

Network Path Problems in DNSSEC's Deployment

Eric Osterweil

Dan Massey

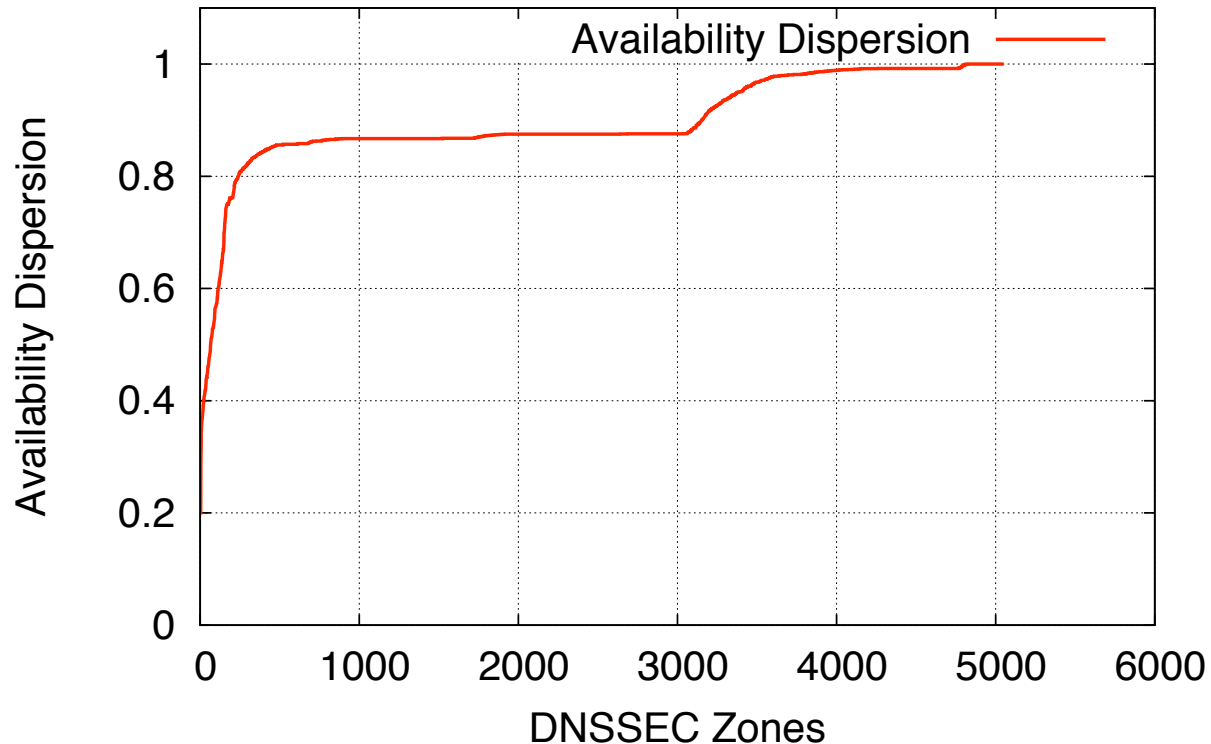
Lixia Zhang

The Network Path and PMTU

- A network path is a sequence of links
 - Each link can only support packets of a certain size (MTU)
- The smallest MTU for a network path is its bottleneck (PMTU)
- Further limited by “middle boxes” (firewalls, NAT, etc.)
 - We overload the term PMTU to apply in these cases too
- Network paths that *do* support large packets may fail to deliver large DNS messages

More Succinct

Availability Dispersion of DNSSEC Zones



- We use a metric (defined in prior work) to quantify the “availability dispersion” of each zone
 - Captures how different each poller’s view of each zone is
- Using a weighted average over time, we see that most zones have suffered dispersion

“dnstunnel”

