

IPv6 Extension Headers Again!

Geoff Huston, Joao Damas


APNIC Labs

RFC 7872 - 2016

- Sending crafted IPv6 packets within a TCP session to web servers and measuring the success rate through receiving a TCP ACK from the server2


Observed Drop Rate

	DST	HBH	FRAG
World IPv6 Day Data Set	11.88%	40.70%	30.51%
Alexa Top 1M	10.91%	45.45%	28.26%



APNIC's Test Rig

- We use an ad campaign to direct browsers to retrieve a URL from our server
- The server sits behind a front end that performs packet mangling (and TCP MSS wrangling)
- The particular packet transform is selected at random for each TCP session

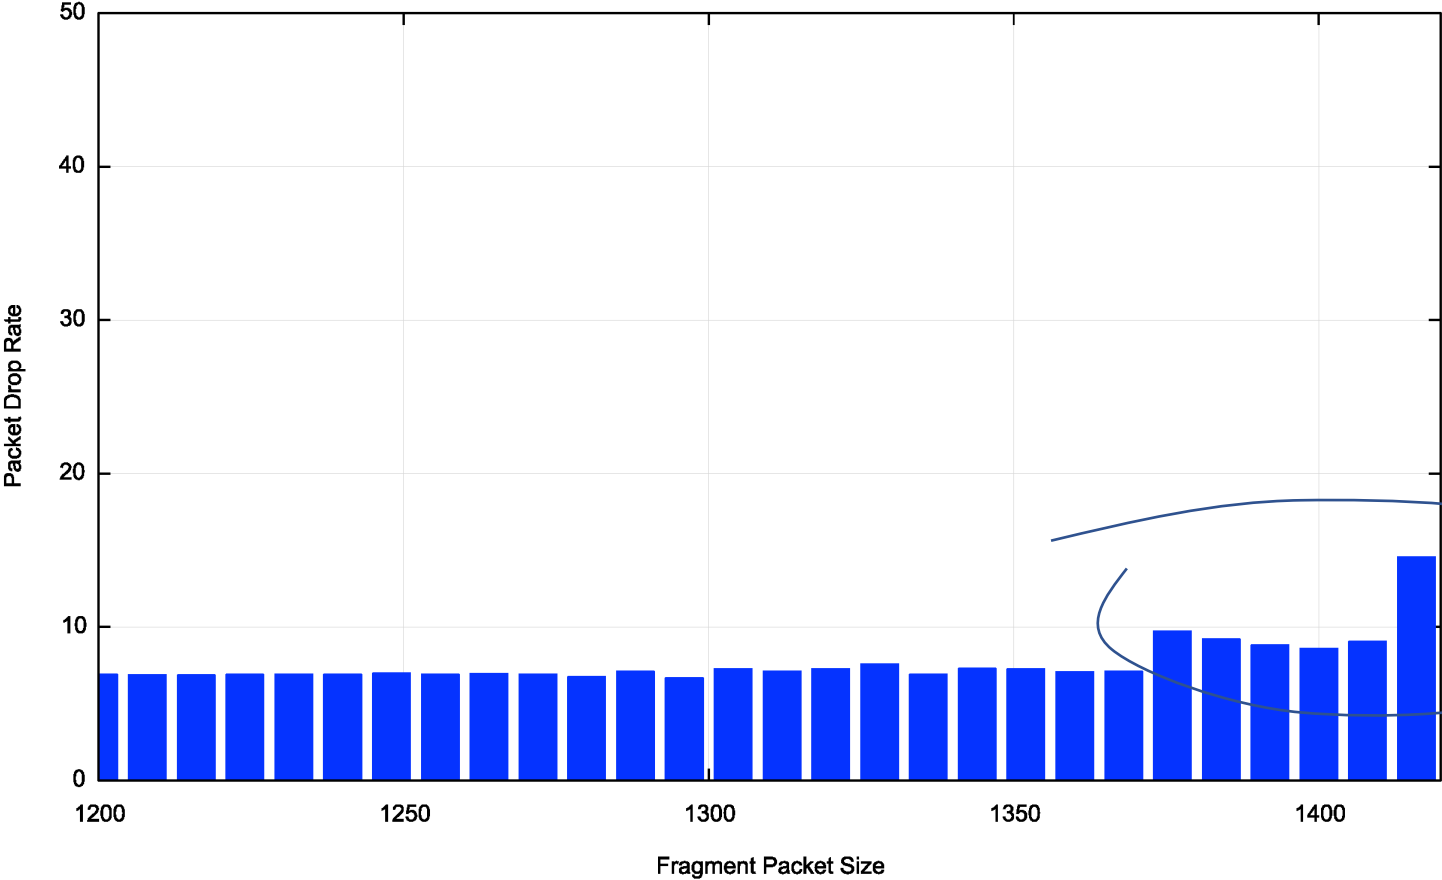


Tests

Random selection of:

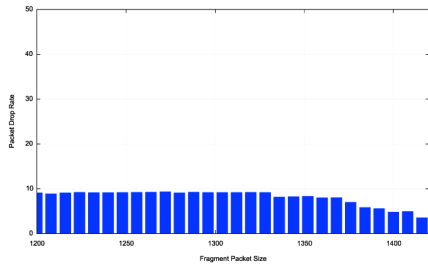
- Fragmentation
 - Vary the size of the initial fragment: 1200, 1208, 1216, ... 1416 octets
- DST Header
 - Use 0x1E type code, filled with 0's
 - Size 8, 16, 32, 64, 128 octets
- HBH header
 - Use 0x1E type code, filled with 0's
 - Size 8, 16, 32, 64, 128 octets

Fragmentation Drop Rate

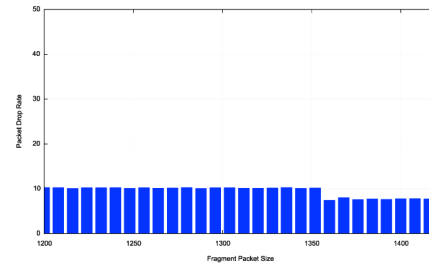


Huh?


Fragmentation around the world




Fragmentation Packet Drop Profile for Europe




Fragmentation Packet Drop Profile for North America




Fragmentation Packet Drop Profile for South America



Fragmentation Packet Drop Profile for Asia



Fragmentation Packet Drop Profile for India



Fragmentation Packet Drop Profile for China

In Europe and North America the drop rate DECREASES for large sized initial frag packets

In South America there is a large drop signal at 1416 bytes


In Asia the drop rate INCREASES at 1380 bytes

India has a small drop rate up to 1380 bytes

China has a large drop rate for all sizes, but it shifts up at 1316 bytes

Fragmentation or EH


- Let's also include an "Atomic Fragment"



Overall Frag Drop rate


Atomic Frag Drop rate

Country variations



United States

China



Why?

