

# Deep Dive into IPv6 Extension Header Testing Across the Internet

IEPG: IETF115

Nalini Elkins, Mike Ackermann: Industry Network Technology Council

Dhruv Dhody, Praneet Kaur: India Internet Engineering Society

Dr. Mohit Tahiliani: NITK Surathkal

Dr. Priyanka Sinha: Zenatix Solutions

Ameya Deshpande: Google

Dr. Ana Custura: University of Aberdeen

# Can IPv6 Extension Headers Be Used on the Internet?

- Controversy for many years
- A number of studies showing that IPv6 extension headers “don’t work”
- Studies (by and large) sent “fake” IPv6 extension headers to Alexa top n sites
- If this is true, our work on our IPv6 Extension Header Destination Option Performance and Diagnostic Metrics (PDM) is really for naught

# Brief explanation of PDM

- RFC8250: IPv6 Performance and Diagnostic Metrics (PDM) Destination Option
- To assess performance problems, this document describes optional headers embedded in each packet that provide sequence numbers and timing information as a basis for measurements. Such measurements may be interpreted in real time or after the fact. This document specifies the Performance and Diagnostic Metrics (PDM) Destination Options header.

# IAB Workshop: Encrypted Mgmt Techniques

- How to manage encrypted networks a big problem
- Our proposal to use PDMv2 (encrypted Dest. Options IPv6 Extension Header) accepted to IAB workshop
- As soon as we have a stable implementation, will try to collocate at various points in the Internet
- Crucial that EH works

# What we did last time

- Used a small hosting service (not one of the “brand-name” ones)
- Locations throughout the world
  1. PDM-Warsaw
  2. PDM-Toronto
  3. PDM-Seattle
  4. PDM-Mumbai
  5. PDM-Melbourne
  6. PDM-Frankfurt

All machines are FreeBSD with a modification to the kernel to send PDM IPv6 Destination option with every packet

# Thanks to...



India Internet Engineering Society



National Institute of Technology  
Karnataka, Surathkal

राष्ट्रीय प्रौद्योगिकी संस्थान  
कर्नाटक, सुरत्कल

In particular, Dr. Mohit Tahiliani  
and his undergraduate students:  
Balajinaidu V, Chinmaya  
Sharma, and Amogh Umesh



Industry Network Technology Council

# Tested large FTP: Toronto to Mumbai (with PDM)

- Connected to **2401:c080:2400:1179:5400:04ff:fe0f:804a.**
- 220----- Welcome to Pure-FTPd [privsep] [TLS] -----
- 220-You are user number 1 of 50 allowed.
- 220-Local time is now 15:12. Server port: 21.
- 220 You will be disconnected after 15 minutes of inactivity.
- 331 User PDMuser OK. Password required
- 230 OK. Current directory is /
- Remote system type is UNIX.
- Using binary mode to transfer files.

- 229 Extended Passive mode OK (|||3353|)
- 150-Accepted data connection
- 150 **27872.0 kbytes to download**
- 100%  
|\*\*\*\*\*  
\*\*\*\*\*  
\*\*\*\*\*  
\*| 27872 KiB 222.31 KiB/s 00:00 ETA
- 226-**File successfully transferred**
- 226 125.107 seconds (measured here), 222.78 Kbytes per second
- 28540928 bytes received in 02:05 (222.31 KiB/s)
- 221-Goodbye. You uploaded 0 and downloaded 27872 kbytes.
- 221 Logout.



