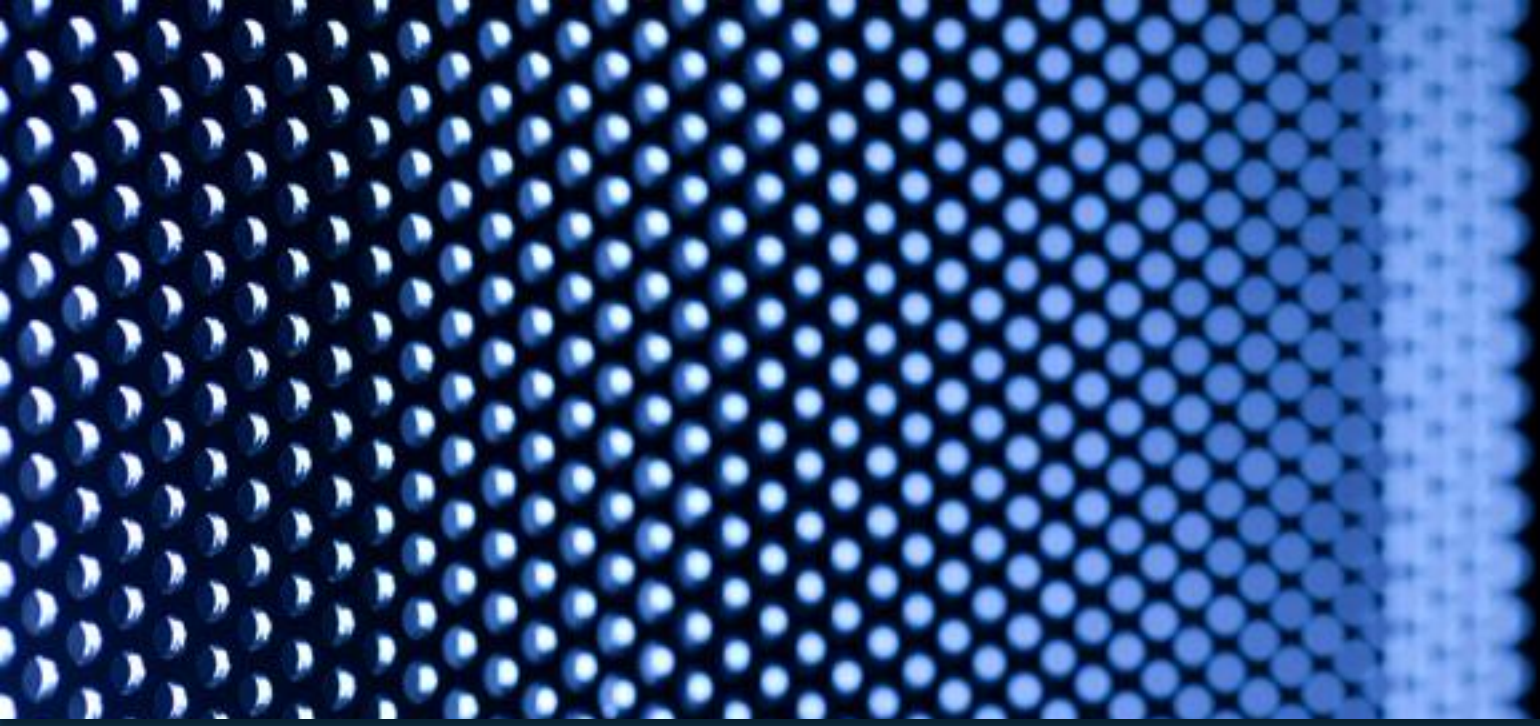

VERISIGN

Traffic Effects of Changing Root Zone Keys

Duane Wessels

DNS-OARC Workshop, Amsterdam

May 9, 2015

Motivation

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

- Resumed work on rolling/changing the root zone Key Signing Key.

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

Background

- 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

Disclaimer

- This work investigates a number of different scenarios, including:
  - A wide range of ZSK lengths
  - Changing the root zone DNSSEC algorithm.

- Verisign is not advocating for ZSK lengths beyond 2048-bits at this time
- Verisign is not advocating for a change to the root zone DNSSEC algorithm at this time.

Status Quo

- Root Zone KSK
  - 2048 bits
  - Rolled: <undef>
  - Signature Validity: 15 days
- Root Zone ZSK
  - 1024 bits
  - Rolled: quarterly (90 days)
  - Signature Validity: 10 days

Scenarios Simulated

- Increasing the root zone ZSK length
  - From 1024 to 1280 ... 4096 bits
- Rolling the root zone KSK
  - Same size and algorithm, just new key
- Changing the root zone KSK/ZSK algorithm
  - From RSA to ECDSA

Experiment Setup

- Create multiple copies of a signed root zone
- Various key sizes, key counts, algorithms
- Serve each root zone with its own named process
  - multiple named processes on loopback addresses
- Capture real root server traffic
- Replay traffic capture
  - qname and qtype
  - DO bit
  - EDNS0 UDP size

Traffic Replay

- For each UDP* query in traffic capture
  - Send UDP query to all named processes
  - Send TCP query to all named processes
- Record
  - Client DO bit
  - Client EDNS UDP size
  - Server RCODE
  - Server TC bit (UDP response)
  - UDP reply size
  - TCP reply size

* Captured TCP queries are ignored under the assumption they might be duplicates of previous UDP queries

Other Zones

- Since most roots also serve arpa, the simulation does as well.
  - With same KSK/ZSK parameters as root.
- Also configured to serve root-servers.net zone
  - Not signed

Sample Replay Output

```
#querynum servnum kskalg kskcnt ksksize zskalg zskcnt zsksize do edns
rcode tc udpsize tcpsize

# WFAWLANConfigSCPD.xml.sitecomwl341. 1
          5    0    8    1   2048    8    1   1024 1   1400 3 0    673    673
          5    1    8    1   2048    8    1   1280 1   1400 3 0    769    769
          5    2    8    1   2048    8    1   1536 1   1400 3 0    865    865
          5    3    8    1   2048    8    1   1792 1   1400 3 0    961    961
          5    4    8    1   2048    8    1   2048 1   1400 3 0   1057   1057
          5    5    8    1   2048    8    1   2304 1   1400 3 0   1153   1153
          5    6    8    1   2048    8    1   2560 1   1400 3 0   1249   1249
          5    7    8    1   2048    8    1   2816 1   1400 3 0   1345   1345
          5    8    8    1   2048    8    1   3072 1   1400 3 1   1026   1441
          5    9    8    1   2048    8    1   3328 1   1400 3 1   1090   1537
          5   10    8    1   2048    8    1   3584 1   1400 3 1   1154   1633
          5   11    8    1   2048    8    1   3840 1   1400 3 1   1218   1729
          5   12    8    1   2048    8    1   4096 1   1400 3 1   1282   1825
...
```

Quick Stats

- Zone File
  - SOA Serial 2015030401
- Input Trace:
  - March 4, 2015
  - 18:10:00 -- 18:20:00 UTC (10 minutes duration)
  - 46,415,453 IP packets captured
  - 23,638,876 DNS UDP queries captured
    - 39,400 queries/second
  - A-root sites: NYC3, LON3, LAX2, FRA1, HKG5

Quick Stats

- DO bit
  - 71% set
  - 29% clear
- RCODEs
  - 41% NOERROR
  - 59% NXDOMAIN
- Queries for root DNSKEY
  - .02 % of all queries
  - 2 out of 10,000

Caveats

- These simulations were done with BIND (9.8.2rc1)
- Other name server software might behave differently

Situations Simulated

Normal operations with different ZSK sizes

- Algorithm remains RSASHA256 (8)
- ZSK length varies from 1024 to 4096
- One RRSIG over all RRSets

- In graphs these are labeled “ZSK RSA xxxx”

ZSK Rollover for different ZSK lengths

- ZSK Rollover occurs quarterly
- For approx 20 day period
- Algorithm remains RSASHA256 (8)
- One RRSIG over all RRSets (pre-publish method)

- In graphs these are labeled “ZSK Roll RSA xxxx”

KSK Rollover

- Algorithm remains RSASHA256 (8)
- KSK length remains 2048-bits
- ZSK length remains 1024-bits
- Two RRSIGs over DNSKEY RRSet
- One RRSIG over other RRSets

- In graphs this is labeled “KSK Roll RSA 2048”

KSK Algorithm Roll

- Algorithm changes for both ZSK and KSK
  - ECDSAP256SHA256 (13)
  - ECDSAP384SHA384 (14)
- Outgoing ZSK length is 1024-bits
- Two RRSIGs over all RRSets

- In graphs this is labeled “KSK to ECDSA-xxx”

