

EVN6: A Framework of Mapping of Ethernet Virtual Network to IPv6 Underlay draft-xls-Intarea-evn6-00

Chongfeng Xie(Presenter) China Telecom

Xing Li CERNET Center/Tsinghua University

Congxiao Bao CERNET Center/Tsinghua University

Mark Smith

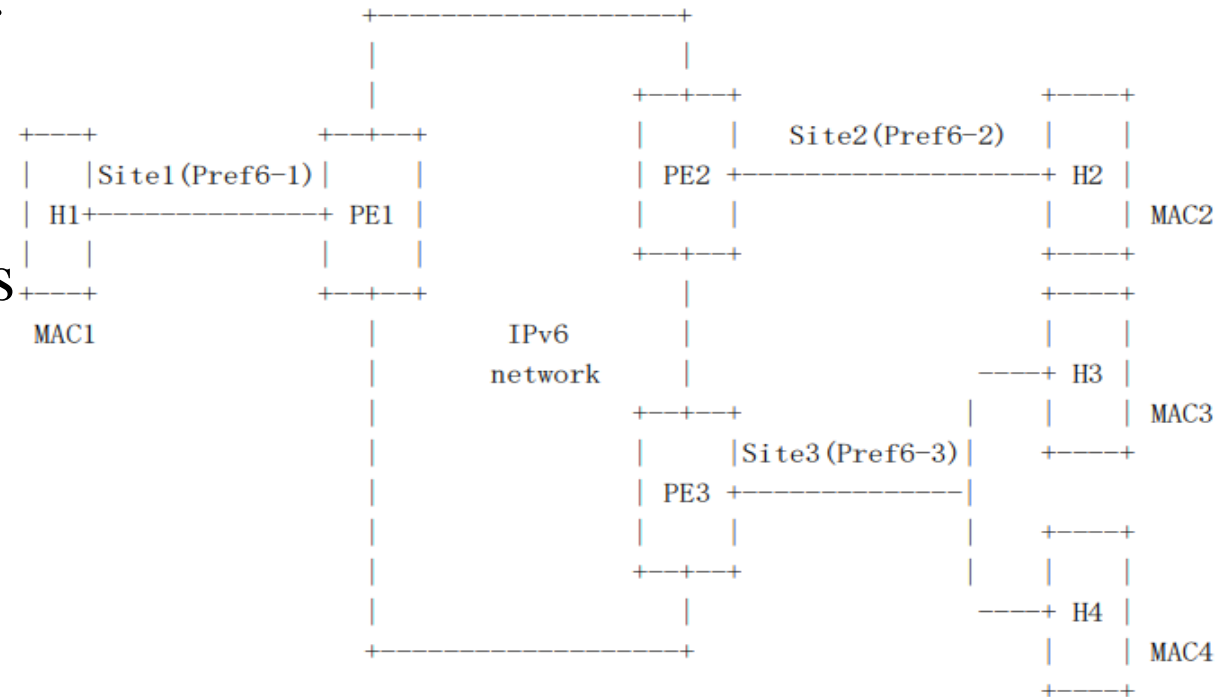
Jibin Sun China Telecom

July 2024

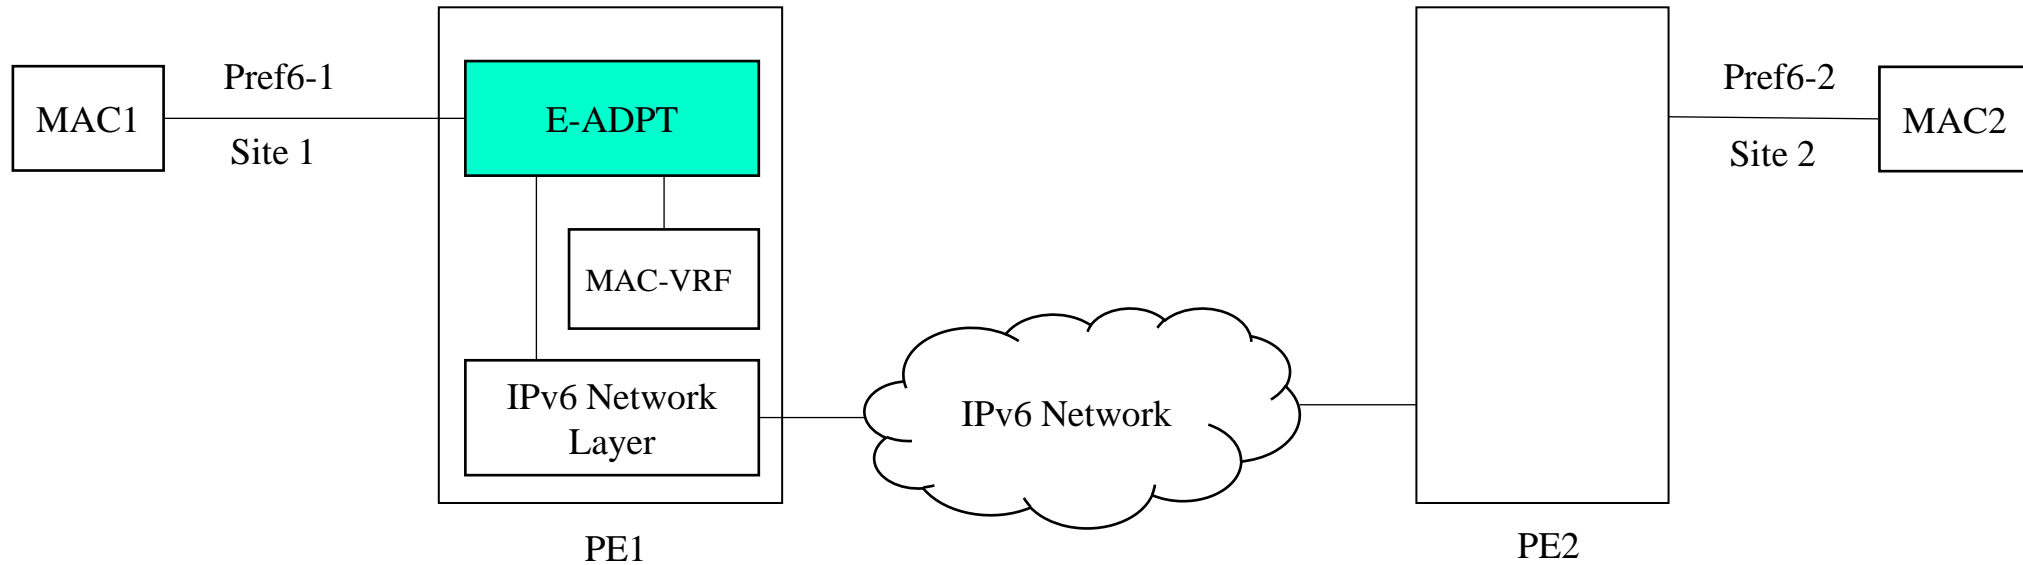
Overview

- EVN6 is Layer-2 network model built on top of the IPv6 underlay to provide connectivity between dispersed customer sites.
 - Ethernet frame is directly copied in the payload of IPv6 packet with encapsulation.
 - Stateless mapping is used to generate IPv6 addresses from the MAC addresses, Virtual Network Identifier and site prefixes.

- This draft was firstly proposed in October 2023 to 6man WG, it has been discussed in 6man and v6ops several rounds.



Overall Architecture



VEI: 32-bits Virtual Ethernet Identifier

Pref6: Site Prefix/ IPv6 Mapping Prefix


MAC-VRF:

| |
|---|
| +-----+ |
| MAC Address VEI Length of Pref6 Pref6 |
| +-----+ |

Pref6:Site Prefix

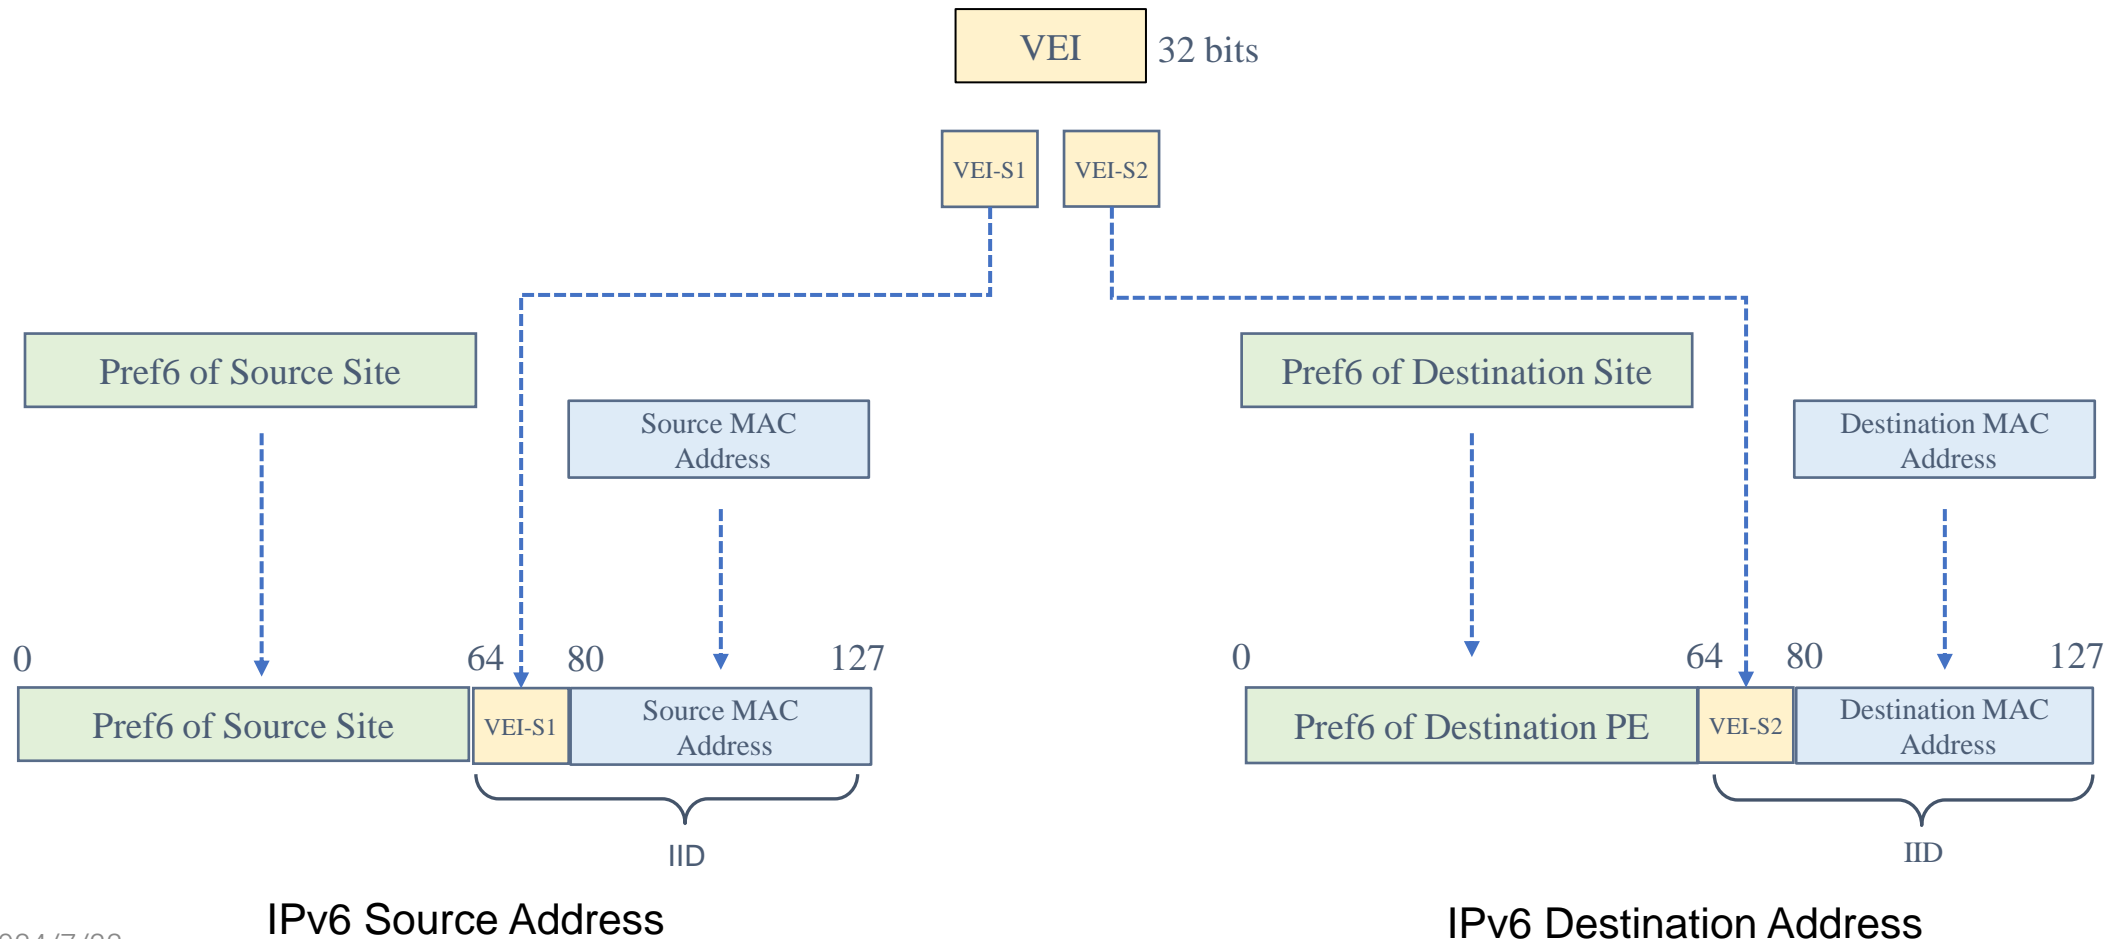
Encapsulation of Ethernet Frame into IPv6 Packet