Apply a display filter ... <Ctrl-/>

No.	Time	Source	Destination	Protocol	PSN This Packet	PSN Last Received	Info
38	2.857775	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	2401:c080:2400:1179:5400:4ff:fe0f:804a	TCP	20489	0	61272 → 53696 [SYN] Seq=0 Win=65535 Len=0 M
39	2.963460	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	2401:c080:2400:1179:5400:4ff:fe0f:804a	TCP	14104	12376	62443 → 21 [ACK] Seq=101 Ack=805 Win=66240
40	3.056635	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	TCP	23911	20489	53696 → 61272 [SYN, ACK] Seq=0 Ack=1 Win=65
41	3.056686	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	2401:c080:2400:1179:5400:4ff:fe0f:804a	TCP	20490	23911	61272 → 53696 [ACK] Seq=1 Ack=1 Win=66240 L
42	3.056735	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	2401:c080:2400:1179:5400:4ff:fe0f:804a	FTP	14105	12376	Request: RETR out.txt
43	3.253255	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23912	20490	IPv6 fragment (off=0 more=y ident=0x73059a8
44	3.253284	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23912	20490	IPv6 fragment (off=1432 more=y ident=0x7305
45	3.253290	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23912	20490	IPv6 fragment (off=2864 more=y ident=0x7305
46	3.253298	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23912	20490	IPv6 fragment (off=4296 more=y ident=0x7305
47	3.253304	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23912	20490	IPv6 fragment (off=5728 more=y ident=0x7305
48	3.253315	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23912	20490	IPv6 fragment (off=7160 more=y ident=0x7305
49	3.253326	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23912	20490	IPv6 fragment (off=8592 more=y ident=0x7305
50	3.253332	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23912	20490	IPv6 fragment (off=10024 more=y ident=0x730
51	3.253341	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23912	20490	IPv6 fragment (off=11456 more=y ident=0x730
52	3.253350	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	FTP-DATA	23912	20490	FTP Data: 14280 bytes (EPASV) (RETR out.txt
53	3.253399	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	2401:c080:2400:1179:5400:4ff:fe0f:804a	TCP	20491	23912	61272 → 53696 [ACK] Seq=1 Ack=14281 Win=519
54	3.266651	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	FTP	12377	14105	Response: 150-Accepted data connection
55	3.372449	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	2401:c080:2400:1179:5400:4ff:fe0f:804a	TCP	14106	12377	62443 → 21 [ACK] Seq=115 Ack=867 Win=66240
56	3.449235	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=0 more=y ident=0x7acf3f8
57	3.449249	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=1432 more=y ident=0x7acf
58	3.449277	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=2864 more=y ident=0x7acf
59	3.449283	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=4296 more=y ident=0x7acf
60	3.449289	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=5728 more=y ident=0x7acf
61	3.449316	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=7160 more=y ident=0x7acf
62	3.449324	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=8592 more=y ident=0x7acf
63	3.449336	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=10024 more=y ident=0x7ac
64	3.449349	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=11456 more=y ident=0x7ac
65	3.449355	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=12888 more=y ident=0x7ac
66	3.449363	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	IPv6	23913	20491	IPv6 fragment (off=14320 more=y ident=0x7ac
67	3.449369	2401:c080:2400:1179:5400:4ff:fe0f:804a	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	FTP-DATA	23913	20491	FTP Data: 17136 bytes (EPASV) (RETR out.txt
68	3.449430	2001:19f0:b001:6ce:5400:4ff:fe0f:806d	2401:c080:2400:1179:5400:4ff:fe0f:804a	TCP	20492	23913	61272 → 53696 [ACK] Seq=1 Ack=31417 Win=490

From PDM IPv6 DOH



# Bottom line

1. PDM-FTP Toronto to Warsaw - worked
2. PDM-FTP Toronto to Seattle - worked
3. PDM-FTP Toronto to Mumbai - worked
4. PDM-FTP Toronto to Melbourne - worked
5. PDM-FTP Toronto to Frankfurt - worked

Traces available for all to look at.

# Deep Dive: Why?

- Find out what is the ACTUAL situation -- do EHs really work?
- If not, then why?
- Is it blocked:
  - At the source?
  - At the destination?
  - In a transit network?
- Then
  - Is it intentional?