Results

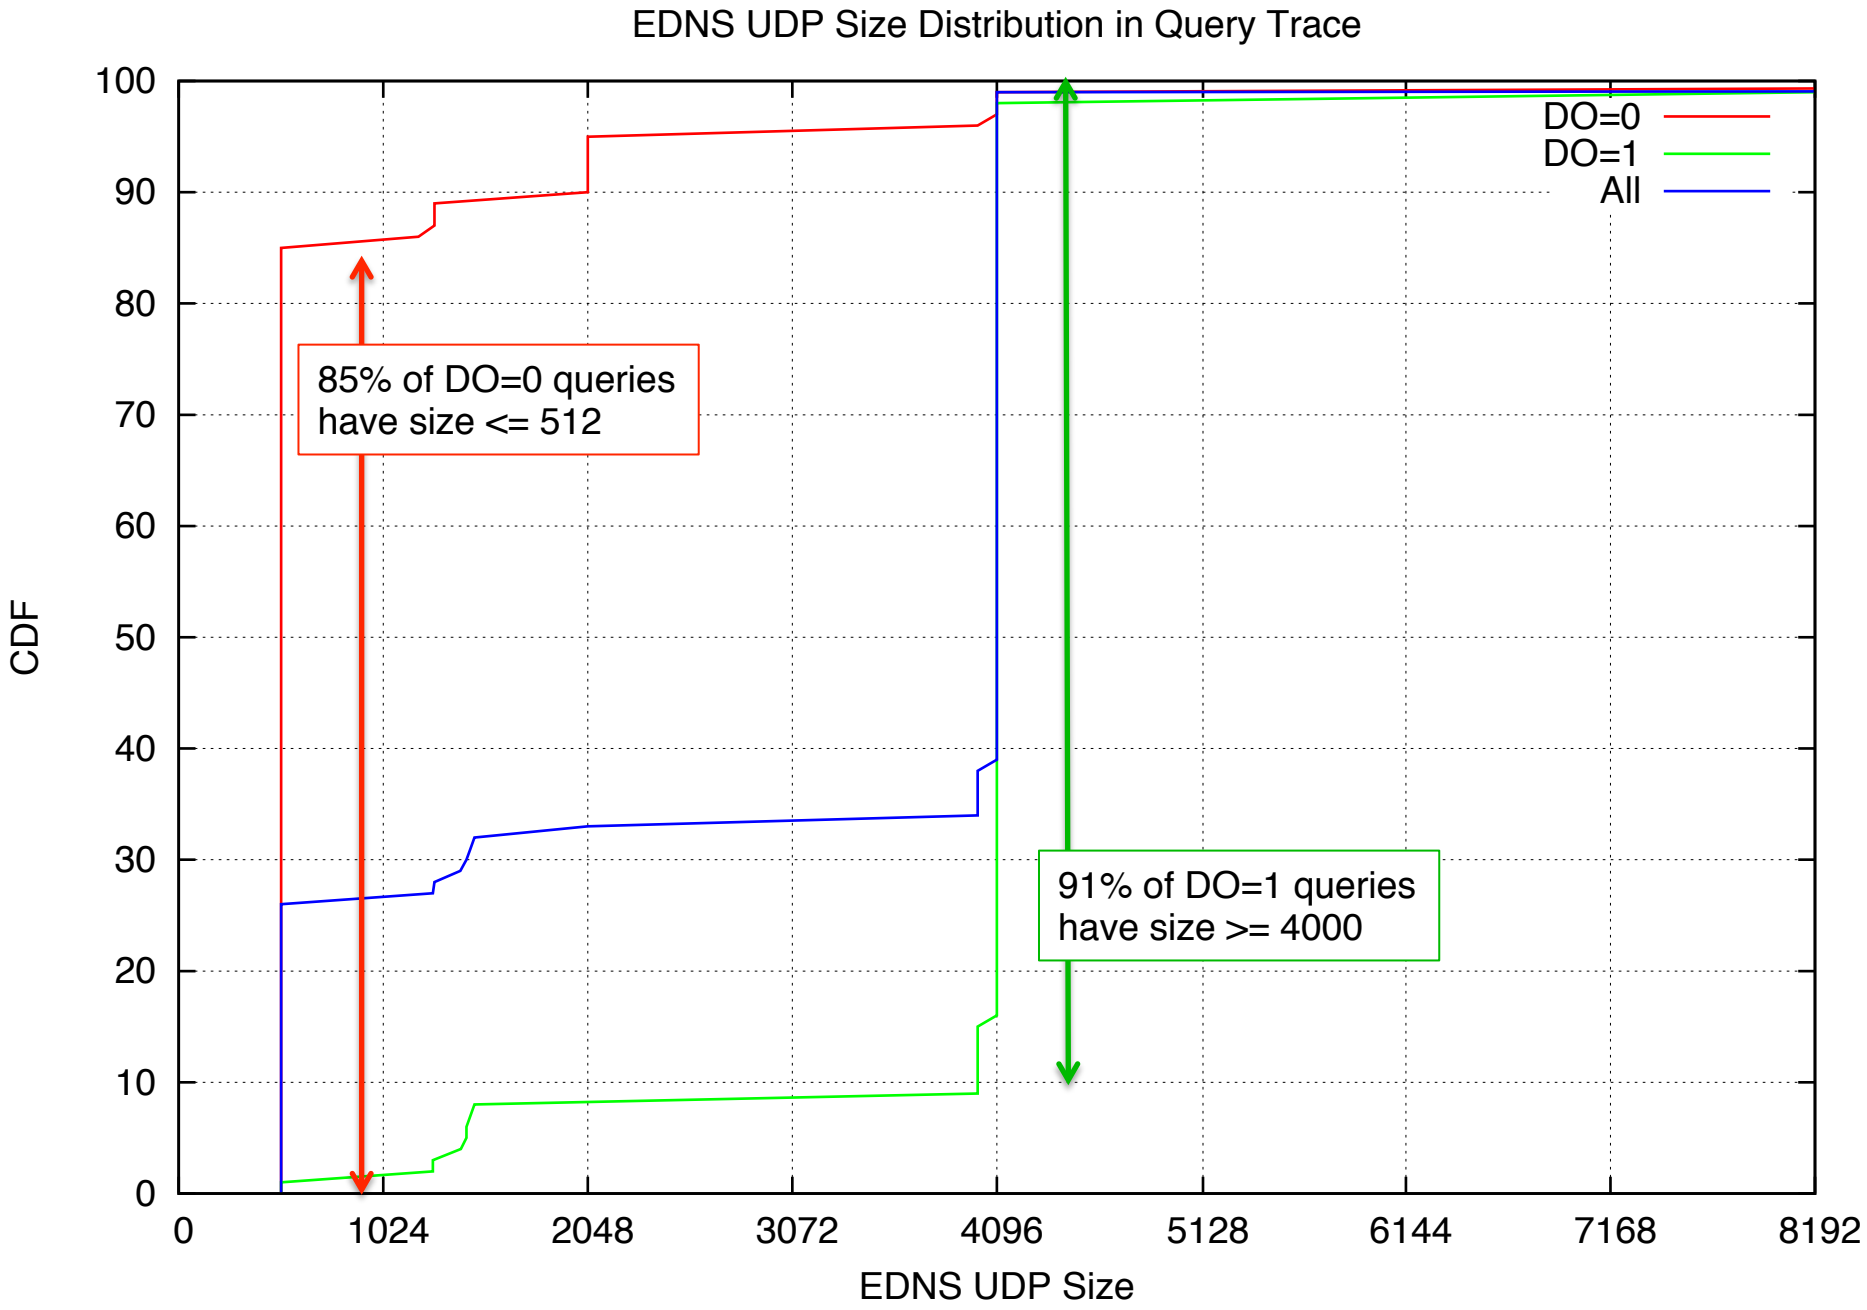
EDNS UDP Size Distribution in Query Trace
CDF
EDNS UDP Size
DO=0
DO=1
All
85% of DO=0 queries
have size <= 512
91% of DO=1 queries
have size >= 4000
0
10
20
30
40
50
60
70
80
90
100
0
1024
2048
3072
4096
5128
6144
7168
8192

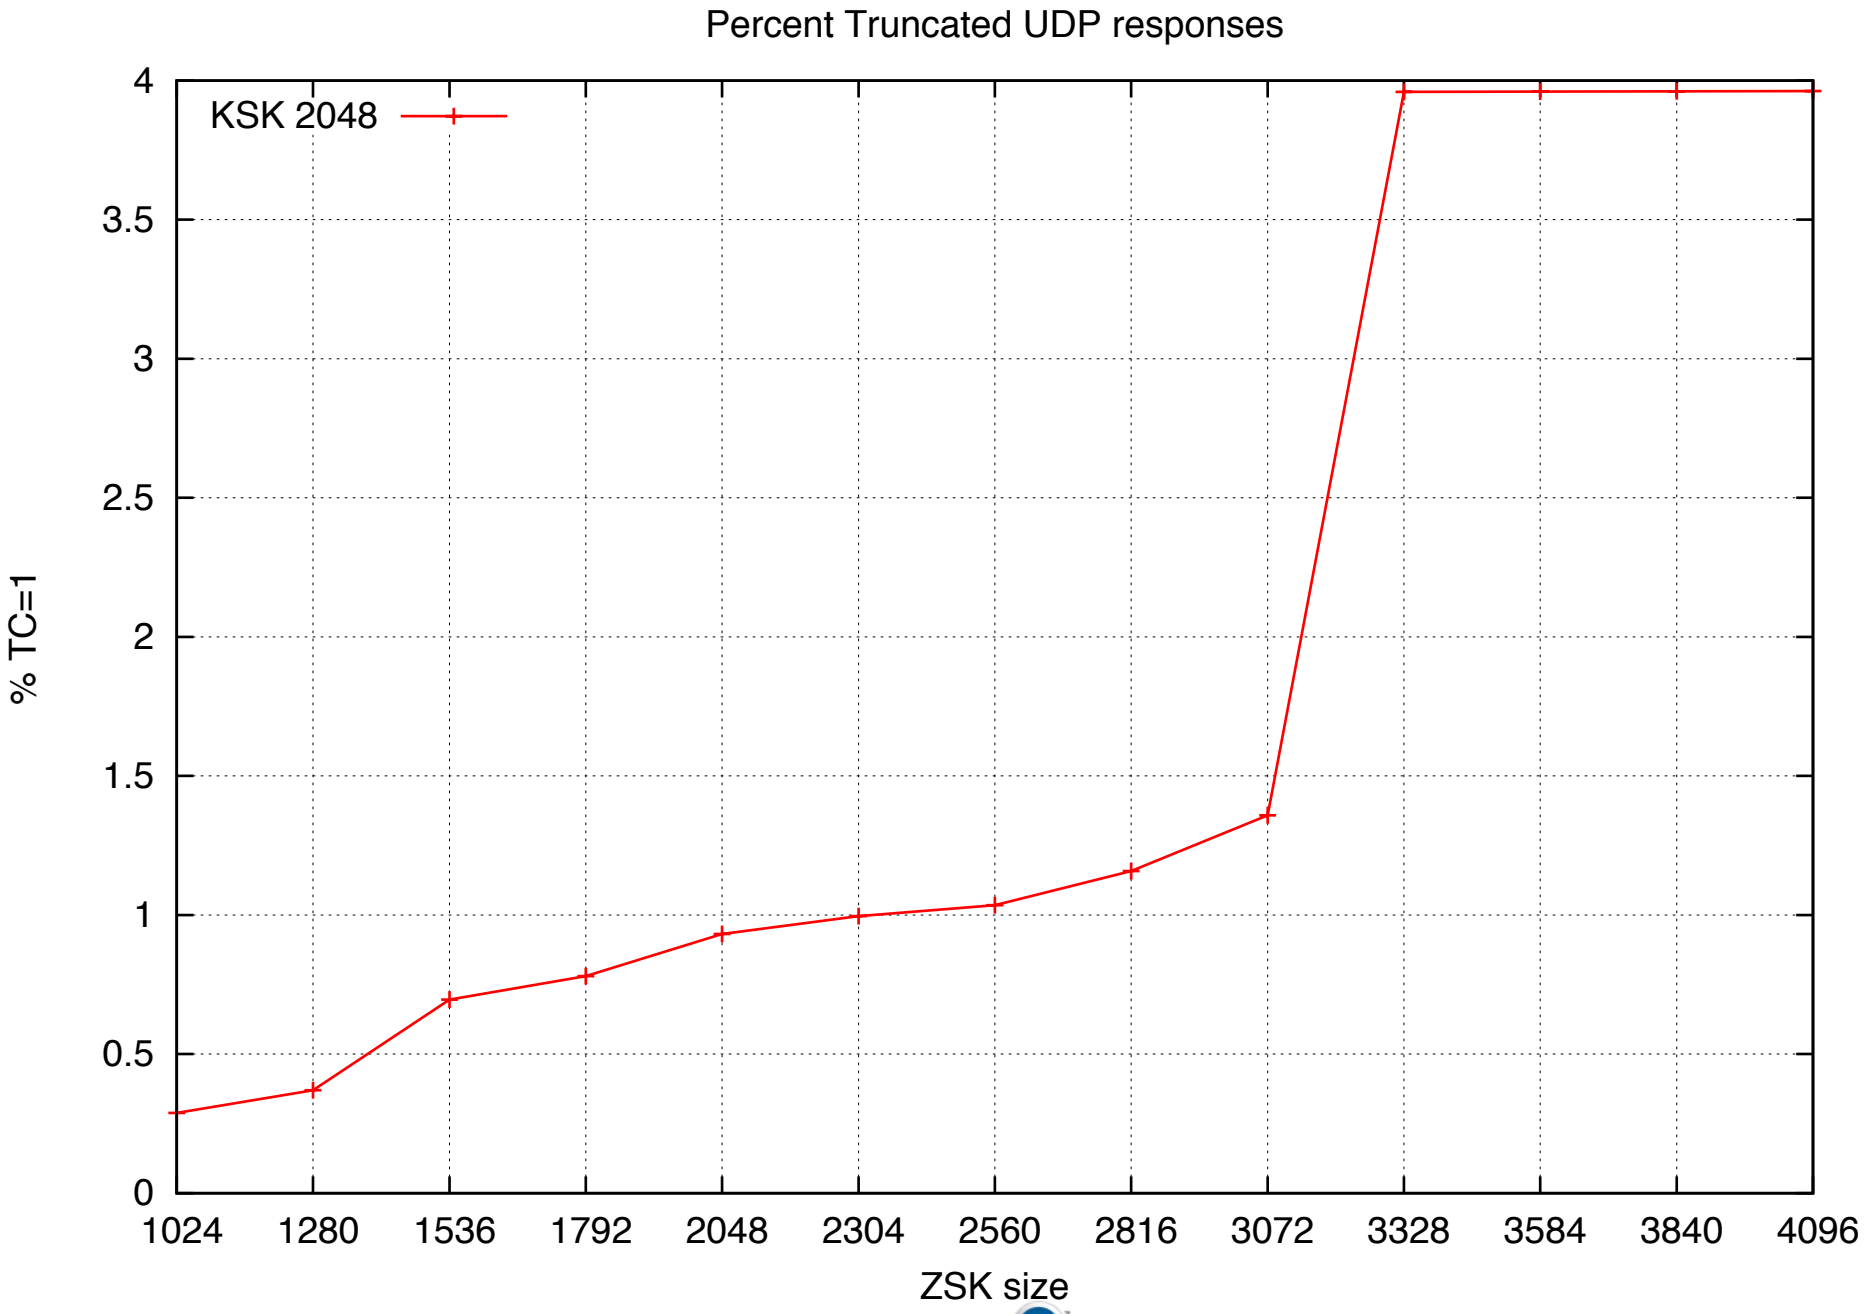
Percent Truncated UDP responses
KSK 2048
% TC=1
ZSK size
0
0.5
1
1.5
2
2.5
3
3.5
4
1024
1280
1536
1792
2048
2304
2560
2816
3072
3328
3584
3840
4096

