

Arista VARP, IPv6 and fall-out

Henrik Lund Kramshøj hlk@nordu.net



NORDUnet is a collaboration between the National Research and Education Networks (NRENs) of the five Nordic countries

Although HSRP and VRRP provide redundancy, they are active-standby FHRR protocols and do not provide a balanced data traffic distribution over Multi Chassis Link Aggregated topologies.

- Virtual Router Redundancy Protocol (VRRP) RFC-5798 is not good enough
- Arista VARP is great

Source for quote:

https://eos.arista.com/active-active-router-redundancy-using-varp/

NORDUnet Overview VARP



How does it work?

- Router 1 has MAC address aa:aa:aa:aa:aa:aa
- Router 2 has MAC address aa:aa:aa:aa:aa:aa
- These MACs are used for the default gateway IP addresses
- Nearest switch/router pick up traffic, and forwards it on Layer 3

NORDUnet How to configure Arista VARP

Site 1:

ip virtual-router mac-address aa:aa:aa:aa:aa:aa interface Vlan123 description NDN-Testnet ip address 109.105.123.2/25 ip virtual-router address 109.105.123.1

Site 2:

ip virtual-router mac-address aa:aa:aa:aa:aa:aa:aa interface Vlan123 description NDN-Testnet ip address 109.105.123.3/25 ip virtual-router address 109.105.123.1 ipv6 address 2001:948:123:4::3/64 ipv6 virtual-router address 2001:948:123:4::1

Yeah, sorry site 1 uses 7150 and did not support IPv6 like we wanted or something ③



- NORDUnet has used Arista switches for some time, before I started
- At some point it was decided to replace Arista 7150 with Arista 7050SX
- something to do with IPv6 support
- Don't mention the war
- We only needed to upgrade hardware in one site, Site 1
- We planned and announced service windows, like usual

NORDUnet Service Window What Happened

- We backed up configs, always nice
- Upgraded hardware 7150 to 7050 very similar config, actually the same
- Made sure the config was in place
- Moved cables, turned up services one-by-one

Everything went smoothly, ... but it wasn't

- Lost IPv6 connectivity to hosts which should send from site 1 to site 2
- We didn't change config, less changes in one go
- We didn't change virtual addresses, we didn't add the IPv6 virtual address to site 1
- We could ping6 the switch, 2001:948:123:4::3 so IPv6 packets go through, WTF

NORDUnet Further debugging

- Client hosts had the right MAC / ARP entries
- IPv4 gateway ARP
- IPv6 gateway NDP, neighbor discovery protocol
- Something was obviously working, but no outside connectivity

The MAC was aa:aa:aa:aa:aa:aa 🙂

Suspicion arose, if the new switch has better IPv6 support it might eat the packets!

Verify and fix

- How to verify?
- Easy add the virtual IPv6 address to the new switches, 7050SX
- Yay! Everything works
- Next, start a support case to confirm with Arista, maybe get them to update documentation?



Arista support topology



<u>Test 1:</u>

NORDUnet Arista support debuggging

Ping results: Ping to physical IP (2001:2::2) on switch do500 was successful from host up441: up441#ping ipv6 2001:2::2 PING 2001:2::2(2001:2::2) 72 data bytes 80 bytes from 2001:2::2: icmp_seq=1 ttl=64 time=0.372 ms 80 bytes from 2001:2::2: icmp_seq=2 ttl=64 time=0.127 ms 20 However, ping to IPv6 virtual address failed from the host: up441#ping ipv6 2001:2::1 PING 2001:2::1(2001:2::1) 72 data bytes --- 2001:2::1 ping statistics ---5 packets transmitted, 0 received, 100% packet loss, time 4003ms

NORDUnet How to really configure Arista VARP!!!!1111

Site 1:

ip virtual-router mac-address aa:aa:aa:aa:aa:aa:aa interface Vlan123 description NDN-Testnet ip address 109.105.123.2/25 ip virtual-router address 109.105.123.1 ipv6 address 2001:948:123:4::2/64 ipv6 virtual-router address 2001:948:123:4::1

Site 2:

ip virtual-router mac-address aa:aa:aa:aa:aa:aa:aa interface Vlan123 description NDN-Testnet ip address 109.105.123.3/25 ip virtual-router address 109.105.123.1 ipv6 address 2001:948:123:4::3/64 ipv6 virtual-router address 2001:948:123:4::1 Arista support later wrote:

After digging deeper, I found that we are treating this as an expected behavior with the current code implementation. I also found RFE40072: 'Packets to VARP MAC should be switched if SVI is down', which has been already filed to change this behavior and have added Nordunet as an interested customer for this feature request. Since it is expected for the Arista switch to consume packets destined to VMAC if the SVI has not been configured on the VLAN

Arista support was very forthcoming, and nice



- When using Arista VARP and both IPv4 and IPv6
- Make sure to add virtual address on ALL routers
- Was the fault adding IPv6 to site 2 before replacing the last hardware?
- Should we have removed the virtual address in site 1 altogether? only L2 switched there
- Was the fault running with mixed hardware for some time?
- Was the fault NOT adding the IPv6 virtual address when replacing the hardware?
- Should we have taken the Arista training before doing this? (spoiler it does not really cover this specific use-case)

Really, help me, what IS the right answer? ③