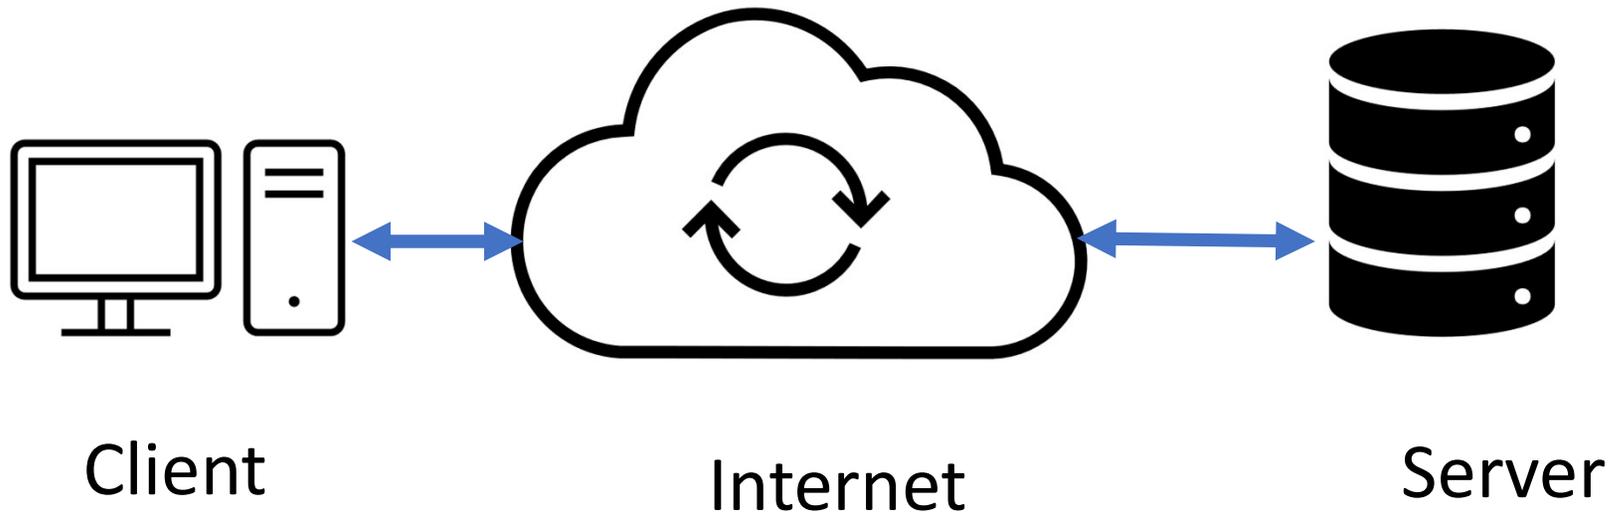
# Why CDNs?

- Will talk more about standalone client – server testing at side meeting (Thursday)
- Started focusing on CDNs
- Many high volume sites on the internet use CDNs

# Let's look at topologies

- Client – Internet – Server
- Client – Internet – CDN Cache Server – CDN network – Origin Server
  - (Internal to CDN may have multiple more complex topologies)
- Client – Internet – Edge of Cloud Provider – Origin server hosted by cloud provider