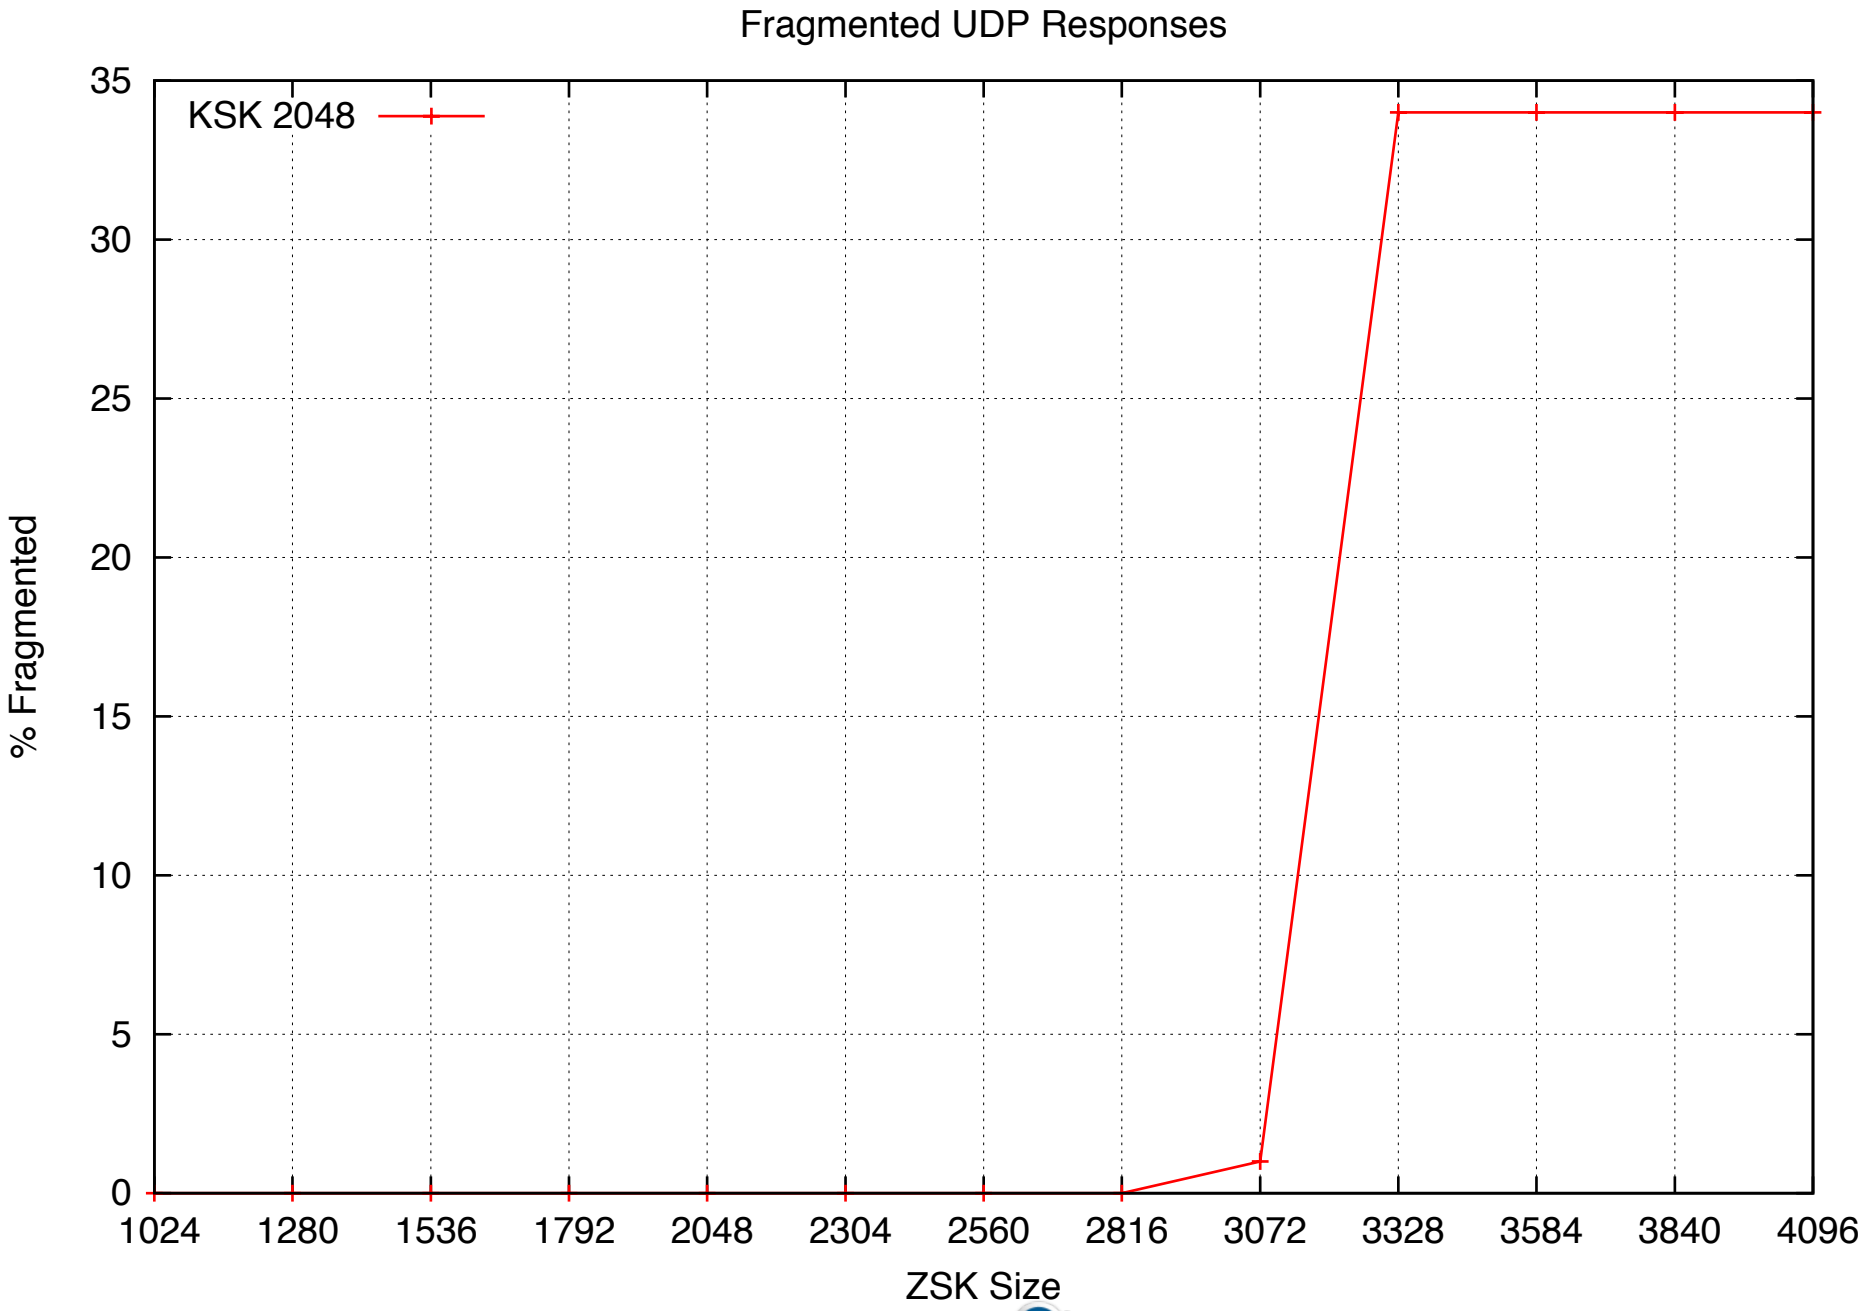
Fragmented UDP Responses
KSK 2048
% Fragmented
0
5
10
15
20
25
30
35
1024
1280
1536
1792
2048
2304
2560
2816
3072
3328
3584
3840
4096
ZSK Size