- <insert favourite theory here>

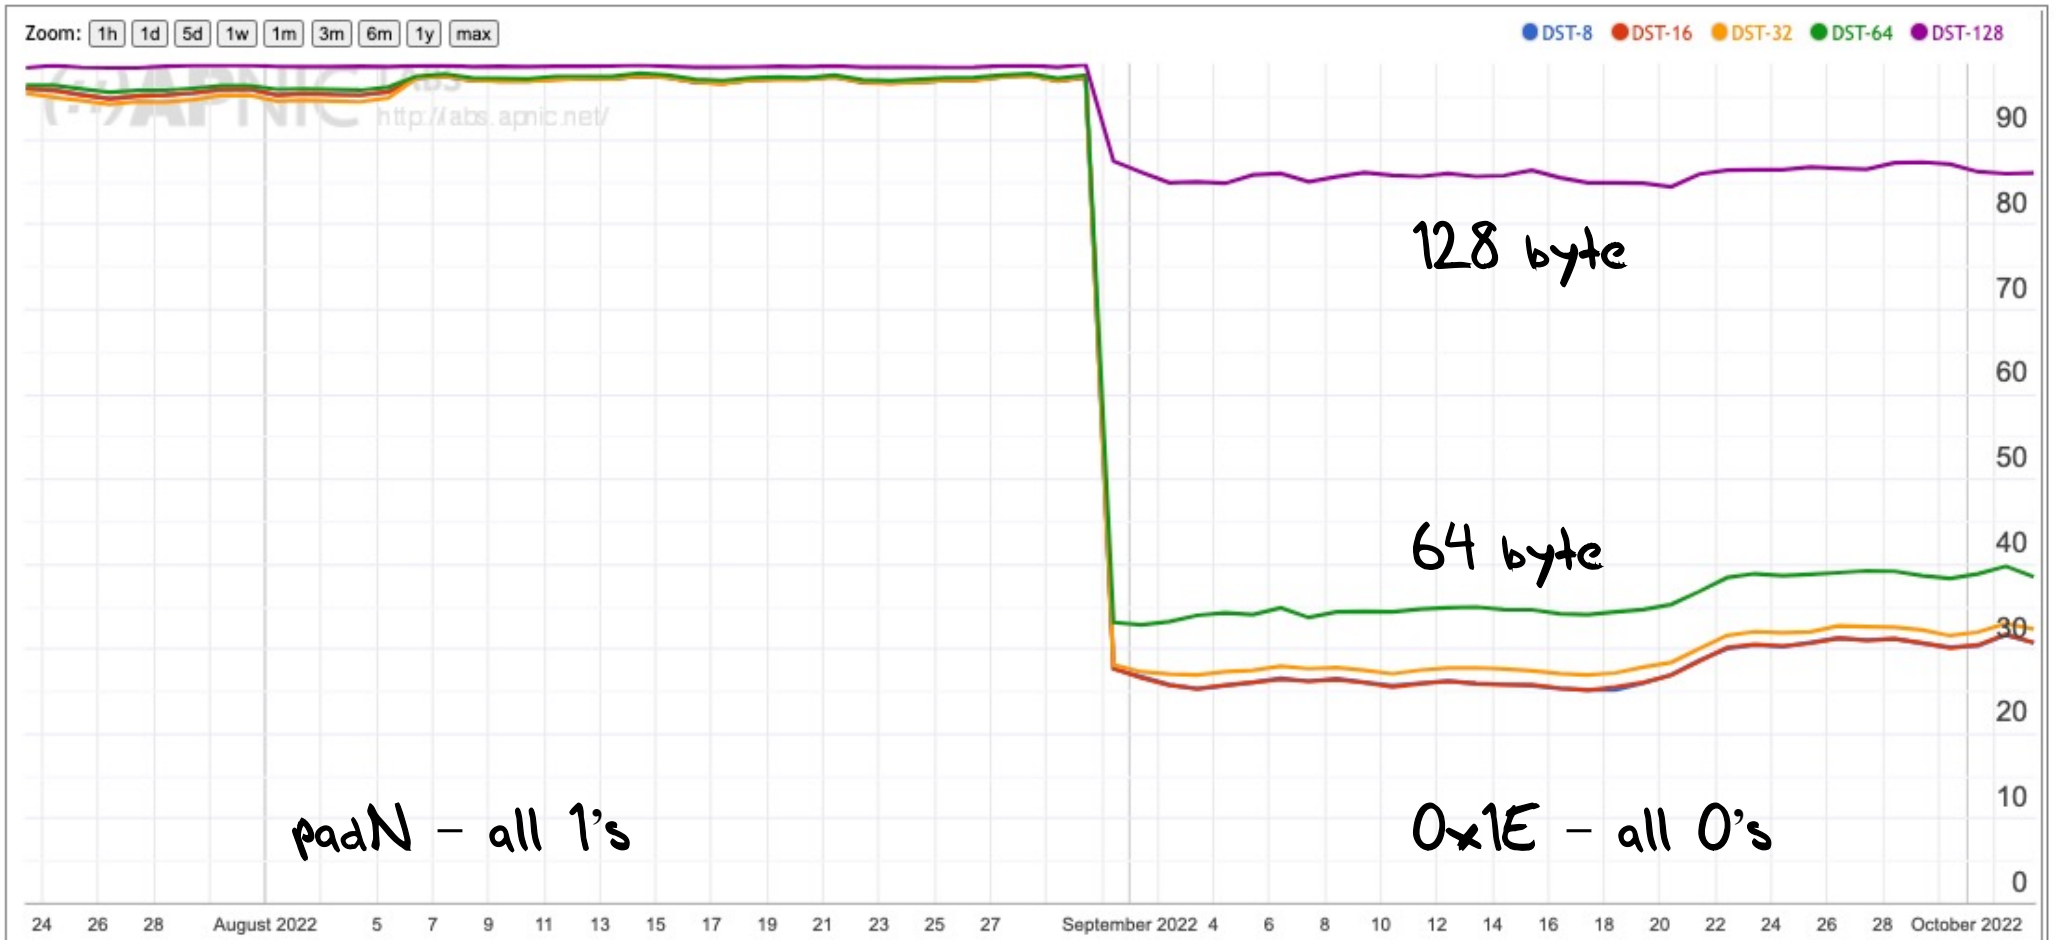
DST EH Drop Rate

DST EH Drop Rate

Linux Filter

```
if (nh[off] == IPV6_TLV_PADN) {
    /* RFC 2460 states that the purpose of PadN is
     * to align the containing header to multiples
     * of 8. 7 is therefore the highest valid value.
     * See also RFC 4942, Section 2.1.9.5.
     */
    padlen += optlen;
    if (padlen > 7)
        goto bad;
    /* RFC 4942 recommends receiving hosts to
     * actively check PadN payload to contain
     * only zeroes.
     */
    for (i = 2; i < optlen; i++) {
        if (nh[off + i] != 0)
            goto bad;
    }
}
```