Take a Look Yourself

```
208.67.177.8 4096B * * *
208.67.177.8 2304B * * *
208.67.177.8 1408B * * *
208.67.177.8 960B 0,179109 (truncated)
208.67.177.8 1184B 0,175534 (truncated)
208.67.177.8 1296B * * *
208.67.177.8 1240B 0,178052 (truncated)
208.67.177.8 1268B * * *
208.67.177.8 1254B * * *
208.67.177.8 1247B 0,175454 (truncated)
208.67.177.8 1250B 0,176776 (truncated)
208.67.177.8 1252B 0,175576 (truncated)
208.67.177.8 1253B * * *
208.67.177.8 1252B 0,175913 (truncated)
208.67.177.8 1253B * * *
208.67.177.8 1252B 0,176101 (truncated)
```

PMTU walking summary:

```
=====
Name Server      Keys   Small  Largest  Optimal
IP               fit?   Buffer  Buffer    Buffer
-----
208.67.177.138  yes    2225   4096     2977
208.67.177.7   no     1252   0         0
208.67.177.139 yes    2225   4096     2977
208.67.177.8   no     1252   0         0
```

dnstunnel questsys.com.

The Nature of the Problem

- There is no silver bullet
 - This is an end-middle-end problem
- Network paths have limitations at *various* points
 - Not always just the sender or receiver side
 - If problem is at the sender/receiver, test and upgrade
- Resolvers specify buffer sizes that exceed PMTUs
 - EDNS0 buffer size \neq PMTU size
 - More intelligent discovery and use of buffer size
- Name servers treat buffer sizes as path-capacity
 - How should they know better?
 - Where possible, reduce set sizes

What to do Tactically

- Download our toolkit that includes dnsfunnel:
<http://vantage-points.org/>



Vantages



- Check your zones' availability at:



<http://secspider.cs.ucla.edu/>

SecSpider the DNSSEC Monitoring Project



[Home](#) | [Blog](#) | [About](#) | [FAQ](#) | [Documentation](#) | [Usage](#) | [Pollers](#) | [GPG Key](#) | [IRL](#)

- ☀ Check out our [BIND formatted Trust Anchors File!](#)
- ☀ Check out our [new blog entry](#) about our enhanced zone drilldown pages: featuring zones' availability (PMTU failure reports) and counts of state RRsets!

Search for zone:

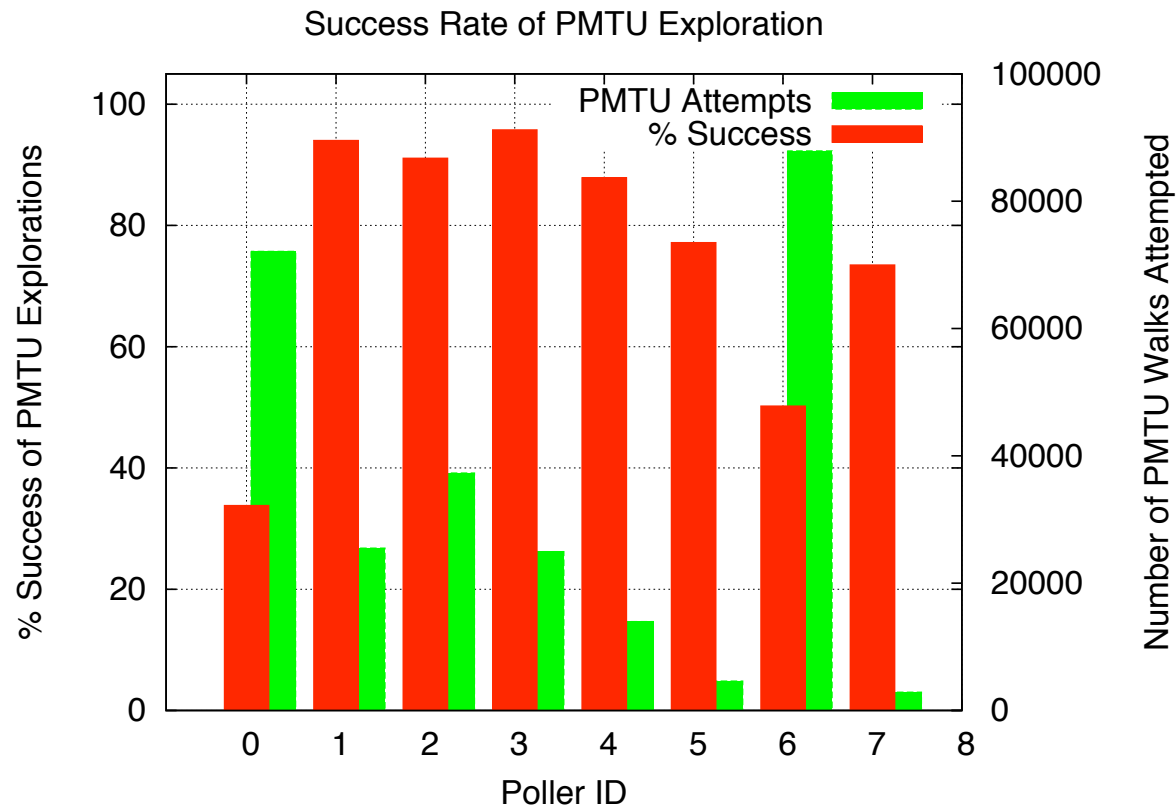
Zone:

To add a zone for monitoring, please submit below:

Thanks

Backup

As Seen From All of Our Pollers



- Green bars indicate the number of times a poller needed to do a PMTU walk
- Red bars indicate the percentage of times a PMTU was able to find a buffer size the allowed DNSKEYs to be received

“dnstunnel”

Take a Look Yourself

```
208.67.177.8 4096B * * *
208.67.177.8 2304B * * *
208.67.177.8 1408B * * *
208.67.177.8 960B 0,179109 (truncated)
208.67.177.8 1184B 0,175534 (truncated)
208.67.177.8 1296B * * *
208.67.177.8 1240B 0,178052 (truncated)
208.67.177.8 1268B * * *
208.67.177.8 1254B * * *
208.67.177.8 1247B 0,175454 (truncated)
208.67.177.8 1250B 0,176776 (truncated)
208.67.177.8 1252B 0,175576 (truncated)
208.67.177.8 1253B * * *
208.67.177.8 1252B 0,175913 (truncated)
208.67.177.8 1253B * * *
208.67.177.8 1252B 0,176101 (truncated)
-----
PMTU walking summary:
=====
Name Server  Keys  Small  Largest  Optimal
IP           fit?  Buffer  Buffer   Buffer
-----
208.67.177.138 yes    2225   4096    2977
208.67.177.7  no     1252    0        0
208.67.177.139 yes    2225   4096    2977
208.67.177.8  no     1252    0        0
```

```
192.134.0.49 4096B 0,033097
-----
192.36.125.2 4096B 0,004157
192.36.125.2 2304B 0,002445
192.36.125.2 1408B 0,002923 (truncated)
192.36.125.2 1856B 0,002845
192.36.125.2 1632B 0,002304
192.36.125.2 1520B 0,002108
192.36.125.2 1464B 0,002897
192.36.125.2 1436B 0,003885
192.36.125.2 1422B 0,005049 (truncated)
192.36.125.2 1429B 0,001806 (truncated)
192.36.125.2 1432B 0,002284
192.36.125.2 1430B 0,002039
192.36.125.2 1429B 0,001743 (truncated)
192.36.125.2 4096B 0,002577
-----
PMTU walking summary:
=====
Name Server  Keys  Small  Largest  Optimal
IP           fit?  Buffer  Buffer   Buffer
-----
192.5.4.1    yes    1430   4096    1854
193.0.0.195 yes    1438   4096    1862
192.134.0.49 yes    1430   4096    2320
192.36.125.2 yes    1430   4096    2320
```

- dnstunnel questsys.com.

- dnstunnel ripe.net