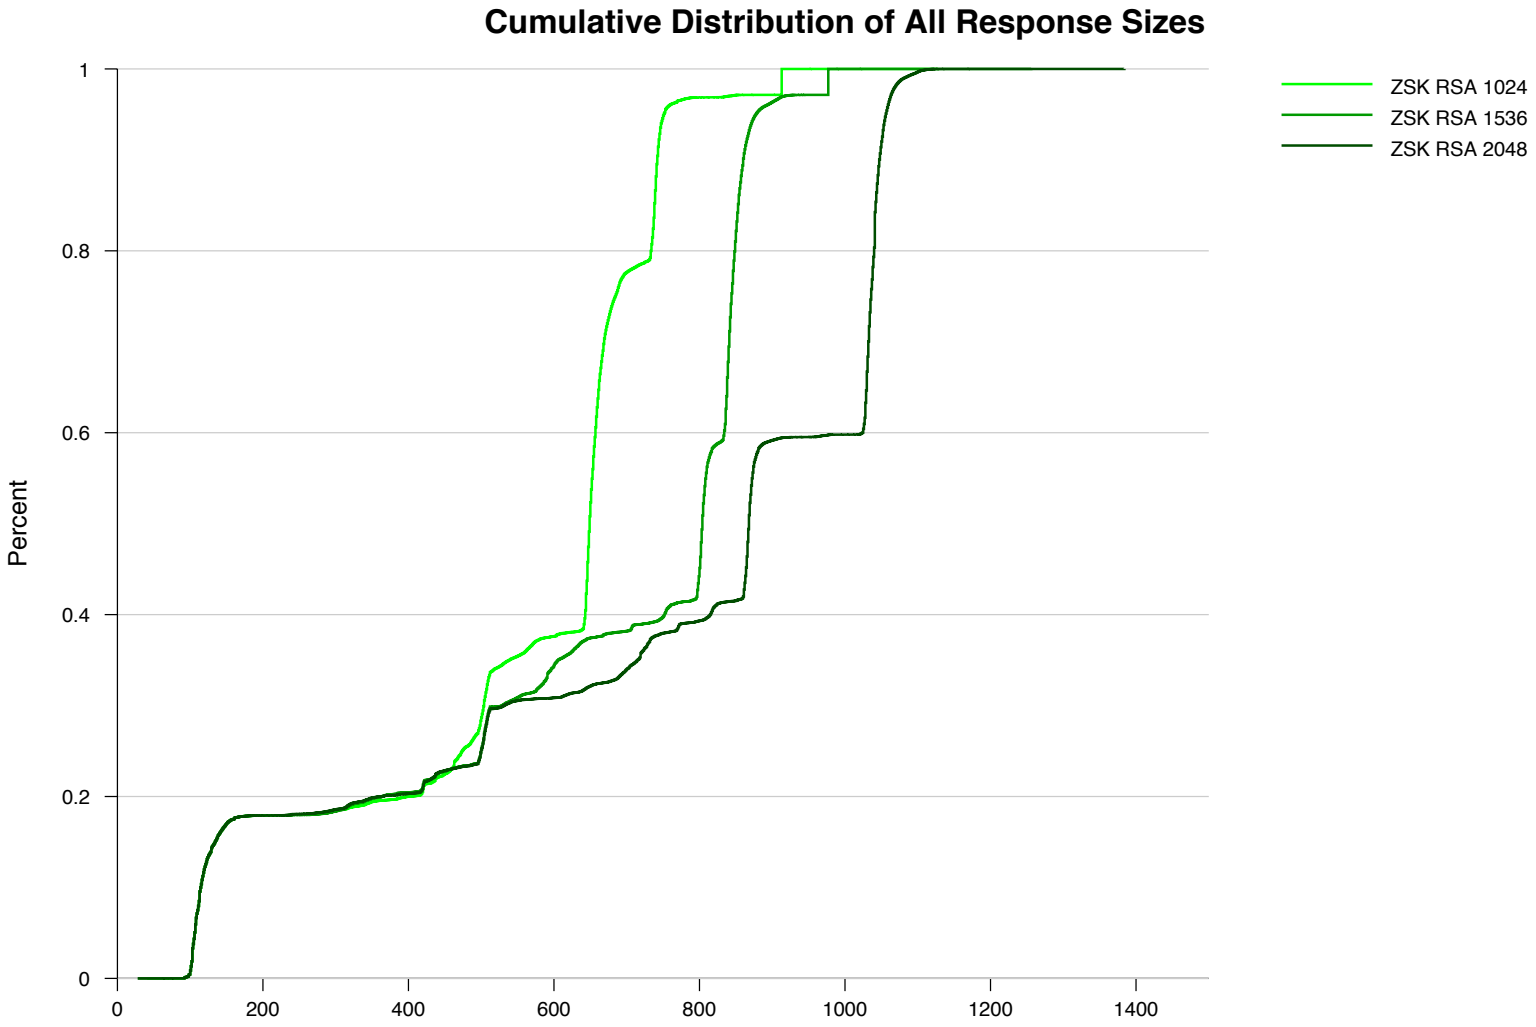
Cumulative Distribution of All Response Sizes
Percent
1
0.8
0.6
0.4
0.2
0
0
200
400
600
800
1000
1200
1400
ZSK RSA 1024
ZSK RSA 1536
ZSK RSA 2048

Cumulative Distribution of All Response Sizes

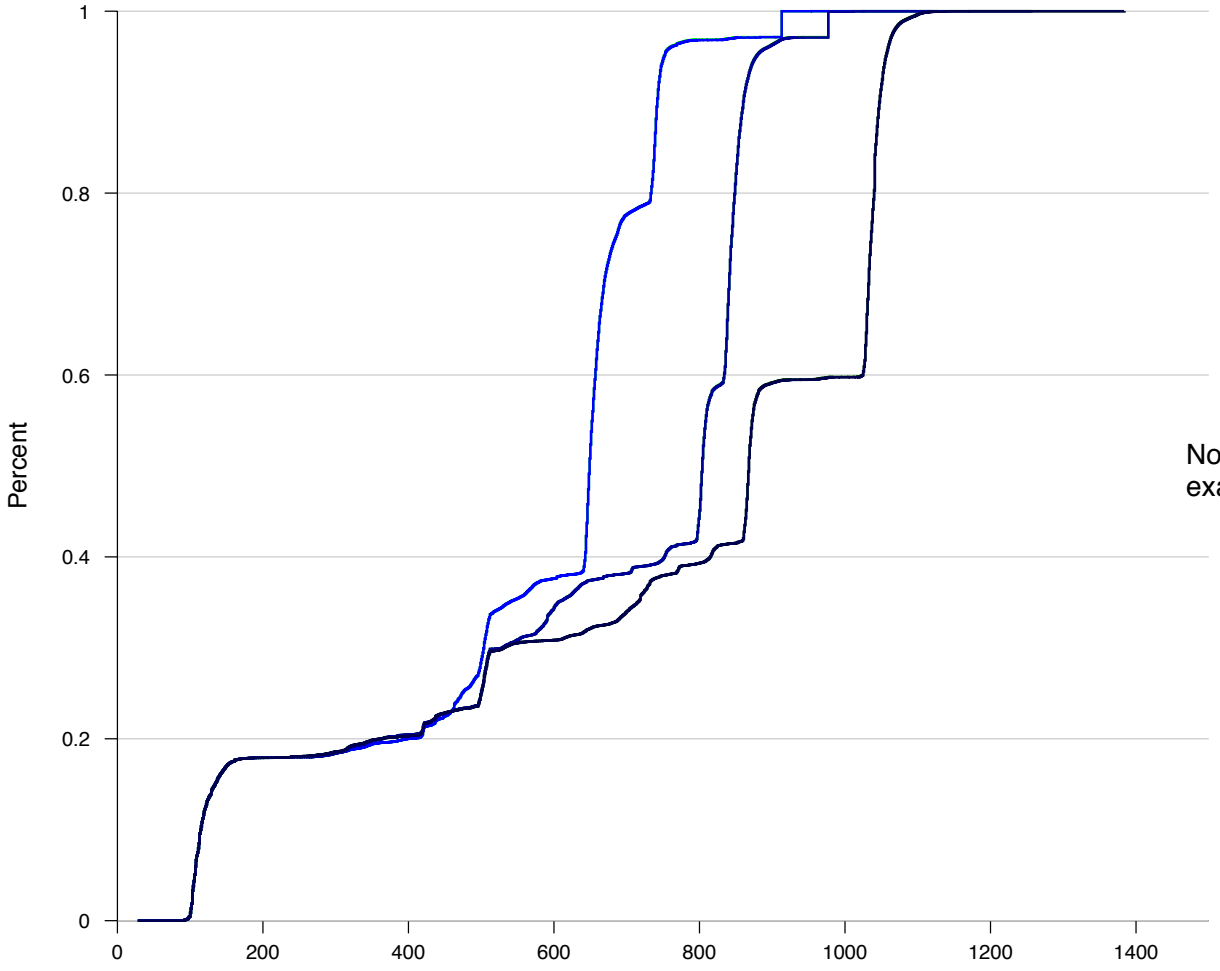

Note that ZSK Roll lines exactly match the non-Roll lines.

