - There is an opt-out mechanism of sorts
- Adversely affects caching
- Could allow someone to enumerate CDN service addresses by “spoofing” wide range of edns-client-subnet values.
- Operationally, seems to require heavy whitelisting

Encrypting DNS Queries

- Encryption can prevent man-in-the-middle eavesdropping.
- TLS for DNS
  - Run DNS over TLS session
  - Requires TCP
  - draft-ietf-dprive-start-tls-for-dns (disclosure: I'm a coauthor)
- DNS over DTLS
  - Datagram Transport Layer Security
  - UDP only
  - draft-wing-dprive-dnsodtls
- Confidential DNS
  - No particular transport requirements (UDP or TCP)
  - Server publishes encryption keys in new ENCRYPT record type
  - draft-wijngaards-dnsop-confidentialdns

DNS over TCP

Why worry about DNS over TCP?

- Connectionless UDP trivially spoofed
- DNS is protocol-of-choice for large scale DDoS attacks
- Increase in DNS response sizes
- Want privacy (i.e., TLS)
- Implementations need to improve their TCP support

Spoofing & Attacks

- Very easy for miscreants to spoof source addresses
- TCP's three way handshake raises the bar
- However, TCP requires other attack protections
  - SYN flooding
  - State exhaustion
  - Bogus resets
  - etc
- Response Rate Limiting (RRL) is a common defense
  - In RRL, some responses "slip" through the rate limiter, but are truncated, perhaps resulting in a TCP connection.
- Also recently revived proposal for "DNS cookies"

Increasing Response Sizes

- Trend for larger and larger DNS responses
  - Adding DNSSEC signatures (i.e., unsigned vs signed)
  - Longer DNSSEC keys (i.e., 1024 -> 2048 bit)
  - New cryptographic/security uses (DANE, TLSA, SMIMEA, OPENPGPKEY, SSHFP)

- As response sizes approach and exceed network MTUs, responses must be either truncated or fragmented.
  - Truncated -> TCP
  - Fragmented -> firewall problems and security vulnerabilities

Performance Improvements

- Connection reuse
  - Reuse connections when possible to avoid TCP handshake
- TCP Fast Open (RFC 7413)
  - Use cookies to avoid subsequent TCP handshakes
  - Similarly: TLS session resumption
- Pipelining
  - Send subsequent queries immediately, rather than wait for previous response
- Out-of-Order Processing
  - Clients: be prepared to receive out-of-order responses over TCP
  - Servers: send responses immediately when ready, possibly out-of-order
- See draft-ietf-dnsop-5966bis (disclosure: I'm a co-author)

DNS-Based Authentication of Named Entities (DANE)

DNS-Based Authentication of Named Entities

- What is a “named entity?”
  - well, things that have names!
  - versus, say, addresses
  - TLS/X509 certificates (servers & services)
  - Email addresses, Jabber IDs (people)
  - ?
- DANE is an IETF working group, and umbrella term for a set of specific protocols.
  - TLSA
  - SMIMEA, OPENPGPKEY
  - SSHFP

DANE TLSA

- RFC 6698, RFC 7218
- Defines TLSA record type
- Associates a domain name and a TLS certificate.
- Constrains or replaces traditional system of PKIX Certificate Authorities.

