

Framework for Multi-domain IPv6-only Underlay Network and IPv4 as a Service

draft-ietf-v6ops-framework-md-ipv6only-underlay

Chongfeng Xie(Presenter) China Telecom

Chenhao Ma China Telecom

Xing Li CERNET Center/Tsinghua University

Gyan Mishra Verizon


Mohamed Boucadair Orange

Thomas Graf Swisscom

v6ops@IETF 117, July 2023

Overview

- This draft was proposed in Jan. 2022, first presented in IETF 113, adopted after IETF 115, the current version is -02.
- Its purpose is to provide end-to-end IPv4 service delivery over multi-domain IPv6-only underlay networks, and eliminate unnecessary IPv4/IPv6 conversions in a scalable way.



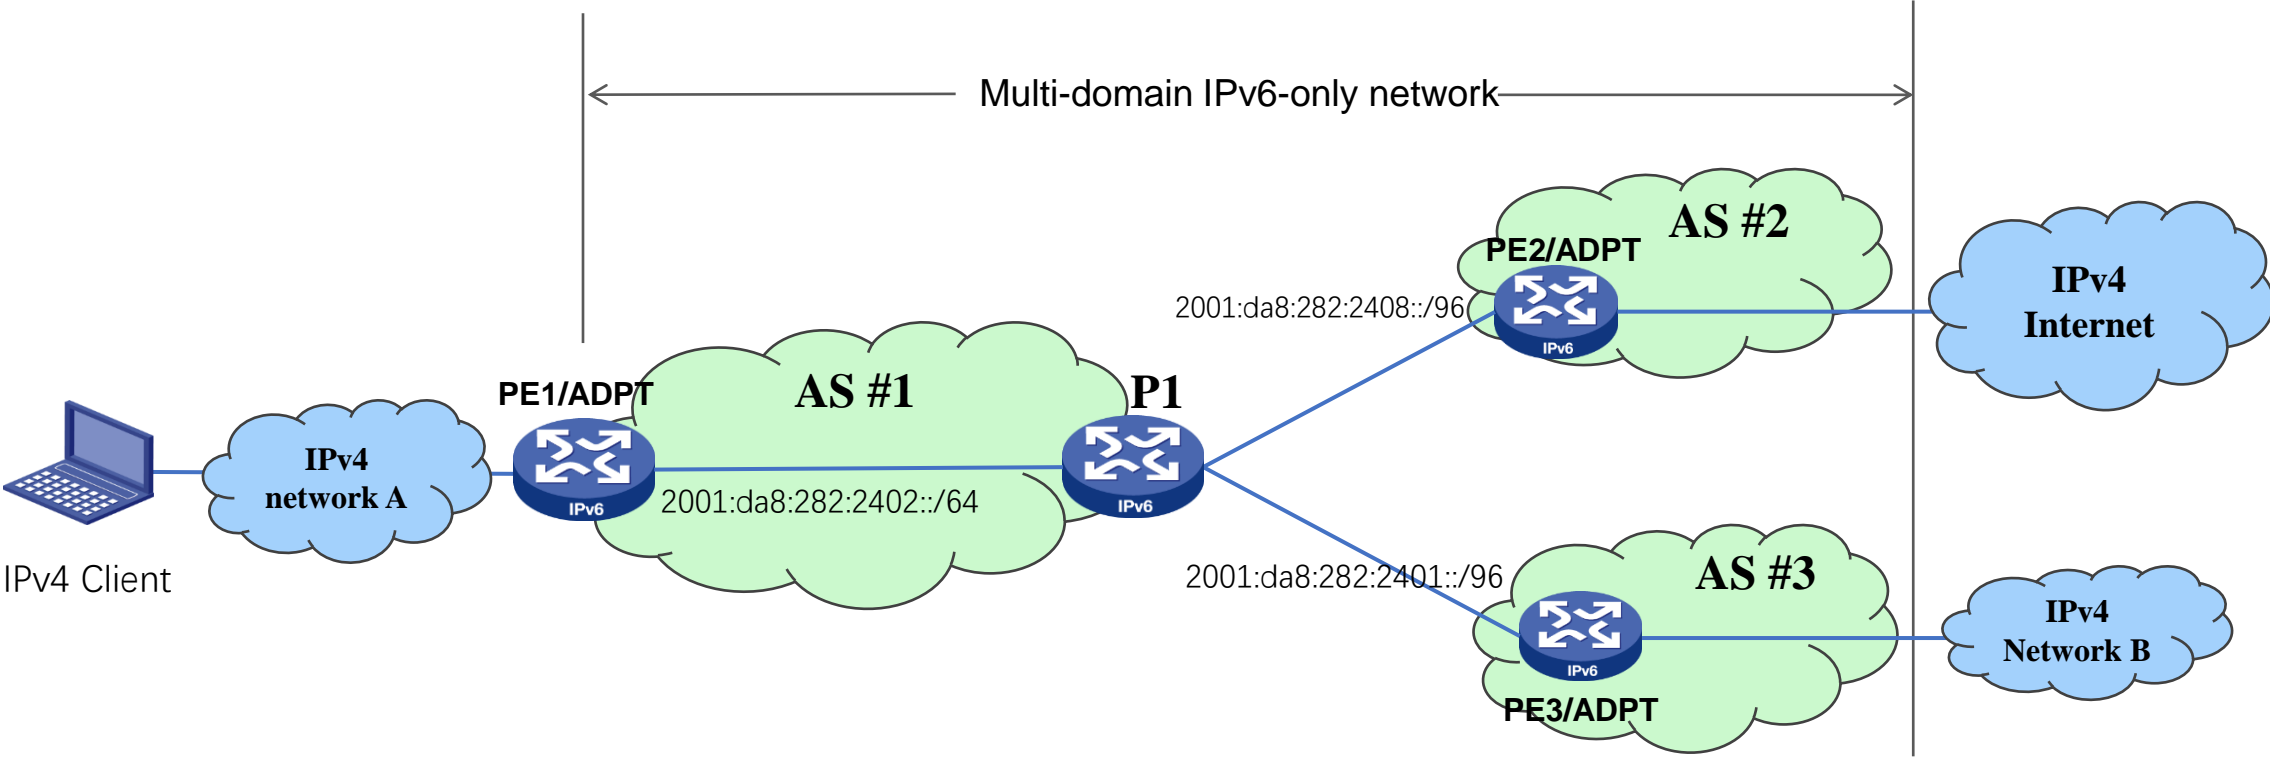
Acknowledgement for your comments

- Comments were received from Brian E. Carpenter, Bob Harold, Dhruv Dhody, Xipeng Xiao, Eduard Metz, Giuseppe Fioccola, Qin Wu, Shuping Peng, Zhenbin Li, Ron Bonica, Cheng Li, Vasilenko Eduard, Jingrong Xie, Aijun Wang, Dhruv Dhody, Nick Buraglio, Linda Dunbar, Guoliang Han, Weiqiang Cheng, Tianran Zhou, Huaimo Chen, etc.
- All are appreciated!

Revisions made since IETF 116

- 1) The following statements are added
 - Section 1: “For example, when broadband users experience abnormal access to services, network operators need to troubleshoot whether it is an IPv4 protocol failure or an IPv6 protocol failure, which increases the workload by at least twice.”
 - Section 4: “Since it is based on the OSPF protocol, it only supports IPv4aaS within a single AS.”
 - Section 6.1: “which needs to be obtained remotely in advance.”
- 2) Section 4: “backbone” → “network”