DST EH drop rate

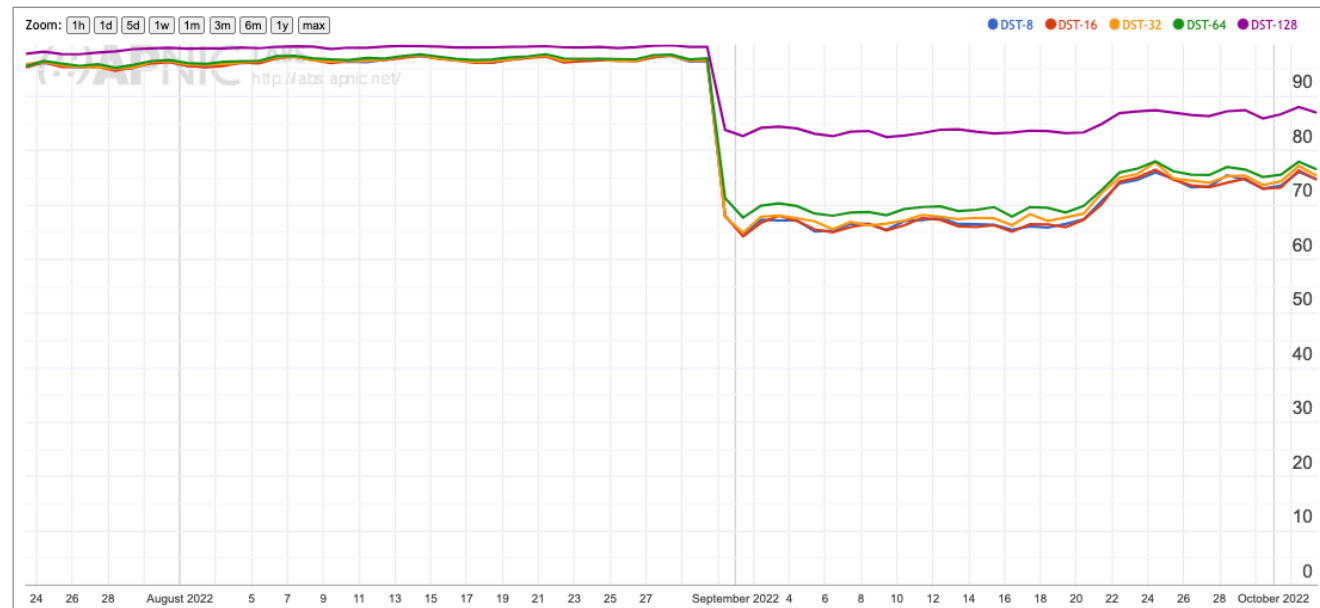


Country Variations

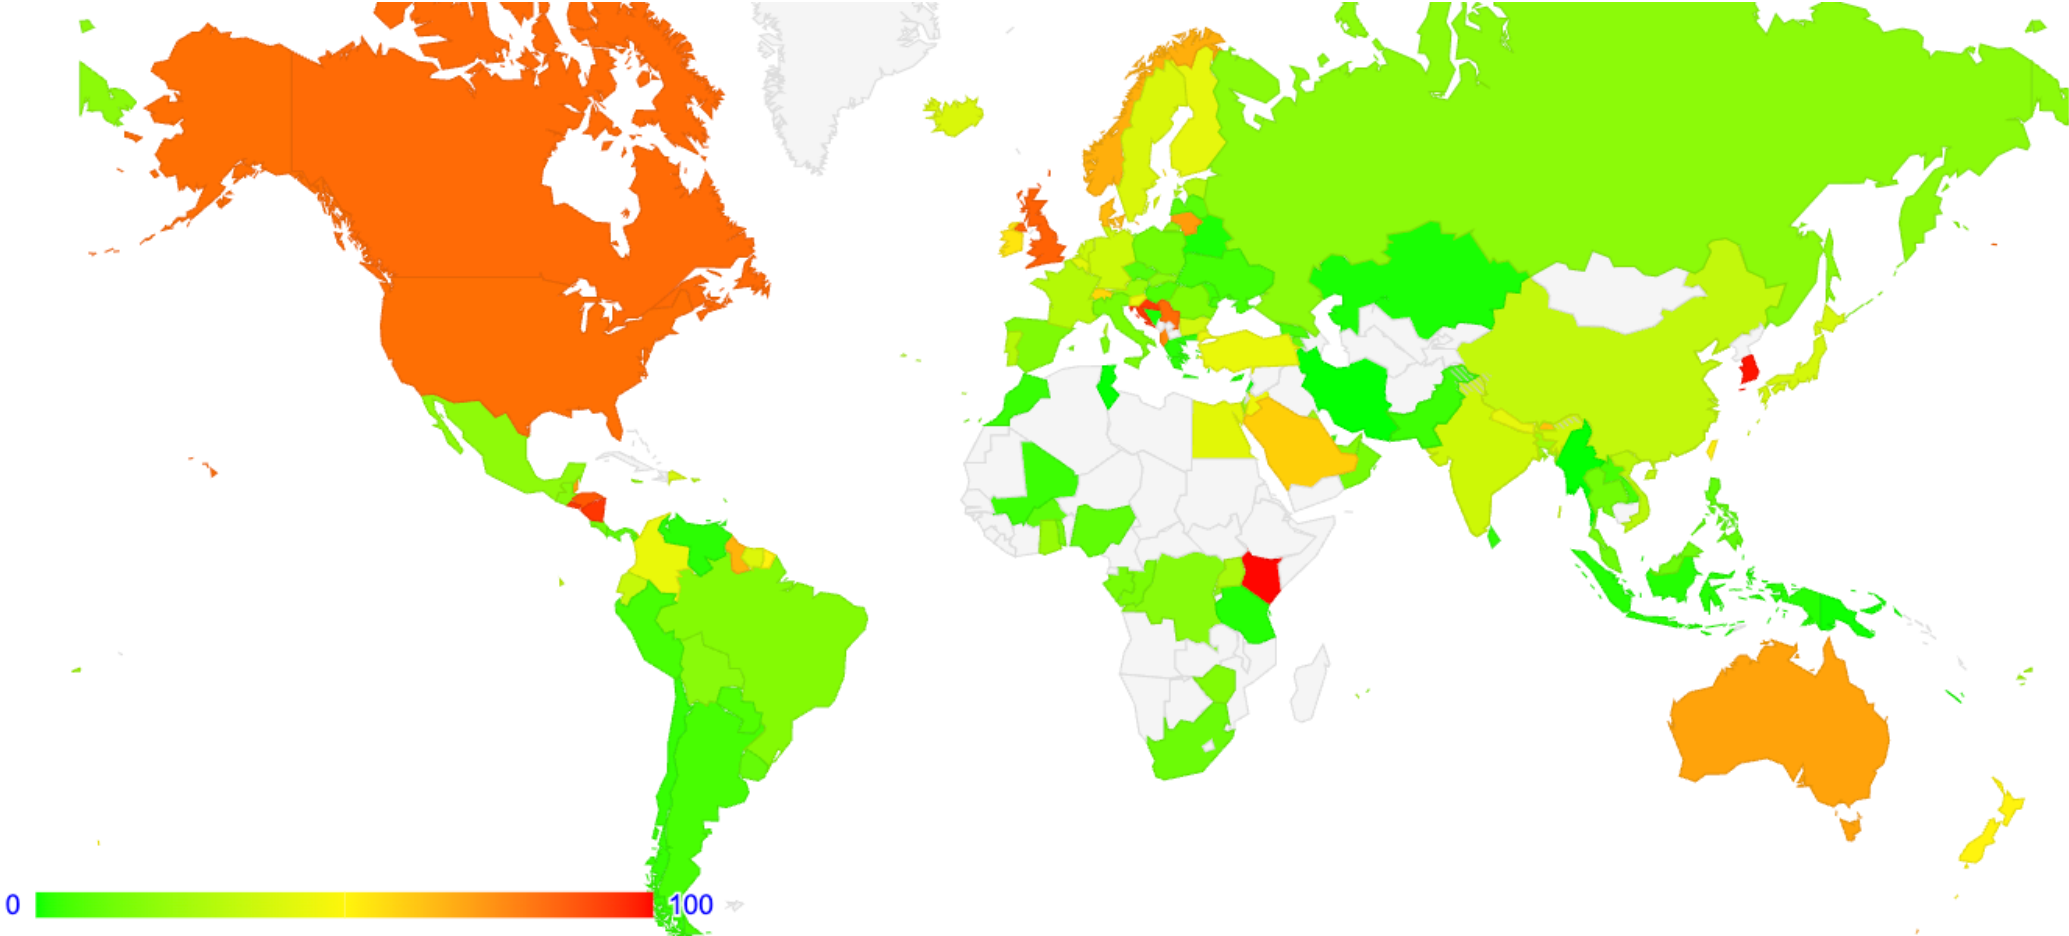