Without DANE TLSA

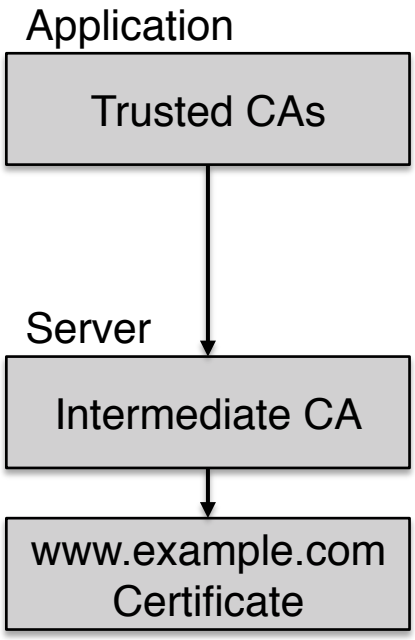

TLSA PKIX-TA CA constraint

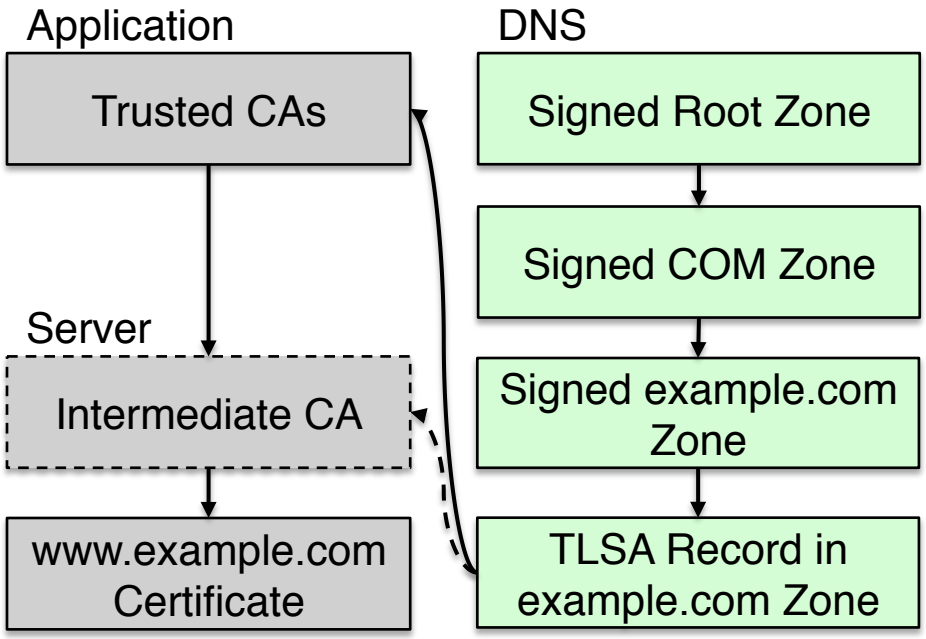

TLSA PKIX-EE Service certificate constraint

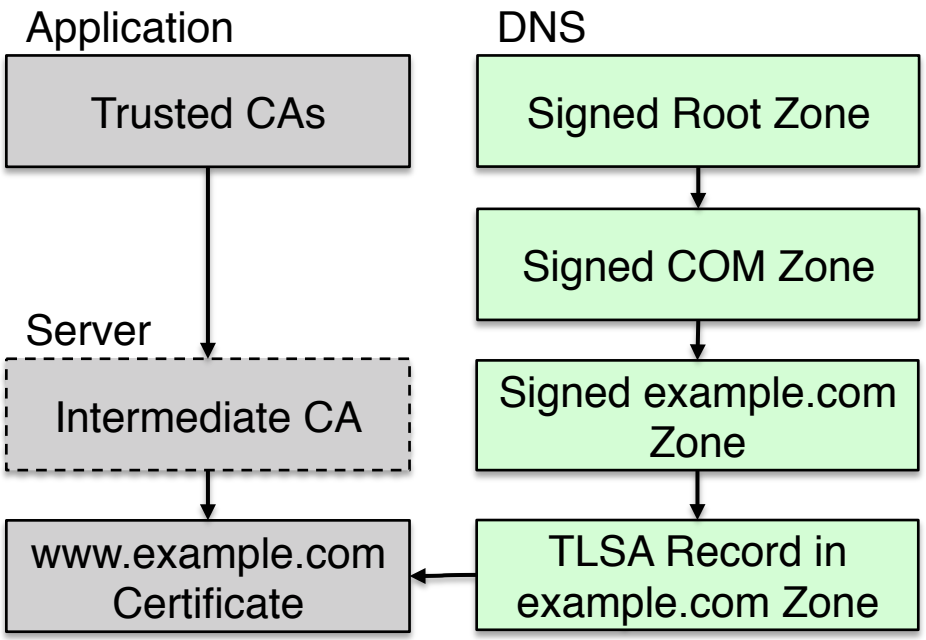

TLSA DANE-TA Trust anchor assertion

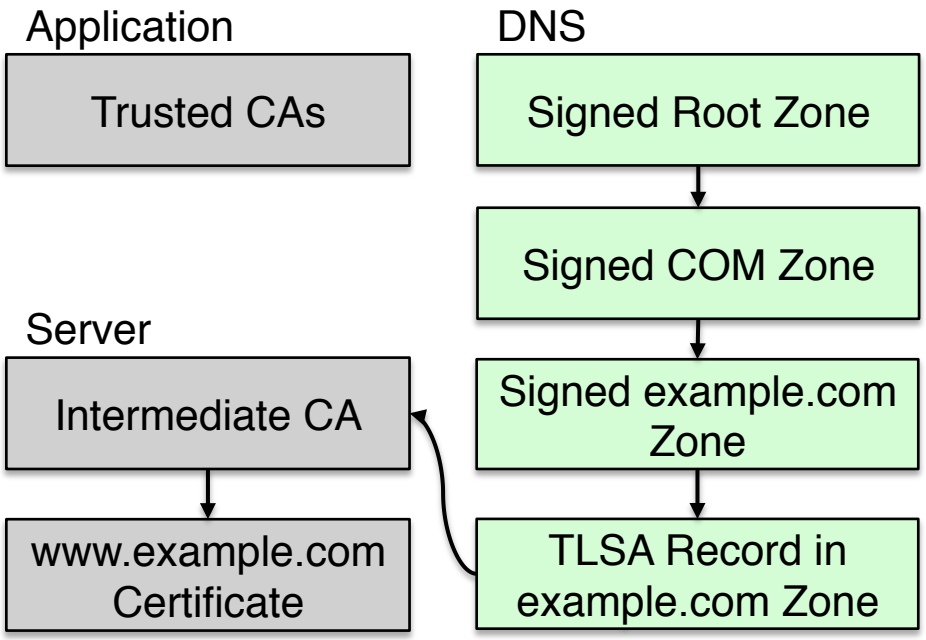

TLSA DANE-EE Domain-issued certificate

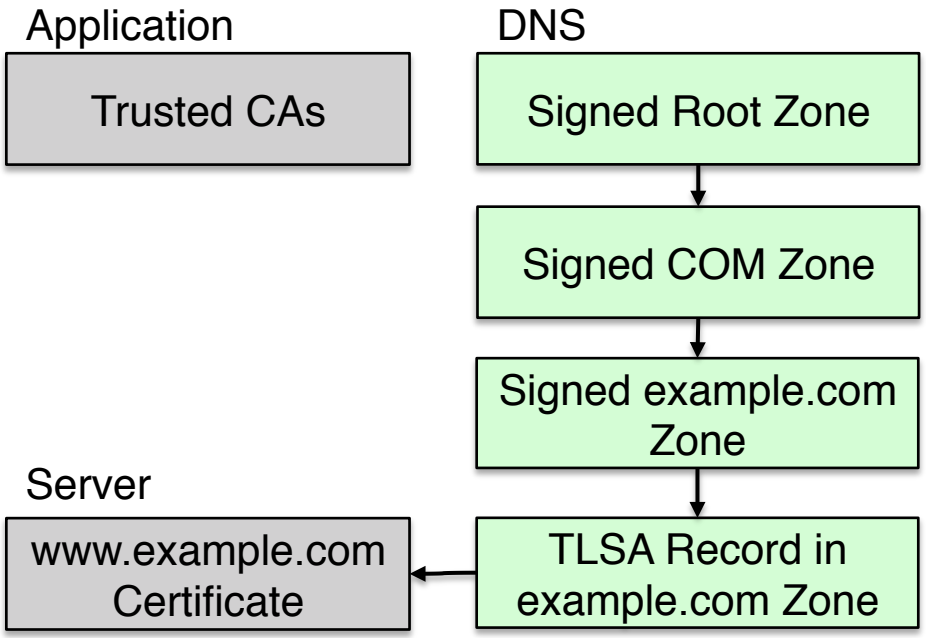

DANE SMIMEA, OPENPGPKEY

- Proposed DNS record types to encode security data associated with email addresses.
- draft-ietf-dane-smime
  - Very similar to TLSA record format
  - Base32 encoding of email address local part into a domain name
  - Susceptible to zone-enumeration
- draft-ietf-dane-openpgpkey
  - Much different than SMIMEA
  - record data is a PGP key “blob”
  - local-part email address is hashed

SSHFP

- SSH server fingerprint
- Pre-dates DANE
- See “VerifyHostKeyDNS” in OpenSSH

```
$ ssh packet-pushers.com
The authenticity of host '[packet-pushers.com] ([173.230.152.222])' can't
be established.
ECDSA key fingerprint is 25:61:95:05:ed:09:a1:69:f9:b1:dd:6b:fe:3b:5c:72.
Matching host key fingerprint found in DNS.
Are you sure you want to continue connecting (yes/no)?
```

Q & A

VERISIGN®