Cumulative Distribution of All Response Sizes

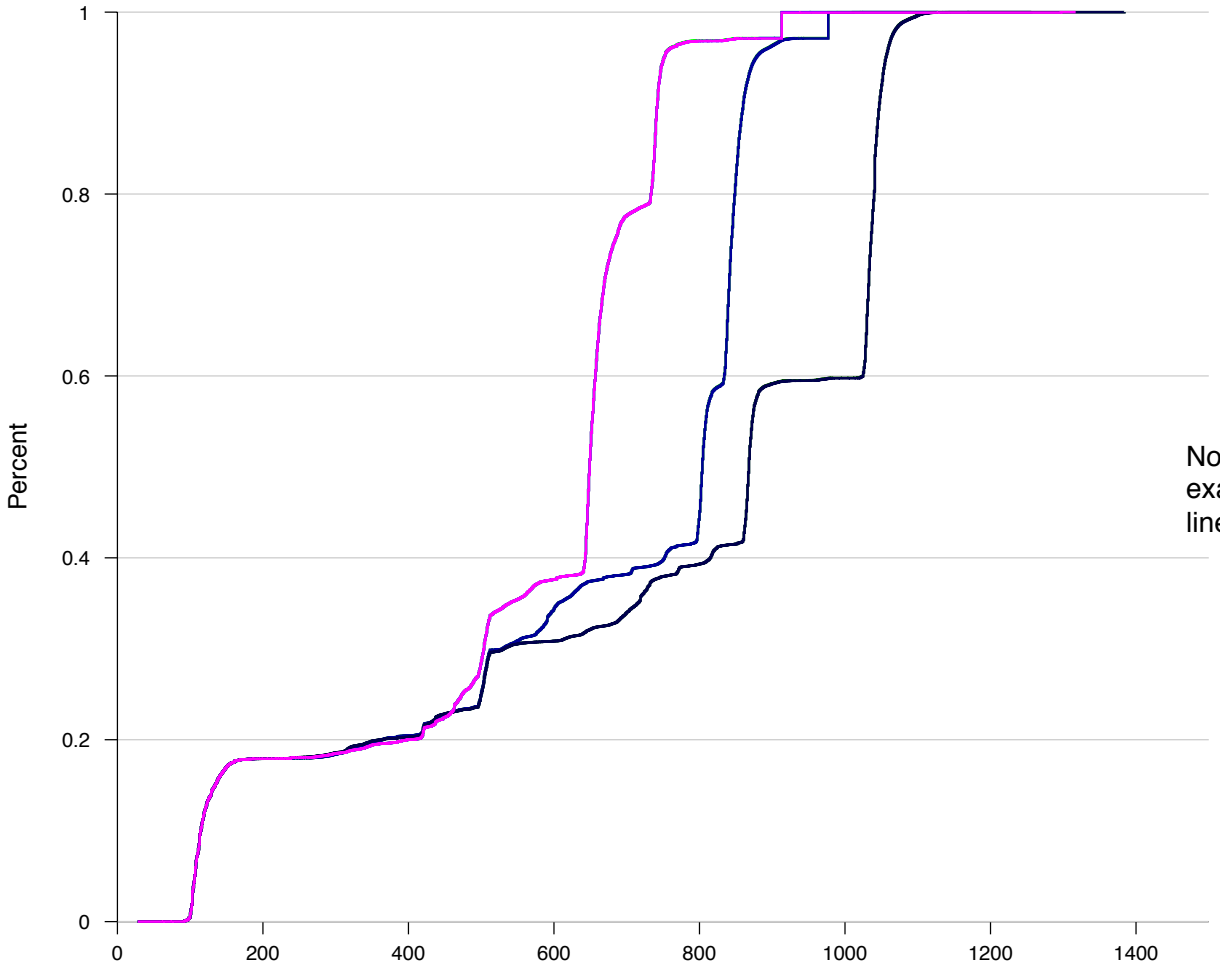

Note that KSK Roll RSA 2048 line exactly matches the other ZSK 1024 lines.

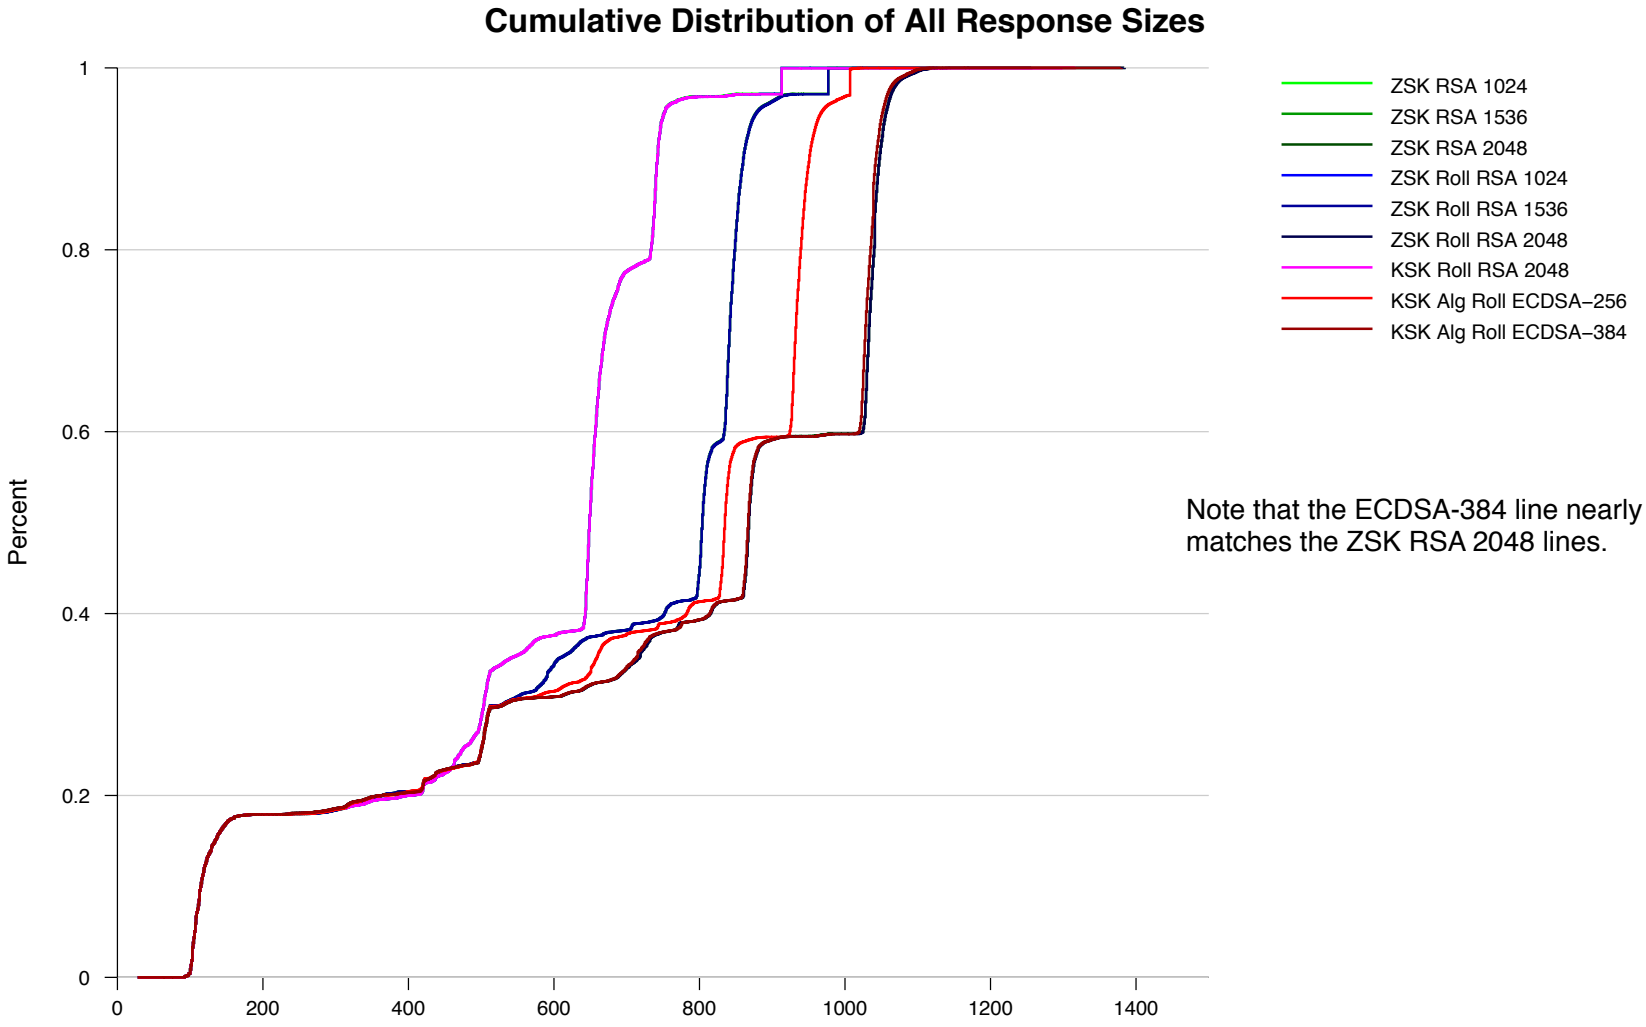

Note that the ECDSA-384 line nearly matches the ZSK RSA 2048 lines.

