

An Update on IPv6 Fragmentation

Geoff Huston
Chief Scientist, APNIC



Initial Tests: 2014 (RFC 7872)

- August 2014 and June 2015
- Sent fragmented IPv6 packets towards “well known” IPv6 servers (Alexa 1M and World IPv6 Launch)
- Drop Rate:



Dataset	DO8	HBH8	FH512
Web servers	10.91% (46.52%/53.23%)	39.03% (36.90%/46.35%)	28.26% (53.64%/61.43%)
Mail servers	11.54% (2.41%/21.08%)	45.45% (41.27%/61.13%)	35.68% (3.15%/10.92%)
Name servers	21.33% (10.27%/56.80%)	54.12% (50.64%/81.00%)	55.23% (5.66%/32.23%)

Table 2: Alexa's Top 1M Sites Dataset: Packet Drop Rate for Different Destination Types, and Estimated (Best-Case / Worst-Case) Percentage of Packets That Were Dropped in a Different AS

APNIC Test - August 2017

- Use APNIC IPv6 measurement platform to test the drop rate of IPv6 packets flowing in the opposite direction (server to client)

	Count	%
Tests	1,675,898	
ACK Fragmented Packets	1,324,834	79%
Fragmentation Loss	351,064	21%



This is an improvement over the RFC 7872 measurement


APNIC Test - 2021

Re-work of the 2017 measurement experiment

- Same server-to-client TCP session fragmentation mechanism
- Uses a middlebox to fragment outgoing packets - drop is detected by a hung TCP session that fails to ACK the sequence number in the fragmented packet
- This time we randomly vary the initial fragmented packet size between 1,200 and 1,416 bytes
- Performed as an ongoing measurement

2021 Fragmentation Drop Rate

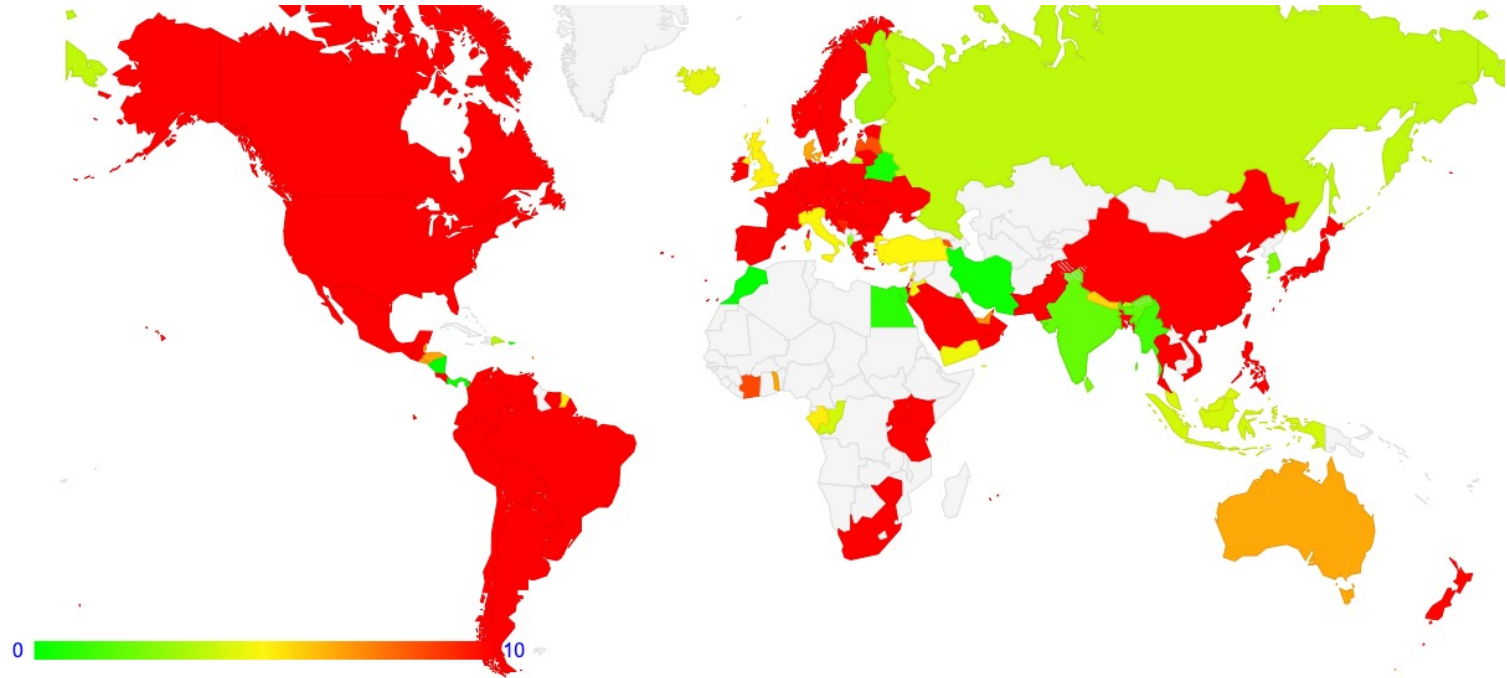
Use of V6FRAG Drop Rate for World (XA)