Simplest: Client – Internet -- Server



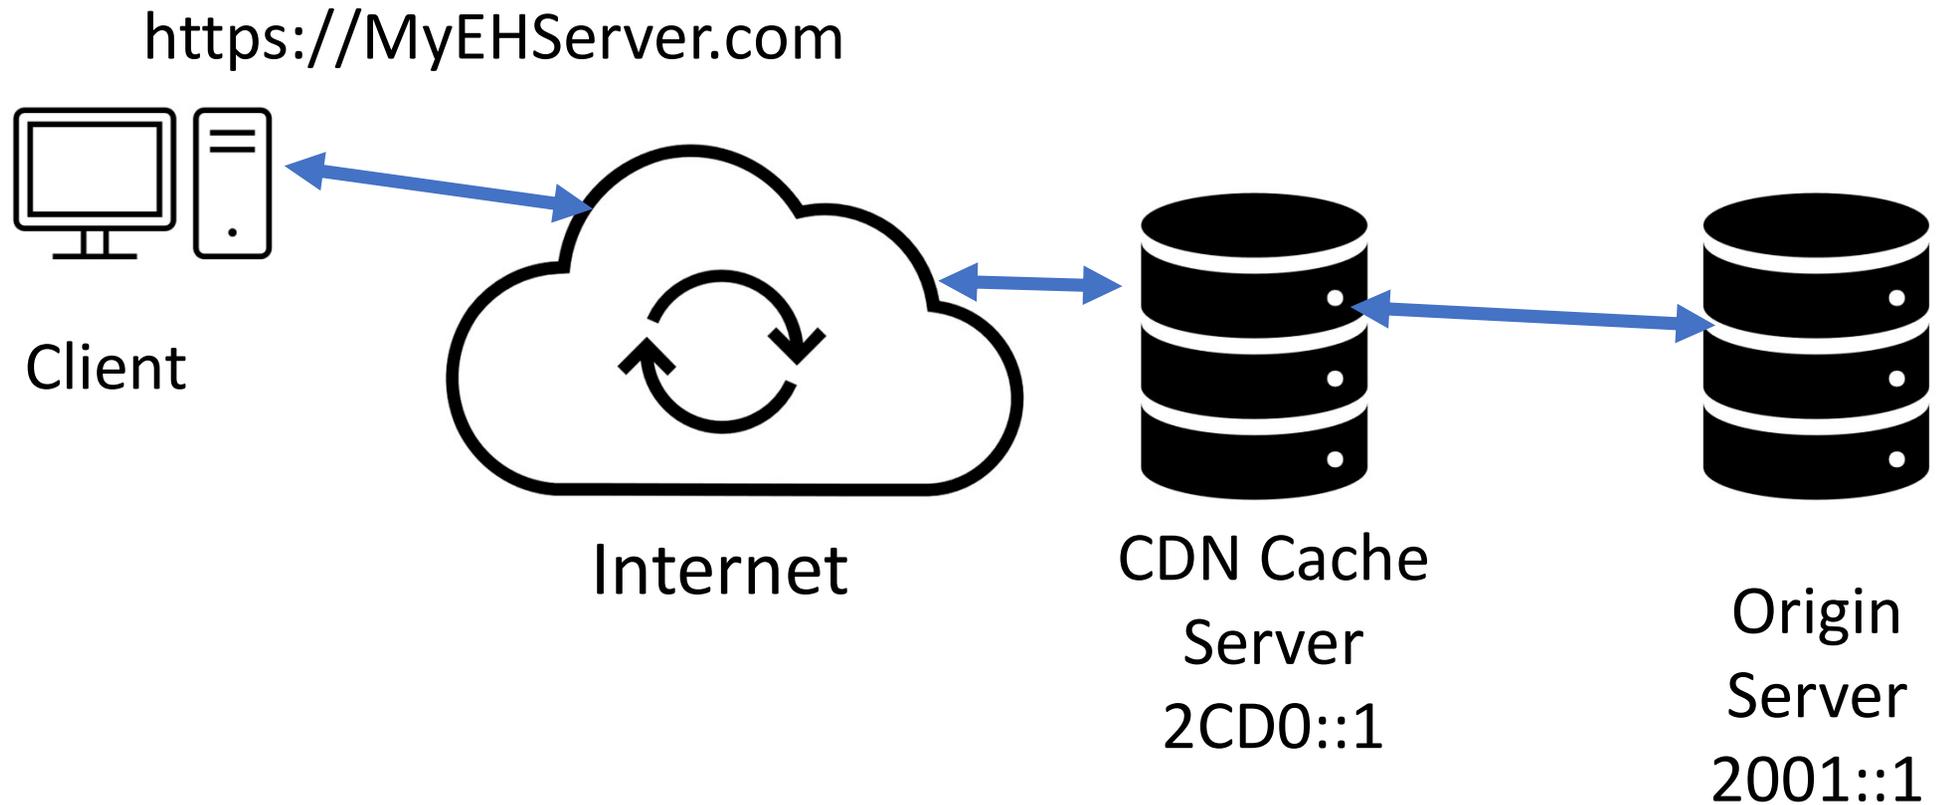
# Move Server Behind CDN

- Our server has a domain name: MyEHServer
- Our server also has an IPv6 address (also IPv4 probably)
- Let's say: 2001::1 and 201.1.1.1 (MyEHServer resolves to these)
  