Greece

United States



Variations in DST drop rate



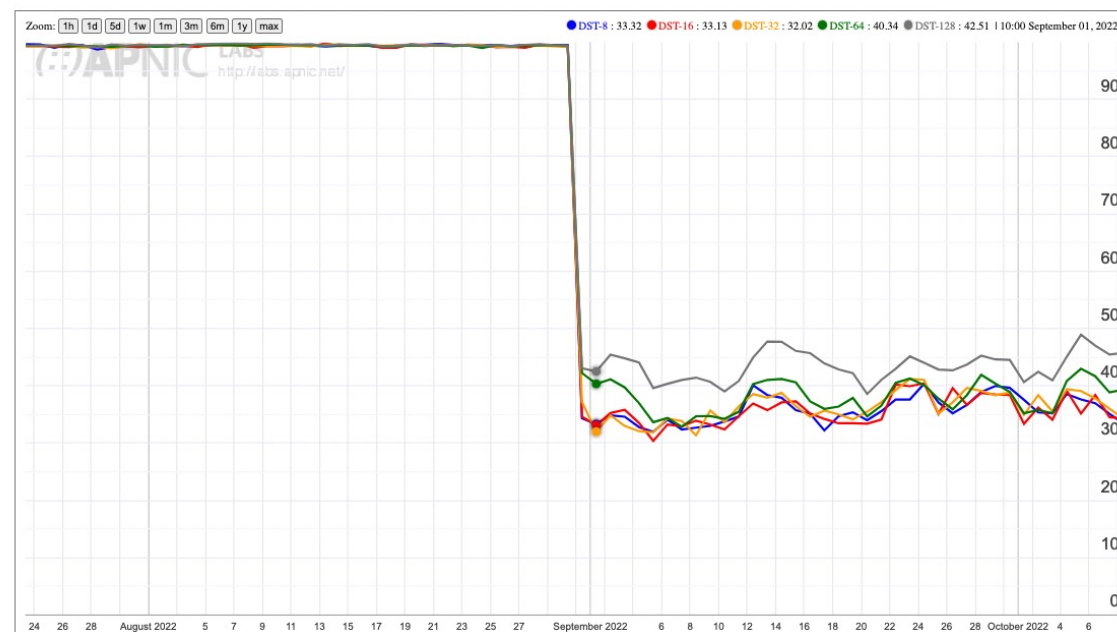
Network Type?