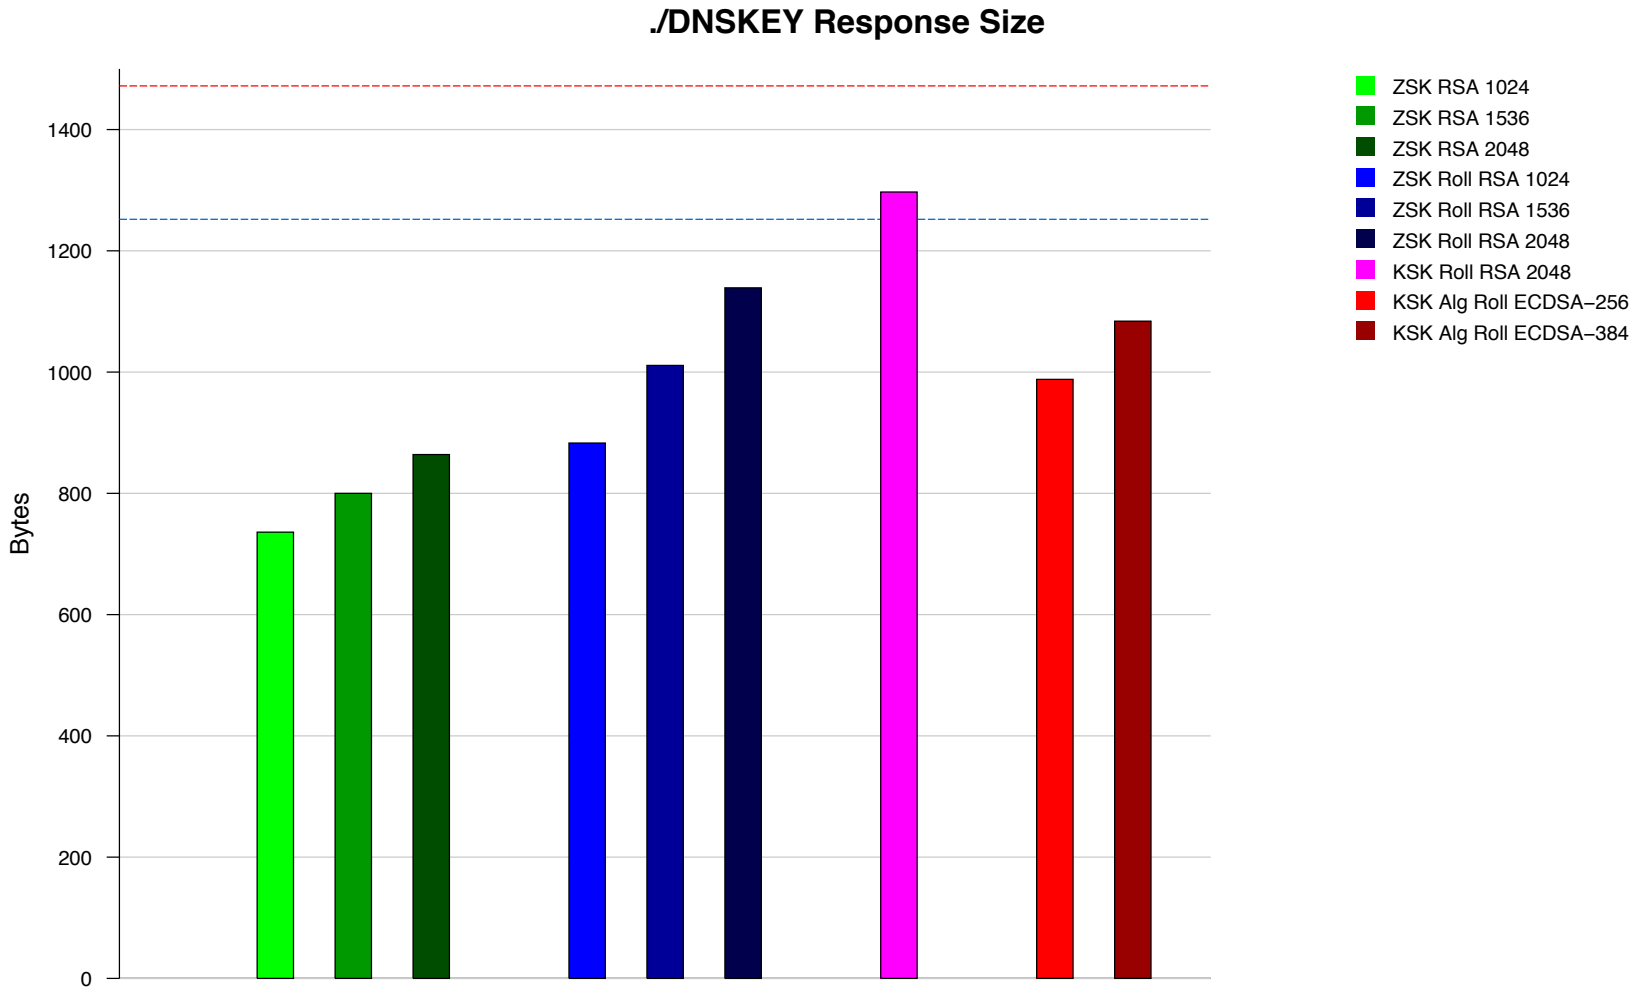
./DNSKEY Response Size
Bytes
0
200
400
600
800
1000
1200
1400
ZSK RSA 1024
ZSK RSA 1536
ZSK RSA 2048
ZSK Roll RSA 1024
ZSK Roll RSA 1536
ZSK Roll RSA 2048
KSK Roll RSA 2048
KSK Alg Roll ECDSA–256
KSK Alg Roll ECDSA–384

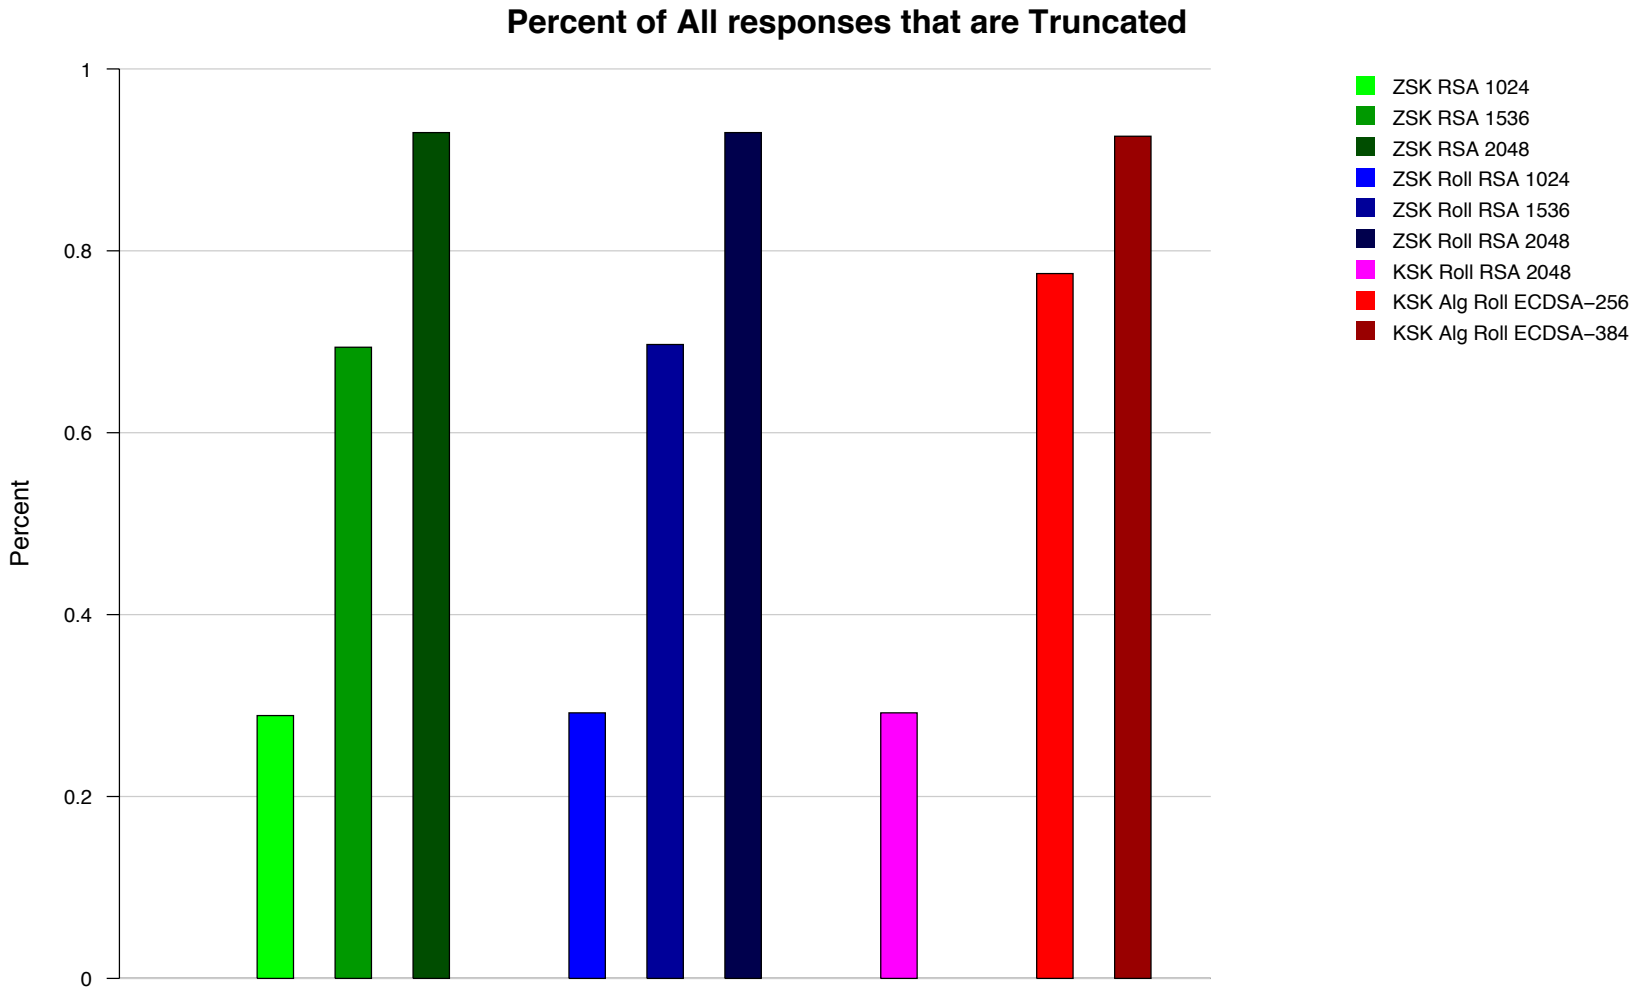
Percent of All responses that are Truncated
Percent
1
0.8
0.6
0.4
0.2
0
ZSK RSA 1024
ZSK RSA 1536
ZSK RSA 2048
ZSK Roll RSA 1024
ZSK Roll RSA 1536
ZSK Roll RSA 2048
KSK Roll RSA 2048
KSK Alg Roll ECDSA–256
KSK Alg Roll ECDSA–384