 - “They are routed through the network to another PE router, after which they leave the **network** and continue their way.”
- 3) “Users” is added in the sentence of section 4
 - “In other words, IPv6-only network should not only carry native IPv6 services, but also allow **users** to reach IPv4-only services.”
- 4) Nit fixed
 - Transpor → transport

System Implementation and Test



Original IPv4 Route vs New IPv6 Route

IPv4 route

```
BGP table version is 11, local router ID is 192.168.66.189, vrf id 0
Default local pref 100, local AS 64589
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop        Metric LocPrf Weight Path
*> 110.242.68.0/24 0.0.0.0          0         32768 i
*> 192.168.8.0/24  0.0.0.0          0         32768 i
*> 192.168.66.173/32
    192.168.66.173          0             0 64520 i
    192.168.66.175/32
    192.168.66.175          0             0 64550 i

Displayed 4 routes and 4 total paths
xlat#
```

110.242.68.0/24 is in IPv4 Internet
(110.242.68.3 is www.baidu.com)

IPv6 route received by P1 router

```
xlat# show bgp ipv6 2001:da8:282:2408::6ef2:4400/120
BGP routing table entry for 2001:da8:282:2408::6ef2:4400/120
Paths: (1 available, best #1, table default)
Advertised to non peer-group peers:
2001:da8:282:ff0c:62fd:f174:f6c3:189 2001:da8:282:ff0c:9ec5:d104:c664:171
64589
2001:da8:282:ff0c:62fd:f174:f6c3:189 from 2001:da8:282:ff0c:62fd:f174:f6c3:189 (192.168.66.189)
Origin IGP, metric 0, valid, external, best (First path received)
Last update: Thu Jul 13 02:30:19 2023
xlat#
```



2001:da8:282:2408:: 6e f2 : 44 00/120
IPv6 Mapping Prefix (/96) 110. 242. 68. 0/24