DST Drop Measurement for AS7922: COMCAST-7922



Fixed Line Provider

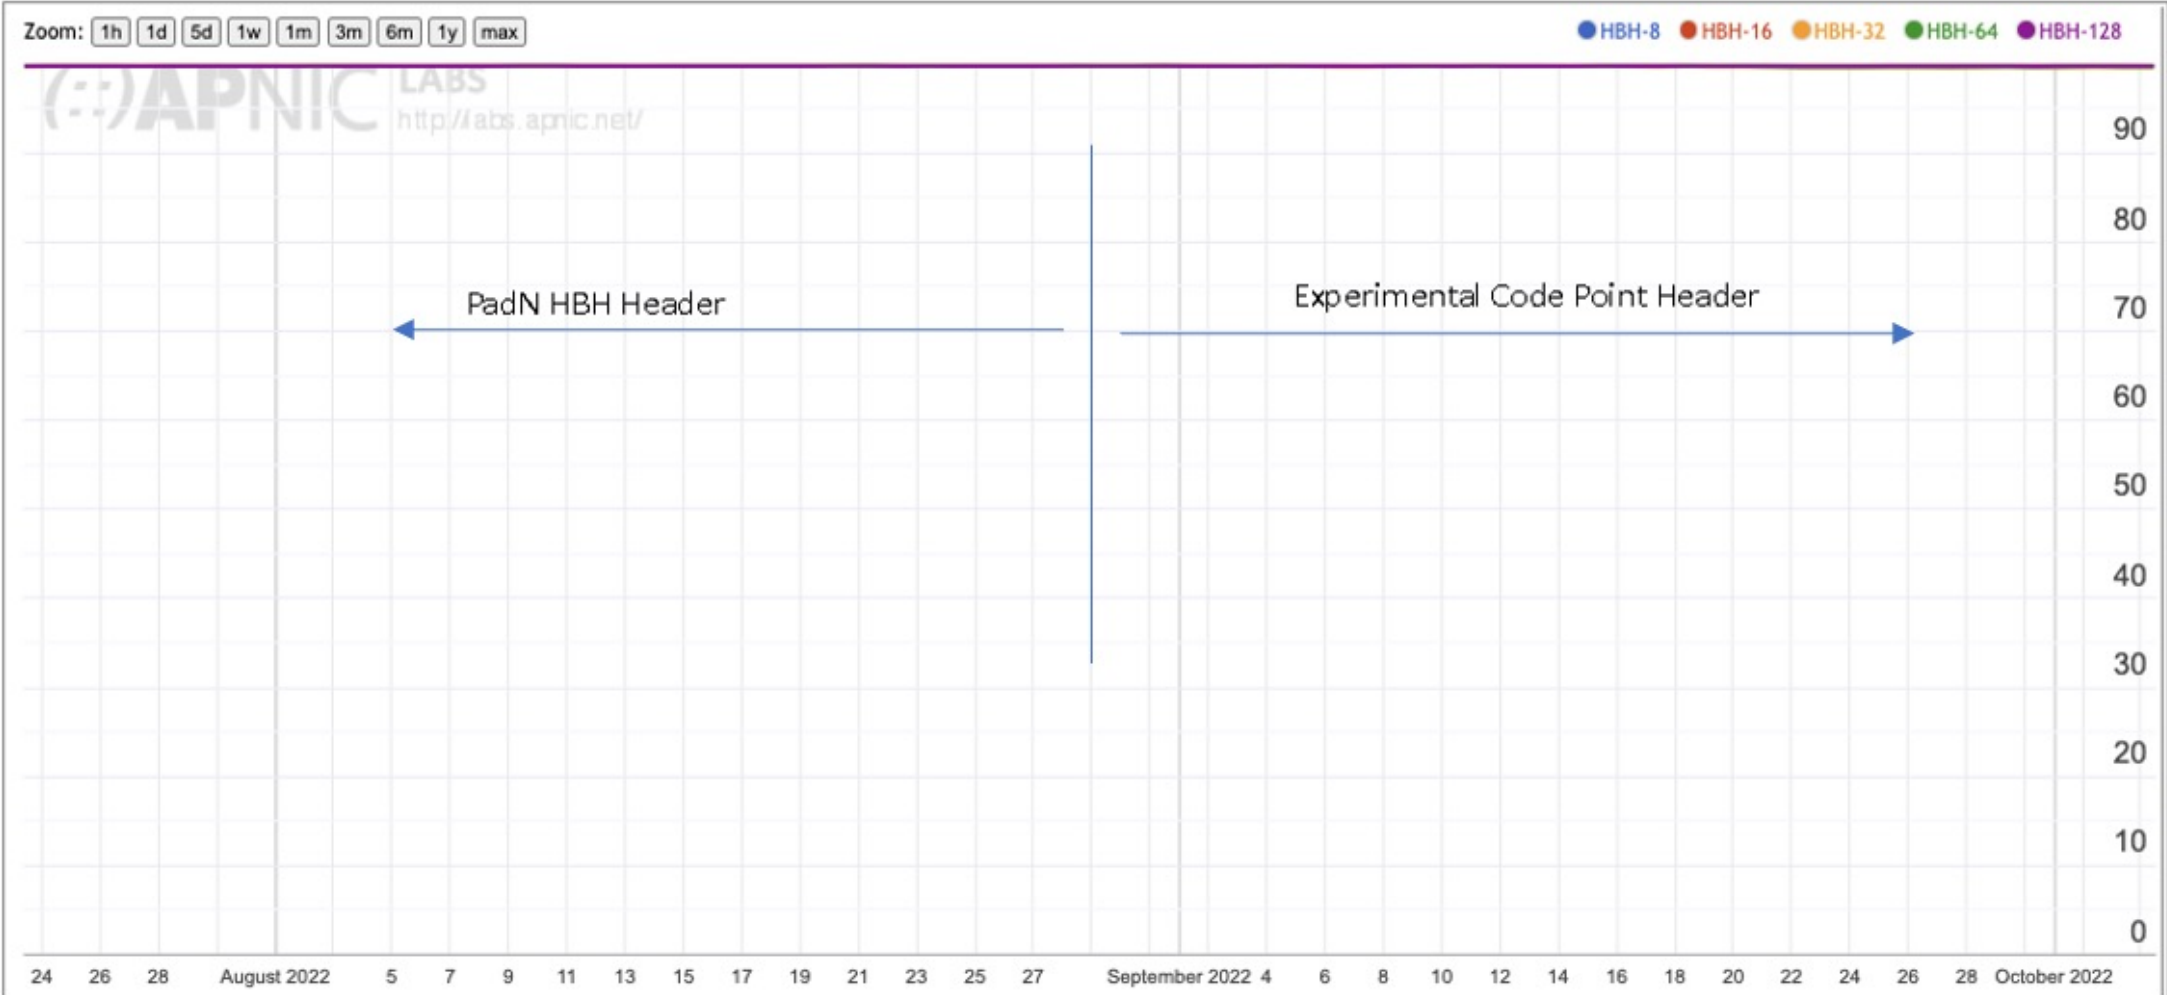
DST Drop Measurement for AS21928: T-MOBILE-AS21928