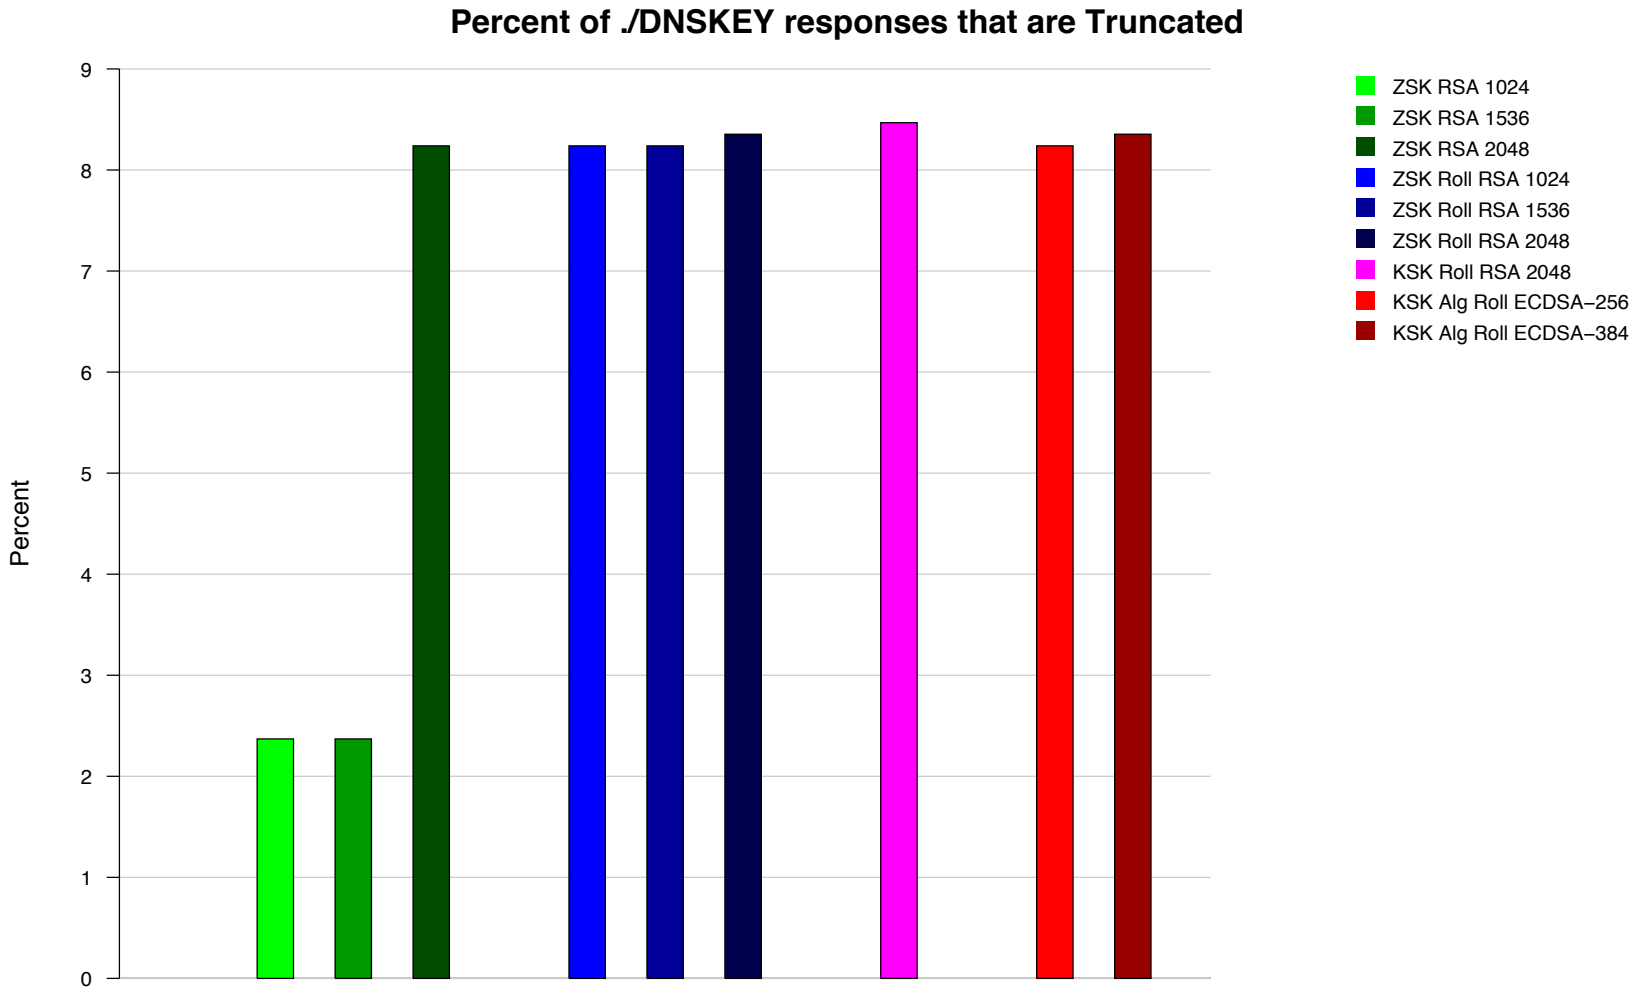
Percent of ./DNSKEY responses that are Truncated
Percent
0
1
2
3
4
5
6
7
8
9
ZSK RSA 1024
ZSK RSA 1536
ZSK RSA 2048
ZSK Roll RSA 1024
ZSK Roll RSA 1536
ZSK Roll RSA 2048
KSK Roll RSA 2048
KSK Alg Roll ECDSA–256
KSK Alg Roll ECDSA–384