IPv6-IPv4 autoconfig

```
xlat# show bgp ipv6 2001:da8:282:2408::6ef2:4400/120
BGP routing table entry for 2001:da8:282:2408::6ef2:4400/120
Paths: (1 available, best #1, table default)
  Advertised to non peer-group peers:
    2001:da8:282:ff0c:62fd:f174:f6c3:189 2001:da8:282:ff0c:9ec5:d104:c664:171
    64589
    2001:da8:282:ff0c:62fd:f174:f6c3:189 from 2001:da8:282:ff0c:62fd:f174:f6c3:189 (192.168.66.189)
      Origin IGP, metric 0, valid, external, best (First path received)
      Last update: Thu Jul 13 02:30:19 2023
xlat#
```

PE1 received IPv6 route `2001:da8:282:2408::6ef2:4400/120` (attribute plen=96)


PE2 map this route to an IPv4-IPv6 translation rule:

`111.242.68.0/24` ---> rule3 (source-IPv6-prefix: `2001:da8:282:2402::/64`; dest-IPv6-prefix: `2001:da8:282:2408::/96`)

```
local IPv6 prefix:2001:da8:282:2402::/64
local map6id:fec0::1/64
local mapid:172.2.0.1/24
local nexthop:2001:da8:282:ff0c:aa4f:44c9:84a0:170
local prefix6 address:64:ff9b::/96
local dnat ns name:dnat66
local dnat ipset name:iviDns64Ipset
index:1      spl:2001:da8:282:2402::/64      dp:2001:da8:282:2401::/96      map6id:fec0::2 mapid:172.2.0.2 peerid:192.168.66.171 pref:73
index:2      spl:2001:da8:282:2402::/64      dp:2001:da8:282:2408::/96      map6id:fec0::3 mapid:172.2.0.3 peerid:192.168.66.189 pref:3
xlat#
```

IPv6 host → IPv4 server visit

Support IPv6 → IPv6, IPv4 → IPv4, IPv6 → IPv4, etc.



ping IPv4 Internet server from IPv6 hosts

windows IPv6 host

```
C:\Users\ivi>ping www.baidu.com

正在 Ping www.a.shifen.com [64:ff9b::6ef2:4403] 具有 32 字节的数据:
来自 64:ff9b::6ef2:4403 的回复: 时间=26ms
来自 64:ff9b::6ef2:4403 的回复: 时间=24ms

64:ff9b::6ef2:4403 的 Ping 统计信息:
    数据包: 已发送 = 2, 已接收 = 2, 丢失 = 0 (0% 丢失),
    往返行程的估计时间(以毫秒为单位):
        最短 = 24ms, 最长 = 26ms, 平均 = 25ms
Control-C
^C
```

MacOS IPv6 host

```
Headers  Cookies  Request  Response  Timings  Security
Filter Headers
GET
Scheme: https
Host: www.baidu.com
Filename: /
Address: [64:ff9b::6ef2:4403]:443
Status: 200 OK
Version: HTTP/1.1
Transferred: 110.51 kB (470.62 kB size)
Request Priority: Highest
Response Headers (846 B)
```

Android IPv6 host

```
$ ping6 www.baidu.com
PING www.baidu.com(64:ff9b::6ef2:4404) 56 data bytes
64 bytes from 64:ff9b::6ef2:4404: icmp_seq=1 ttl=42 time
=18.7 ms
64 bytes from 64:ff9b::6ef2:4404: icmp_seq=2 ttl=42 time
=43.7 ms
64 bytes from 64:ff9b::6ef2:4404: icmp_seq=3 ttl=42 time
=19.2 ms
64 bytes from 64:ff9b::6ef2:4404: icmp_seq=4 ttl=42 time
=30.1 ms
64 bytes from 64:ff9b::6ef2:4404: icmp_seq=5 ttl=42 time
=17.2 ms
64 bytes from 64:ff9b::6ef2:4404: icmp_seq=6 ttl=42 time
=20.5 ms
64 bytes from 64:ff9b::6ef2:4404: icmp_seq=7 ttl=42 time
=38.7 ms
64 bytes from 64:ff9b::6ef2:4404: icmp_seq=8 ttl=42 time
=29.5 ms
64 bytes from 64:ff9b::6ef2:4404: icmp_seq=9 ttl=42 time
=36.1 ms
^C
--- www.baidu.com ping statistics ---
9 packets transmitted, 9 received, 0% packet loss, time
8021ms
rtt min/avg/max/mdev = 17.247/28.249/43.798/9.263 ms
$
```

(termux)

Next Steps

- Collect comments and suggestions, make further improvement to this document.
- Plan to demo in IETF 118.

Thank you !
Q&A