This is a significant improvement over 2017 data

Since 2017 there are 10x the number of IPv6 users and the fragmentation drop rate has come down by 2/3 - we appear to be getting better at handling IPv6 fragments!

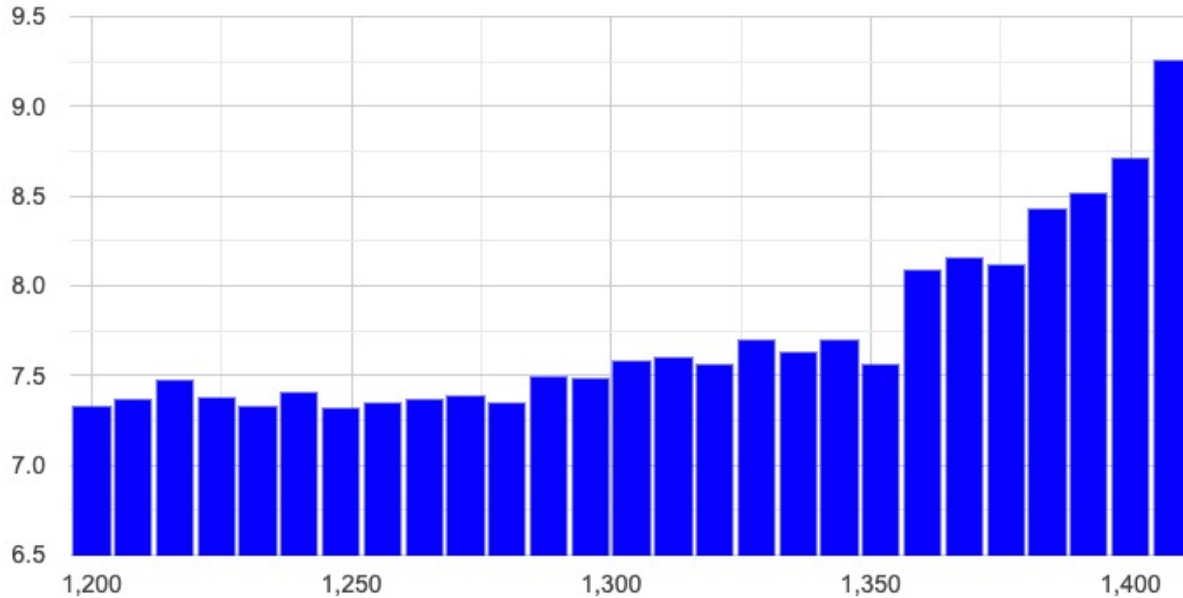
2021 Fragmentation Drop Rate



More recent IPv6 deployments appear to be a lot better than more mature ones

Drop Rate by Size


V6FRAG Drop Rates by Frag Size




This is unexpected. At a total IPv6 packet size of 1408 bytes we did not expect to see higher packet drop rates for this packet size, as there is still an IP encapsulation budget of 92 bytes

Drop Size Profile by Region


North America



Europe



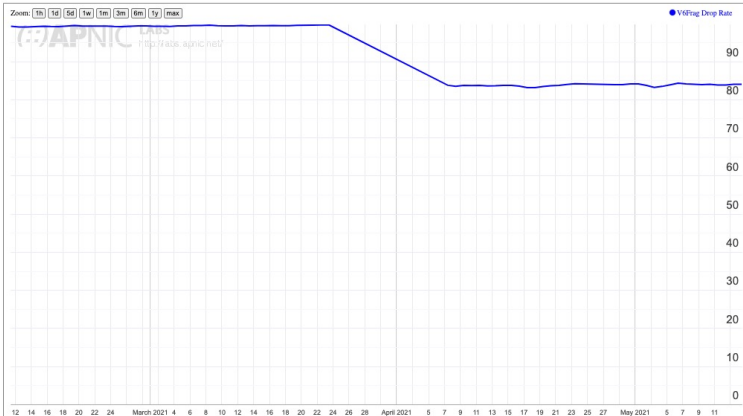
Asia



Why?