- The Ethernet frame is directly placed in the payload of IPv6 packet by the E-ADPT of ingress PE.
- The field of 'Next header' in IPv6 header is set as 143, to indicate that the payload is Ethernet frame.



IPv6 Address Generation by Mapping

Stateless mapping is used to generate IPv6 addresses from the MAC addresses, Virtual Network Identifier and site prefixes.



MAC-VRF and Control Layer

- The data in MAC-VRF should be available before encapsulation of Ethernet frame, so the sites needs to pre-send MAC/Pref6 mapping of each hosts to other sites.
- The exchange of MAC/Pref6 mapping can be implemented at the existing control layer. This has been out of the scope of this document.

Revisions made since IETF 119(1/2)

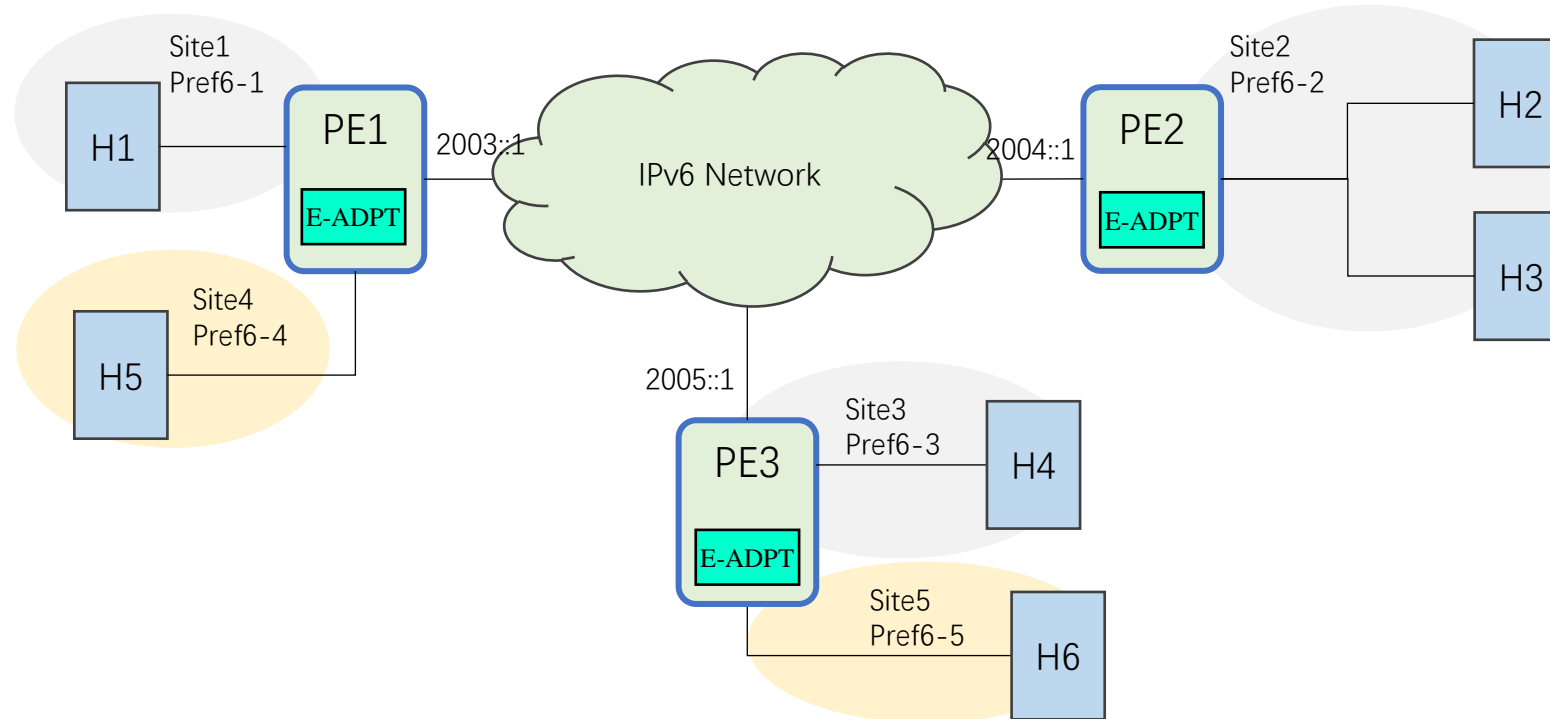
- Based on the suggestion of Eric Vyncke, multicast in EVN6 is added in section 5.1, which illustrates how to process multicast traffic.
- The section of broadcast is revised in section 5.2.
- Jibin Sun is added as one co-author.
- Several editorial changes.