- To move behind a CDN, you have to give the CDN authority to resolve MyEHServer
- Let's give the CDN IPv6 addresses starting with 2CD0::/64 ( 2CD0::1, 2CD0::2, etc)
- After CDN move, MyEHServer will resolve to some CDN cache server address (2CD0::1 for example)



We will now refer to our server as the “Origin Server”

# With CDN Topology



So, first ...

- What is the CDN actually resolving to?
- Remember: To move behind a CDN, you have to give the CDN authority to resolve MyEHServer
- We *\*assume\** it is resolving to the IPv6 address. Of course, why wouldn't we think that? We are IETFers...
- So, I took a trace on both sides (my client machine which is also EH enabled and the origin server behind the CDN which I control)



From our Extension Header



ipv6.addr == 2606:4700:130:436c:6f75:6466:6c61:7265

No.	Time	Source	Destination	Protocol	Length	PSN This Packet	Hop Li
1	0.000000	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TCP	110	23977	
4	0.001507	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TCP	86		
5	0.001546	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TCP	90	23978	
6	0.015248	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TLSv1.2	607	23979	
7	0.016439	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TCP	74		
8	0.019725	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TLSv1.2	4563		
9	0.019757	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TCP	90	23980	
10	0.021177	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TLSv1.2	216	23981	
11	0.022281	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TCP	74		
12	0.022445	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TLSv1.2	125		
13	0.022634	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TLSv1.2	201	23982	
14	0.023756	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TCP	74		
15	0.267409	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TLSv1.2	1472		
16	0.267481	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TCP	90	23983	
17	0.267500	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TLSv1.2	1472		
18	0.267507	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TCP	90	23984	
19	0.267612	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TLSv1.2	2870		
20	0.267650	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TCP	90	23985	
21	0.267659	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TLSv1.2	5666		
22	0.267675	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TCP	90	23986	
23	0.267682	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TLSv1.2	4173		
24	0.267705	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TCP	90	23987	
25	0.372902	2606:4700:130:436c:6f75:6466:6c61:7265	2001:19f0:b002:392:5400:4ff:fe1f:9900	TLSv1.2	1472		
26	0.372948	2001:19f0:b002:392:5400:4ff:fe1f:9900	2606:4700:130:436c:6f75:6466:6c61:7265	TCP	90	23988	

Frame 4: 86 bytes on wire (688 bits) · 86 bytes captured (688 bits)

## From Client Side

- I see that packets with EH are getting through
- But, I only see my EH data on the CLIENT side.
- Why not the server side?
- So, I look at the server trace



Apply a display filter ... <Ctrl-/>

No.	Time	Source	Destination	Protocol	Length	Info
5	14.723375	79.124.62.130	70.34.248.166			
6	14.723399	70.34.248.166	79.124.62.130	TCP	54	
7	14.763039	79.124.62.130	70.34.248.166	TCP	54	
8	15.617507	108.162.241.132	70.34.248.166	TCP	66	
9	15.617560	70.34.248.166	108.162.241.132	TCP	66	
10	15.725303	108.162.241.132	70.34.248.166	TCP	54	
11	15.725417	108.162.241.132	70.34.248.166	HTTP	424	
12	15.726336	70.34.248.166	108.162.241.132	TCP	14654	
13	15.834054	108.162.241.132	70.34.248.166	TCP	54	
14	15.834086	70.34.248.166	108.162.241.132	HTTP	13867	
15	15.834242	108.162.241.132	70.34.248.166	TCP	54	
16	15.941765	108.162.241.132	70.34.248.166	TCP	54	
17	17.901406	89.248.165.63	70.34.248.166	TCP	54	
18	17.901436	70.34.248.166	89.248.165.63	TCP	54	
19	17.928300	89.248.165.63	70.34.248.166	TCP	54	
20	18.143159	180.149.126.177	70.34.248.166	TCP	54	
21	18.143190	70.34.248.166	180.149.126.177	TCP	54	
22	20.792643	70.34.248.166	108.162.241.132	TCP	54	
23	20.900360	108.162.241.132	70.34.248.166	TCP	54	
24	20.900405	70.34.248.166	108.162.241.132	TCP	54	
25	21.820791	162.0.208.183	70.34.248.166	TCP	54	
26	21.820820	70.34.248.166	162.0.208.183	TCP	54	
27	21.971194	162.0.208.183	70.34.248.166	TCP	54	