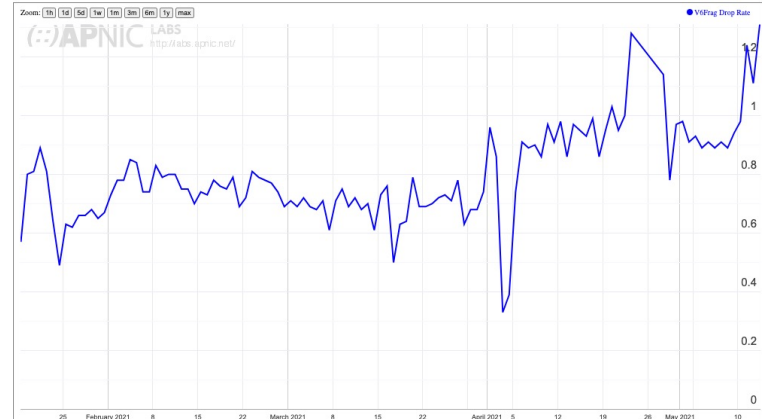
- Drop patterns vary across service providers, so there are probably contributory factors from network equipment and configurations

V6Frag Drop Measurement for AS852: ASN852, Canada (CA)



80% Drop

V6Frag Drop Measurement for AS45609: BHARTI-MOBILITY-AS-AP Bharti Airtel Ltd. AS for GPRS Service, India (IN)



1% Drop

Why?

Other potential factors that could contribute:

- Local security policies
- IPv6 EH may trigger “slow path” processing in network equipment that could lead to higher drop rates
- IPv6 Path MTU woes!

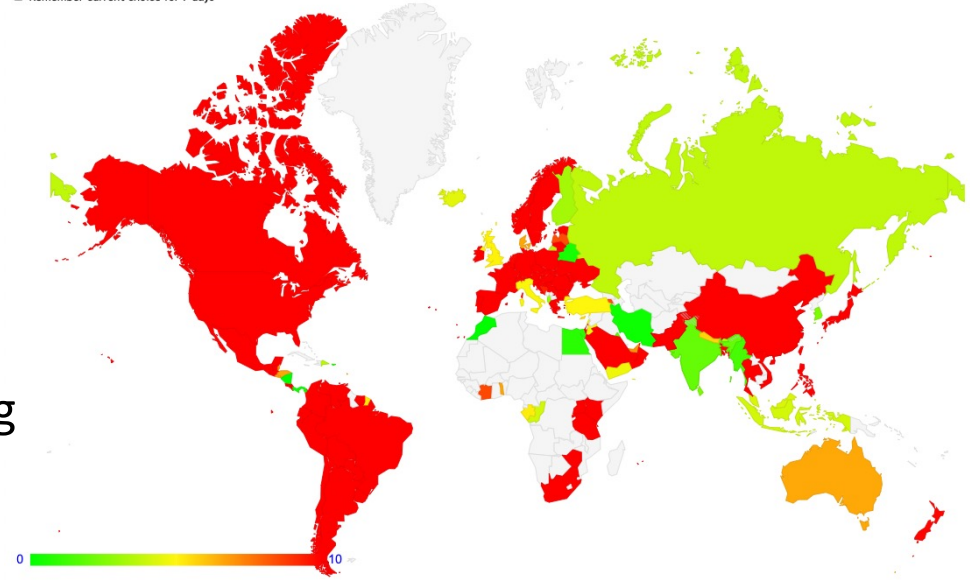
Daily Report

<https://stats.labs.apnic.net/v6frag>

IPv6 Fragmentation Drop Rate by country (%)

[Click here for a zoomable map](#)

Remember current choice for 7 days



7 day average (08/05/2021 - 14/05/2021)

Window (Days)

Code	Region	Frag Drop Rate	V6 Samples
XA	World	8.10%	18,509,740
XC	Americas	14.14%	4,474,221
XE	Europe	13.84%	1,569,881
XB	Africa	12.41%	143,540
XF	Oceania	7.13%	85,580
XG	Unclassified	5.26%	20,677
XD	Asia	5.11%	12,215,841



Summary

- The network is improving it's handling of fragmented packets
- In 5 years its gone from *unusably bad* to *tolerably poor*
- Recent IPv6 deployments appear to show more robust handling of IPv6 packets
- Smaller frags appear to be more robust than larger ones (if you are going to fragment a packet, prefer smaller fragment sizes)

That's it!