Revisions made since IETF 119(2/2)

Justification and Benefits Analysis

- Forward efficiency
 - L2 frame is encapsulated into IPv6 packet without extra encapsulation headers, encapsulation and processing cost can be reduced.
- Delivery flexibility
 - Service can be provisioned as long as access to IPv6 Internet is available.
 - There is no specific requirement for the interworking between ISPs, it can be easily deployed in multi-operator environment
- Secure
 - As there is no pre-configured static tunnel endpoint address, the risk of DDOS attack can be reduced.
- Traffic load-balancing(ECMP)
 - Load balancing can be implemented based on the source IPv6 addresses, different host within the same site has different IPv6 addresses.

System Implementation and Lab Test



Eth. Virtual Network 1 VEI: 1000

| Site Mapping Prefix: | Host Address: |
|----------------------|----------------|
| Pref6-1: 2001::/64 | H1: 2401::1/64 |
| Pref6-2: 2002::/64 | H2: 2401::2/64 |
| Pref6-3: 2003::/64 | H3: 2401::3/64 |
| | H4: 2401::4/64 |

Eth. Virtual Network 2 VEI: 2000

| Site Mapping Prefix: | Host Address: |
|----------------------|----------------|
| Pref6-4: 2004::/64 | H5: 2402::1/64 |
| Pref6-5: 2005::/64 | H6: 2402::2/64 |

Show MAC-VRF/MMIDs of PE1, PE2 and PE3

PE1

```
root@ivi:~/install# ./evn6_ctrl -o 3 -d evn61
Prefix6:2002::/64 VEI:1000 MAC:00:00:00:00:00:00
Prefix6:2003::/64 VEI:1000 MAC:00:00:00:00:00:00
Prefix6:2003::/64 VEI:1000 MAC:02:22:33:11:22:04
Prefix6:2002::/64 VEI:1000 MAC:02:22:33:11:22:03
Prefix6:2002::/64 VEI:1000 MAC:02:22:33:11:22:02
root@ivi:~/install#
```

PE2

```
root@ivi:~/install# ./evn6_ctrl -o 3 -d evn61
Prefix6:2001::/64 VEI:1000 MAC:00:00:00:00:00:00
Prefix6:2003::/64 VEI:1000 MAC:00:00:00:00:00:00
Prefix6:2003::/64 VEI:1000 MAC:02:22:33:11:22:04
Prefix6:2001::/64 VEI:1000 MAC:02:22:33:11:22:01
root@ivi:~/install#
```