**First shock, why am I only seeing IPv4? What did I do wrong?**

Frame 11: 424 bytes on wire (3392 bits) 424 bytes captured (3392 bits)

OnPDMWarsawBehind...e.pcap

File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

Apply a display filter ... <Ctrl-/>

No.	Time	Source	Destination	Protocol	Length	PSN This Packet	Hop Li
11	15.725417	108.162.241.132	70.34.248.166	HTTP	424		

> Frame 11: 424 bytes on wire (3392 bits), 424 bytes captured (3392 bits)  
> Ethernet II, Src: fe:00:04:0f:80:59 (fe:00:04:0f:80:59), Dst: 56:00:04:0f:80:59 (56:00:04:0f:80:59)  
> Internet Protocol Version 4, Src: 108.162.241.132, Dst: 70.34.248.166  
> Transmission Control Protocol, Src Port: 37590, Dst Port: 80, Seq: 1, Ack: 1, Len: 370

> Hypertext Transfer Protocol

- > GET / HTTP/1.1\r\n
- Host: www.exthdrtest.com\r\n
- Connection: Keep-Alive\r\n
- Accept-Encoding: gzip\r\n
- X-Forwarded-For: 2001:19f0:b002:392:5400:4ff:fe1f:9900\r\n
- CF-RAY: 743e5b69cd125467-YYZ\r\n
- X-Forwarded-Proto: https\r\n
- CF-Visitor: {"scheme":"https"}\r\n
- User-Agent: curl/7.50.1\r\n
- Accept: \*/\*\r\n
- CF-Connecting-IP: 2001:19f0:b002:392:5400:4ff:fe1f:9900\r\n
- CF-IPCountry: CA\r\n
- CDN-Loop: \r\n
- \r\n
- [Full request URI: <http://www.exthdrtest.com/>]
- [HTTP request 1/1]
- [Response in frame: 14]

**No... it is forwarded from IPv6!!!!**

# From CDN Provider #1 Documentation

## Configure IPv6 compatibility

If your hosting provider supports IPv6 for your origin web server, [redacted]'s IPv6 Compatibility allows you to route IPv6 connections through [redacted]'s global network when proxying AAAA DNS records. When both IPv4 and IPv6 connections are available, [redacted] prefers IPv4.

Domains on **Enterprise** plans can toggle IPv6 compatibility within the [redacted] dashboard:

1. Login to your [redacted] account.
2. Select the appropriate domain.
3. Click the **Network** app.
4. Toggle **IPv6 Compatibility** *Off* or *On*.

**But, CDN  
provider #1 is  
better than ...**

## Email from CDN Provider #2

Hi Nalini,

Thank you for the quick conversation on the phone, I was able to find some current documentation on how we support IPv6 - but as it turns out you were correct in that we do not support it at the origin. This is something we have on the product roadmap for  but there is no current timetable on implementation.

## CDN Provider #3

- Very cooperative
- Having trouble navigating how to actually get behind the CDN!  
(Will have results next time.)
- They indicate that IPv6 to origin is supported but EH to origin not supported
- We hope to work with them to support EH to origin

# More Breaking News as it Happens ...

- Will continue testing
- Working on cloud providers
- Thoughts?
- Questions?

**Side meeting on EH testing:  
Thursday! Join us.**

Ana's Portion!