Percent of All responses that are Fragmented

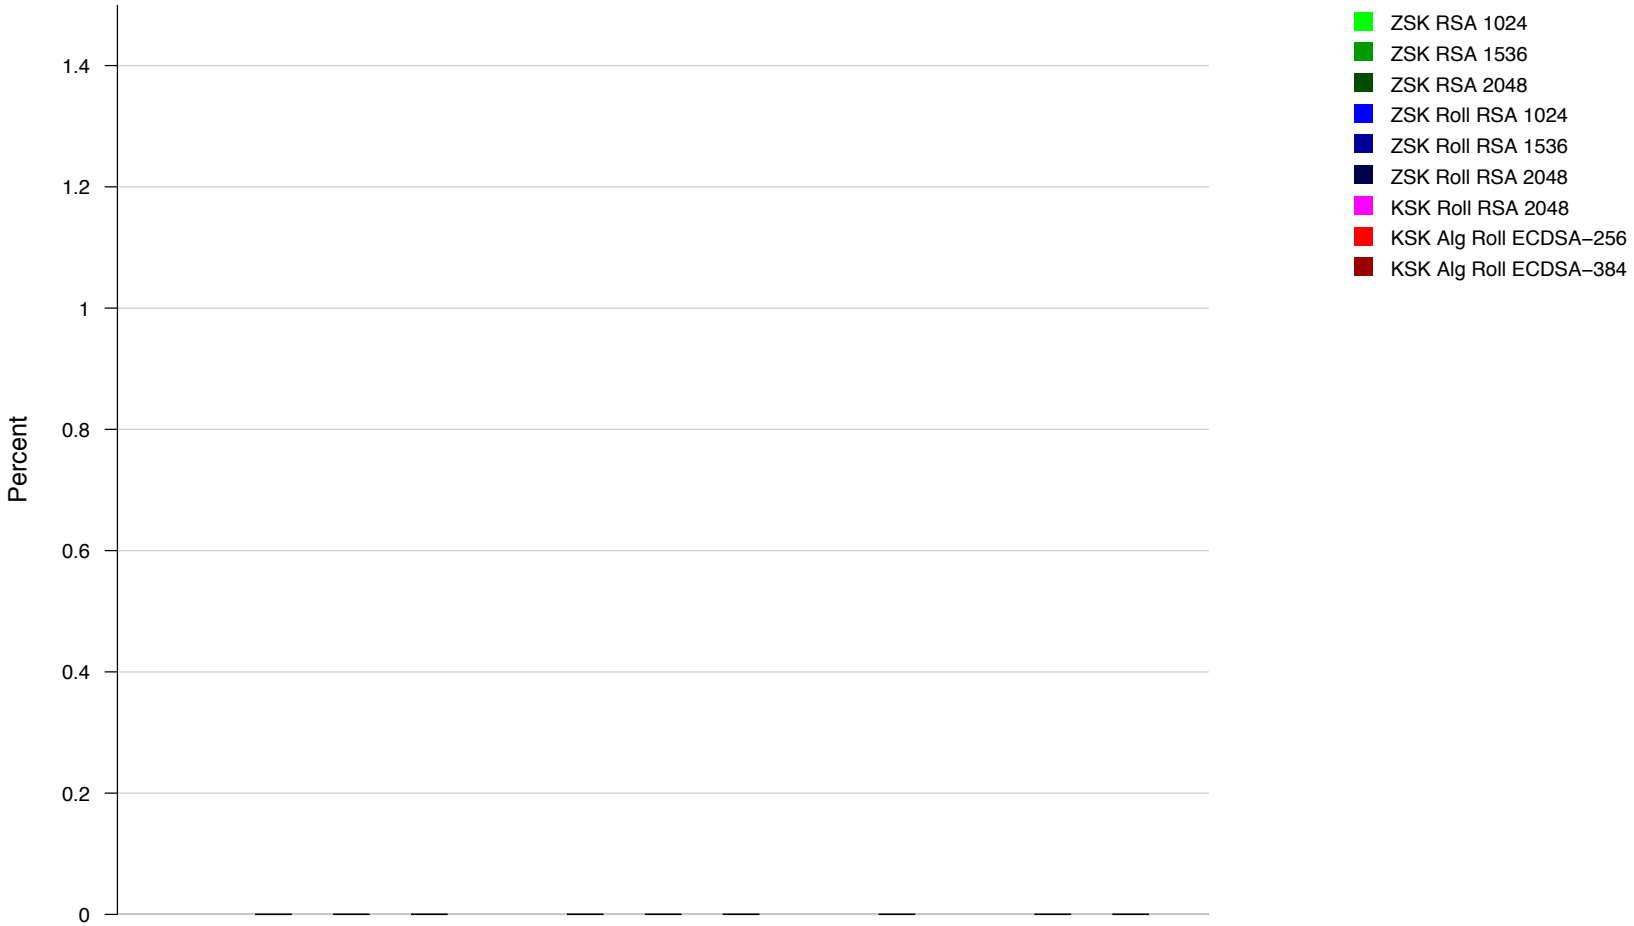

Percent of ./DNSKEY responses that are Fragmented

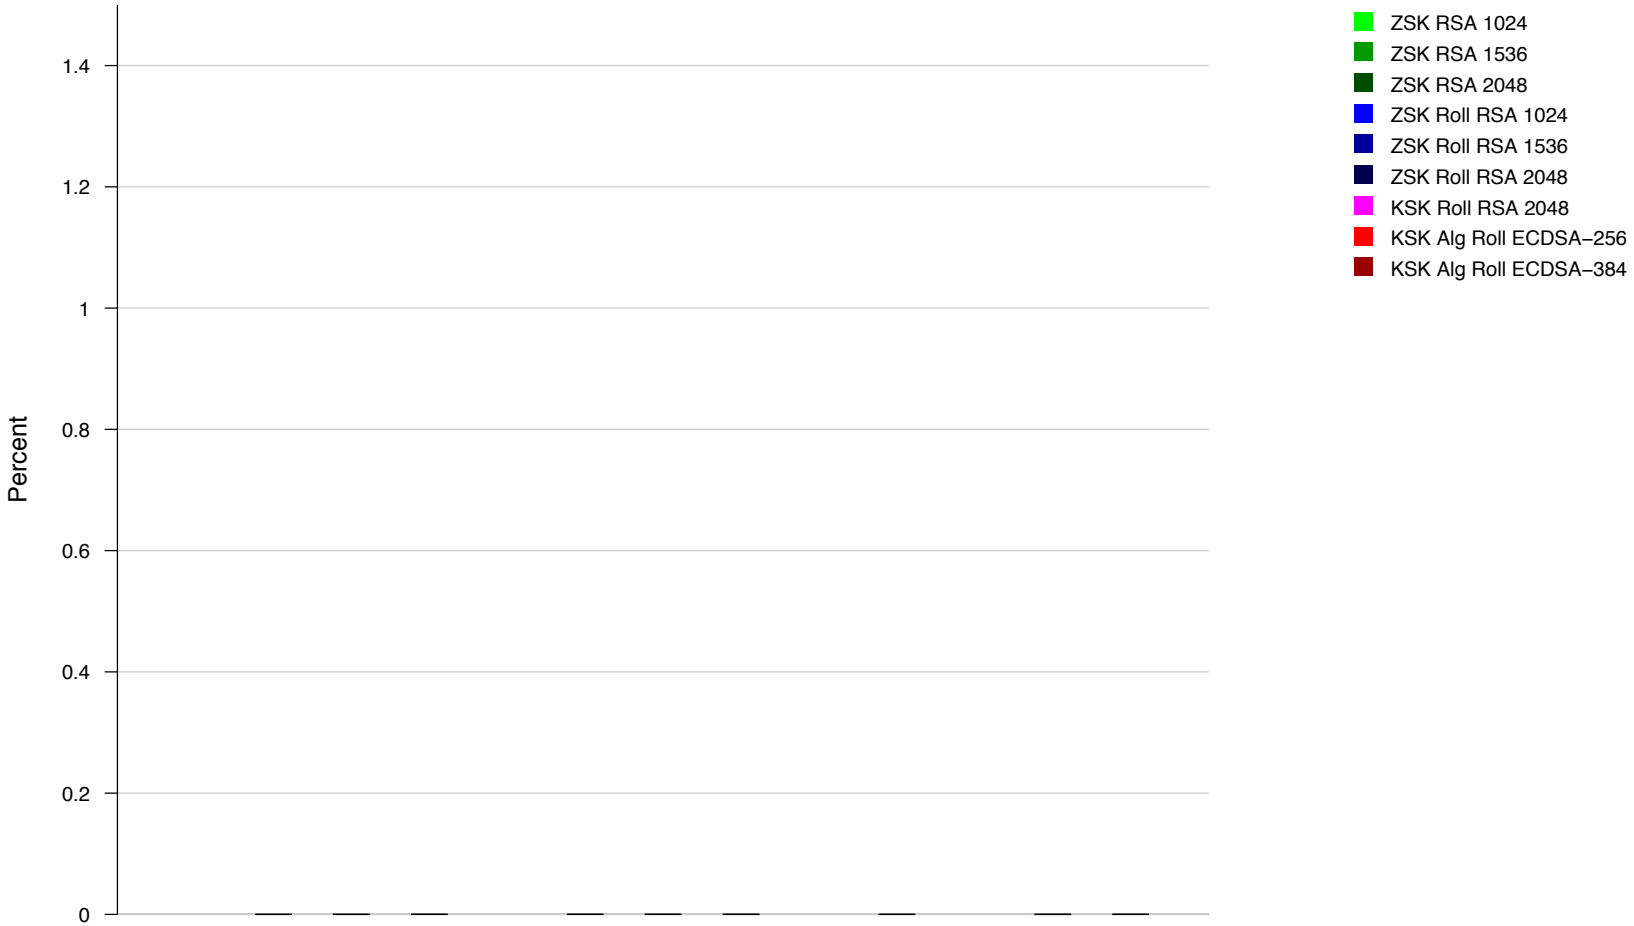

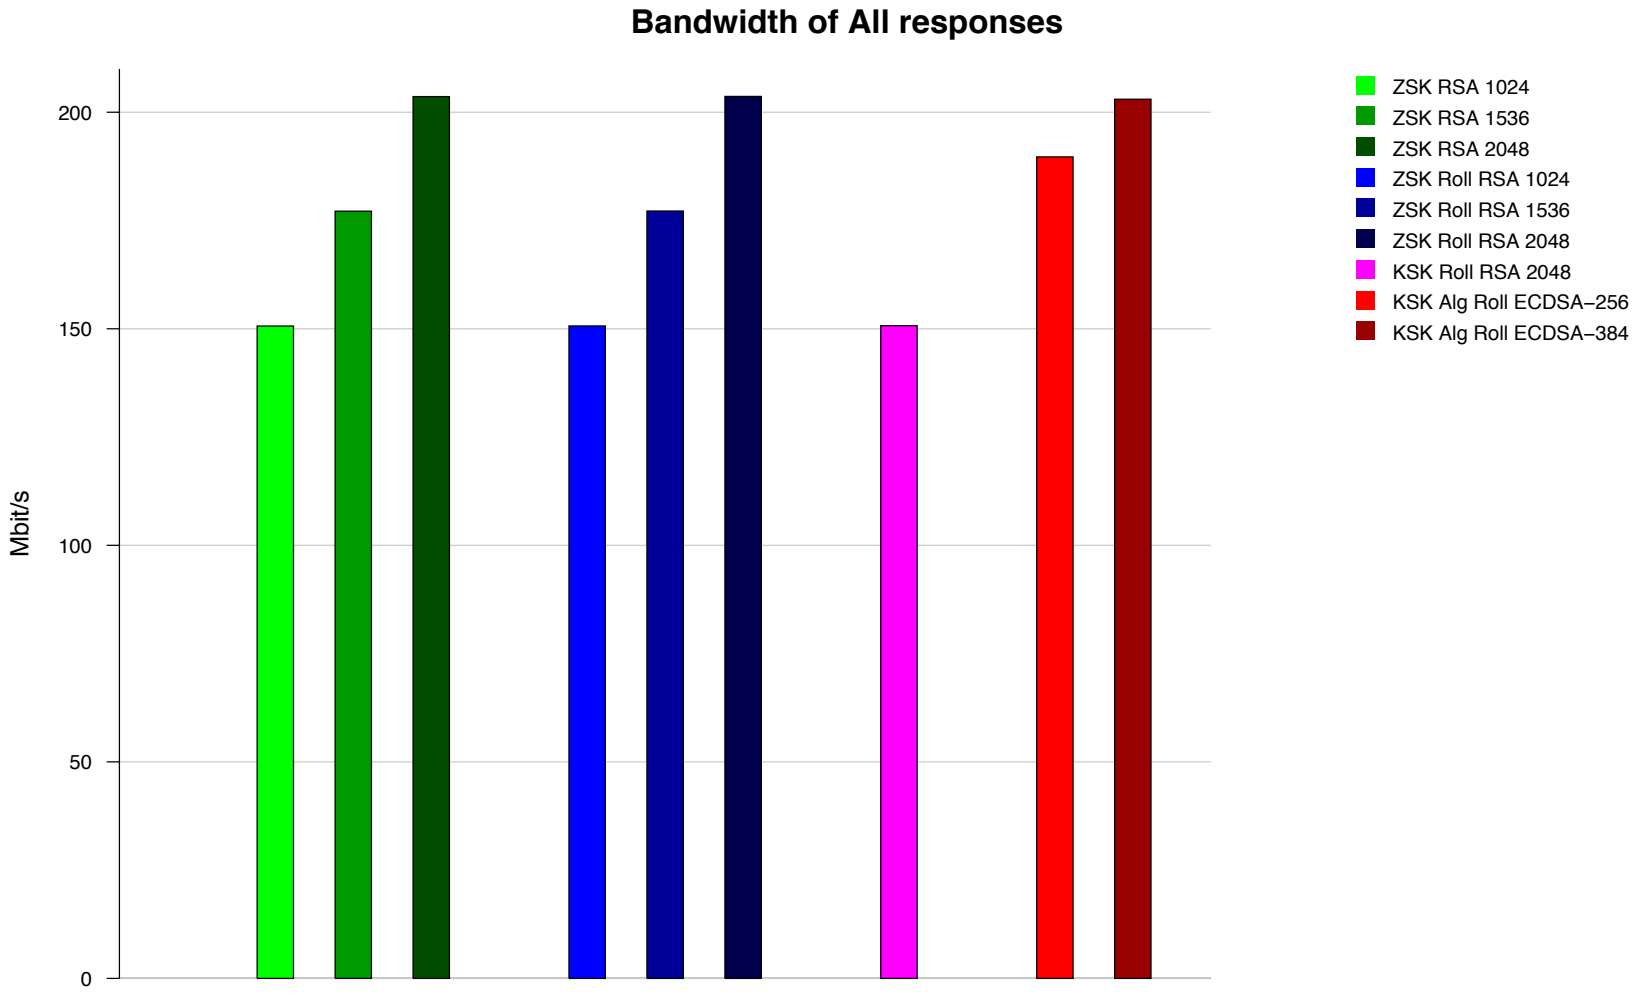
Bandwidth of All responses
Mbit/s
0
50
100
150
200
ZSK RSA 1024
ZSK RSA 1536
ZSK RSA 2048
ZSK Roll RSA 1024
ZSK Roll RSA 1536
ZSK Roll RSA 2048
KSK Roll RSA 2048
KSK Alg Roll ECDSA–256
KSK Alg Roll ECDSA–384

Summary

- Scenarios simulated here indicate:
- Modest increases in truncation (leading to TCP)
- No UDP fragmentation at 1500 byte MTU
- Up to 35% Increase in root server response bandwidth

Questions?

VERISIGN®