PE3

```
root@ivi:~/install# ./evn6_ctrl -o 3 -d evn61
Prefix6:2001::/64 VEI:1000 MAC:00:00:00:00:00:00
Prefix6:2002::/64 VEI:1000 MAC:00:00:00:00:00:00
Prefix6:2002::/64 VEI:1000 MAC:02:22:33:11:22:03
Prefix6:2002::/64 VEI:1000 MAC:02:22:33:11:22:02
Prefix6:2001::/64 VEI:1000 MAC:02:22:33:11:22:01
root@ivi:~/install#
```

Ping Host of the Same EVN6 Instance

ping H2, H3, H4 from H1

```
root@ivi:~/install# docker exec H1 ping 2401::2
PING 2401::2(2401::2) 56 data bytes
64 bytes from 2401::2: icmp_seq=1 ttl=64 time=0.571 ms
^C
-----
root@ivi:~/install# docker exec H1 ping 2401::3
PING 2401::3(2401::3) 56 data bytes
64 bytes from 2401::3: icmp_seq=1 ttl=64 time=0.537 ms
^C
-----
root@ivi:~/install# docker exec H1 ping 2401::4
PING 2401::4(2401::4) 56 data bytes
64 bytes from 2401::4: icmp_seq=1 ttl=64 time=0.543 ms
64 bytes from 2401::4: icmp_seq=2 ttl=64 time=0.561 ms
^C
root@ivi:~/install#
```

File Transfer within EVN6 Instance

Copy File from H1 to H2 and H4

```
root@b96f30b59a4c:~# echo helloworld_test > test
root@b96f30b59a4c:~# scp test docker@[2401::2]:
docker@2401::2's password:
test
100% 15 21.1KB/s 00:00
root@b96f30b59a4c:~# scp test docker@[2401::4]:
The authenticity of host '2401::4 (2401::4)' can't be established.
ED25519 key fingerprint is SHA256:CoqpQRKhEhC/qs4n88Um+QBvoriwOKf7KNQjXNYOVyU.
This host key is known by the following other names/addresses:
  ~/.ssh/known_hosts:1: [hashed name]
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '2401::4' (ED25519) to the list of known hosts.
docker@2401::4's password:
test
100% 15 24.2KB/s 00:00
root@b96f30b59a4c:~# █
```

Snapshot of Frame-mapped IPv6 Packet

| | | | | | |
|------|-----------|---------|---------|--------|--|
| 1013 | 17.968350 | 2401::5 | 2401::4 | TCP | 148 41704 → 22 [SYN] Seq=0 Win=64800 Len=0 MSS=1440 SACK_PERM TSval=2921916548 TSecr=0 WS=128 |
| 1014 | 17.968811 | 2401::4 | 2401::5 | TCP | 148 22 → 41704 [SYN, ACK] Seq=0 Ack=1 Win=64260 Len=0 MSS=1440 SACK_PERM TSval=2821078719 TSecr=2921916549 |
| 1015 | 17.968877 | 2401::5 | 2401::4 | TCP | 140 41704 → 22 [ACK] Seq=1 Ack=1 Win=64896 Len=0 TSval=2921916549 TSecr=2821078719 |
| 1016 | 17.969088 | 2401::5 | 2401::4 | SSHv2 | 181 Client: Protocol (SSH-2.0-OpenSSH_8.9p1 Ubuntu-3ubuntu0.7) |
| 1017 | 17.969435 | 2401::4 | 2401::5 | TCP | 140 22 → 41704 [ACK] Seq=1 Ack=42 Win=64256 Len=0 TSval=2821078720 TSecr=2921916549 |
| 1019 | 18.017993 | 2401::4 | 2401::5 | ICMPv6 | 172 Echo (ping) request id=0x000a, seq=4400, hop limit=64 (reply in 1020) |
| 1020 | 18.018043 | 2401::5 | 2401::4 | ICMPv6 | 172 Echo (ping) reply id=0x000a, seq=4400, hop limit=64 (request in 1019) |
| 1021 | 18.021660 | 2401::4 | 2401::5 | SSHv2 | 181 Server: Protocol (SSH-2.0-OpenSSH_8.9p1 Ubuntu-3ubuntu0.7) |
| 1022 | 18.021697 | 2401::5 | 2401::4 | TCP | 140 41704 → 22 [ACK] Seq=42 Ack=42 Win=64896 Len=0 TSval=2921916602 TSecr=2821078772 |
| 1023 | 18.021942 | 2401::5 | 2401::4 | SSHv2 | 1676 Client: Key Exchange Init |
| 1024 | 18.024693 | 2401::4 | 2401::5 | SSHv2 | 1252 Server: Key Exchange Init |
| 1025 | 18.024967 | 2401::5 | 2401::4 | SSHv2 | 188 Client: Elliptic Curve Diffie-Hellman Key Exchange Init |


```

> Frame 1017: 140 bytes on wire (1120 bits), 140 bytes captured (1120 bits)
> Ethernet II, Src: HanzsungTech_08:3c:27 (04:2b:58:08:3c:27), Dst: HanzsungTech_08:3b:85 (04:2b:58:08:3b:85)
✓ Internet Protocol Version 6, Src: 2003::3e8:222:3311:2204, Dst: 2001::222:3311:2201
  0110 .... = Version: 6
  > .... 0000 0000 .... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
  .... 0000 0000 0000 0000 = Flow Label: 0x000000
  Payload Length: 86
  Next Header: Ethernet (143)
  Hop Limit: 63
  Source Address: 2003::3e8:222:3311:2204
  Destination Address: 2001::222:3311:2201
  [Destination Teredo Server IPv4: 0.0.0.0]
  [Destination Teredo Port: 64989]
  [Destination Teredo Client IPv4: 204.238.221.254]
✓ Ethernet II, Src: MS-NLB-PhysServer-32_02:33:11:22:04 (02:22:33:11:22:04), Dst: MS-NLB-PhysServer-32_02:33:11:22:01 (02:22:33:11:22:01)
  > Destination: MS-NLB-PhysServer-32_02:33:11:22:01 (02:22:33:11:22:01)
  > Source: MS-NLB-PhysServer-32_02:33:11:22:04 (02:22:33:11:22:04)
    Type: IPv6 (0x86dd)
  > Internet Protocol Version 6, Src: 2401::4, Dst: 2401::5
  > Transmission Control Protocol, Src Port: 22, Dst Port: 41704, Seq: 1, Ack: 42, Len: 0
  