Mobile Provider

HBH EH drop rate

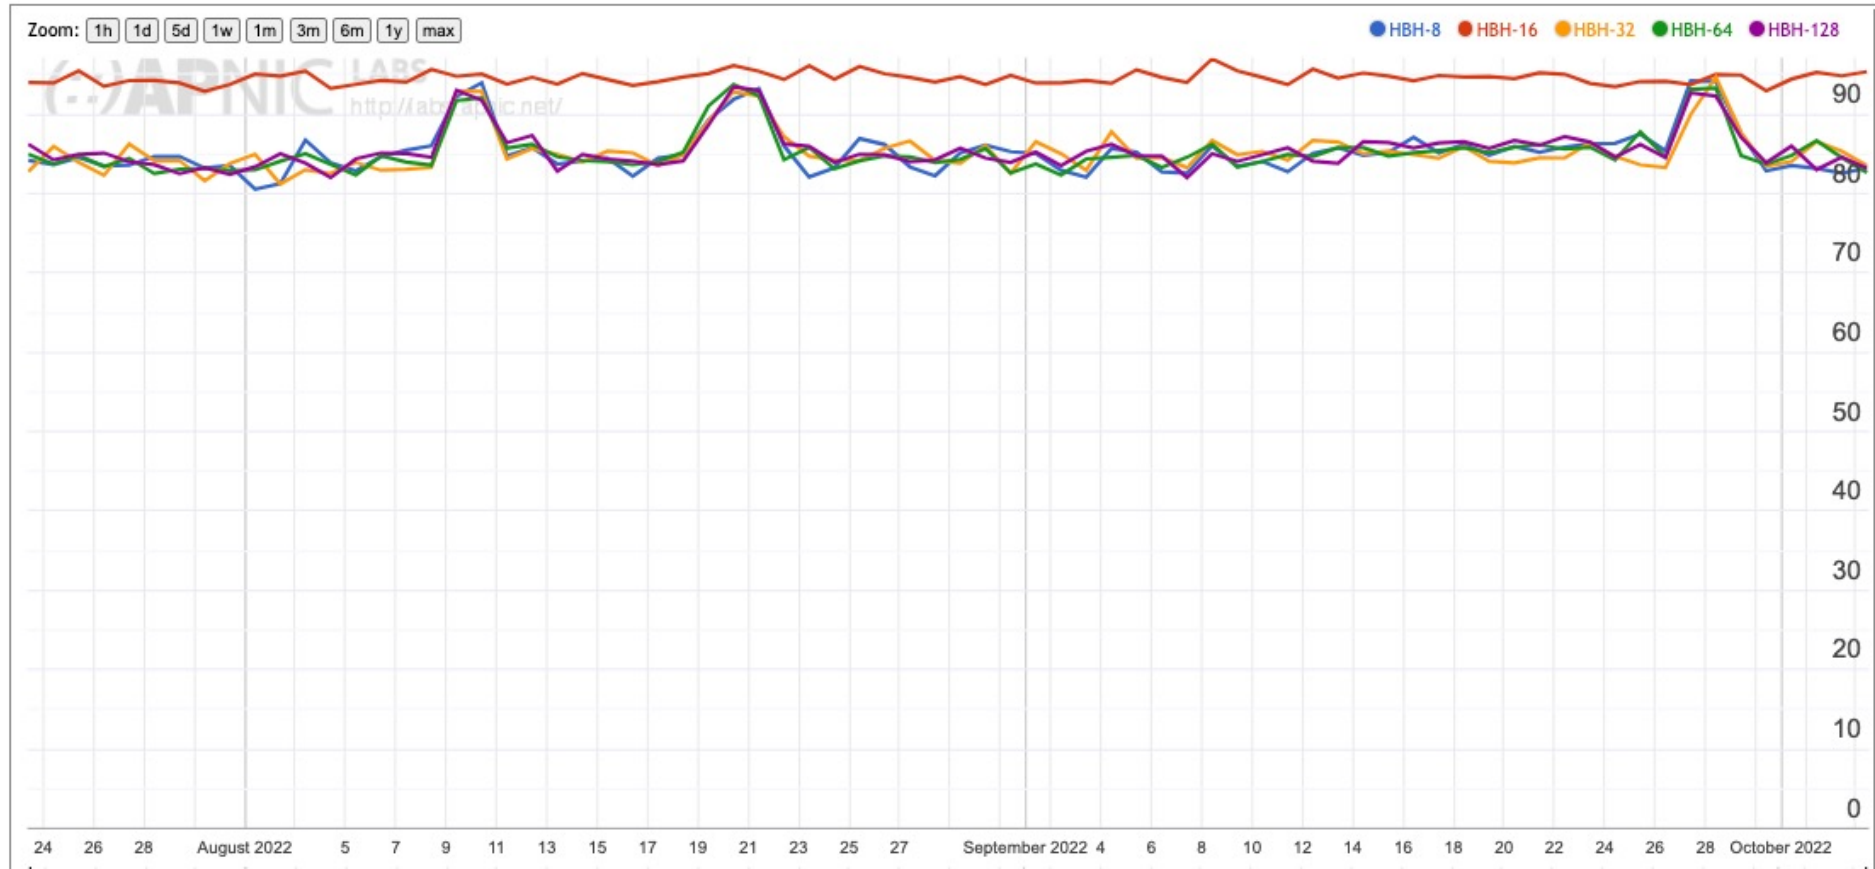
HBH EH drop rate



Everybody Everywhere



With one exception



HBH Option Drop Rate for ETISALAT-MIR (AS36992), Egypt

What can we say here?

“It depends”

- The IPv6 transition strategy
- The ISP equipment and config
- The CPE being used
- The mobile platform(s) being used
- End devices being used
- Private Relays / Proxies
- <insert reason>

What can we say?

- There are issues with networks, hosts, CPE, and the IPv6 selected transition technology
- Depending on where you test, and what you test (end-to-end, end-to-infrastructure, infrastructure-to-end, mobile ends, fixed line ends) its likely that you are going to see differing results for fragmentation and DST EHs
- We can't see any HBH of any size with one exception (and the single exception tends to suggest that the packets are able to leave the source network we are using and enter the transit networks)

Thanks!