```

| | |
|------|---|
| 0000 | 04 2b 58 08 3b 85 04 2b 58 08 3c 27 86 dd 60 00 |
| 0010 | 00 00 00 56 8f 3f 20 03 00 00 00 00 00 00 03 e8 |
| 0020 | 02 22 33 11 22 04 20 01 00 00 00 00 00 00 00 00 |
| 0030 | 02 22 33 11 22 01 02 22 33 11 22 01 02 22 33 11 |
| 0040 | 22 04 86 dd 60 05 4b 3d 00 20 06 40 24 01 00 00 |
| 0050 | 00 00 00 00 00 00 00 00 00 00 00 04 24 01 00 00 |
| 0060 | 00 00 00 00 00 00 00 00 00 00 00 05 00 16 a2 e8 |
| 0070 | b2 5f 71 90 0f 6f 11 ad 80 10 01 f6 c7 1c 00 00 |
| 0080 | 01 01 08 0a a8 26 3e c0 ae 28 e8 85 |

Test of Isolation between different EVN6 Instances: Ping H6 (VEI 2000) from H1 (VEI 1000)

```
root@ivi:~# docker exec -it H1 bash
root@b96f30b59a4c:/# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
37: eth0@if38: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default
    link/ether 02:22:33:11:22:01 brd ff:ff:ff:ff:ff:ff link-netnsid 0
    inet 172.18.0.2/16 brd 172.18.255.255 scope global eth0
        valid_lft forever preferred_lft forever
    inet6 2401::5/64 scope global nodad
        valid_lft forever preferred_lft forever
    inet6 fe80::22:33ff:fe11:2201/64 scope link
        valid_lft forever preferred_lft forever
root@b96f30b59a4c:/# ping 2402::2
PING 2402::2(2402::2) 56 data bytes
From 2401::5 icmp_seq=1 Destination unreachable: Address unreachable
From 2401::5 icmp_seq=2 Destination unreachable: Address unreachable
From 2401::5 icmp_seq=3 Destination unreachable: Address unreachable
C
--- 2402::2 ping statistics ---
5 packets transmitted, 0 received, +3 errors, 100% packet loss, time 4075ms
```

Next Steps

- Comments and suggestions are welcome, and make further refinement to improve the document.
- Authors would like to ask for WG adoption of this document.

Thank you!
Q&A