

AOS-CX 10.10 Update
June 2022

Anycast SVI Gateway (IPv4 and IPv6)

Vincent Giles
Technical Marketing Engineering

aruba

a Hewlett Packard
Enterprise company



Definitions

Acronyms

- SVI **Switch Virtual Interface** : VLAN L3 interface (https://en.wikipedia.org/wiki/Switch_virtual_interface)
 - VXLAN **Virtual eXtensible LAN**
 - VTEP **VXLAN Tunnel End Point**
 - VSX VTEP VTEP function hosted on a VSX cluster for dual-homing capability
 - non-VTEP VSX Plain VSX used in the context of non-VXLAN network
 - VTEP Active-Gateway **Anycast VTEP Virtual IP address** used for distributed default-gateway among all the VTEPs hosting the same subnet, configured with active-gateway command in the SVI context
 - VSX Active-Gateway **VSX Virtual IP address** used for active-active default-gateway shared between VSX primary and VSX secondary, configured with active-gateway command in the SVI context
- Anycast SVI Gateway Anycast SVI IP address used for VTEP or VSX SVI IP address.
Prerequisite: Must be configured after active-gateway command is entered in the SVI context

Anycast SVI Gateway: SVI IP address identical to Active-Gateway IP address

VSX usage versus VTEP usage

VSX Active-Gateway (non-VTEP)

1. Active-gateway command is used to configure a shared virtual IP address (VIP) and a shared virtual MAC address (VMAC) on the Layer3 VLAN interface on the VSX Primary and the VSX Secondary.
2. Support for both IPv4 and IPv6 addresses.
(same virtual MAC for v4 and v6 active-gw per interface VLAN)
3. Support for anycast SVI IPv4 gateway = active-gateway IPv4, since 10.09.0010.
Support for anycast SVI IPv6 gateway = active-gateway IPv6, **since 10.10, for all VSX platforms.**
4. For non-VTEP VSX, anycast SVI gateway is usually not a key requirement as only 2 additional IP addresses are consumed per subnet (impact on IPAM space is very limited).
5. Unicast SVI IP (SVI IP != active-gateway IP) remains supported and **is the best practice** for troubleshooting purpose (ping from different SVI source IP from VSX primary and secondary).
6. IPv4 GARP is sent every 120 seconds.
IPv6 Router-Advertisement is sent every 200-600 seconds (IPv6 Ethernet broadcast Hello is sent every 180 seconds.)
7. Same active-gateway virtual MAC can be used on all SVIs.

VTEP Active-Gateway (standalone, VSF or VSX)

1. Active-gateway command is used to configure EVPN anycast distributed gateway on multiple VTEPs.
2. Support for both IPv4 and IPv6 addresses.
(same virtual MAC for v4 and v6 active-gw per interface VLAN)
3. Support for anycast SVI IPv4 gateway = active-gateway IPv4, since 10.09.0010.
Support for anycast SVI IPv6 gateway = active-gateway IPv6, since 10.09.0010.
4. Anycast SVI gateway provides great IPAM efficiency. No need to provision a distinct IP address per SVI, per physical switch of the EVPN VXLAN fabrics. **This is the best practice.**
5. Unicast SVI IP (SVI IP != active-gateway IP) remains supported; although dedicated SVI IP per switch may quickly lead to IP Address Management operational issue.
6. IPv4 GARP is sent every 120 seconds.
IPv6 Router-Advertisement is sent every 200-600 seconds.
7. Same active-gateway virtual MAC can be used on all SVIs.

Anycast SVI Gateway IPv6 address

Details & Caveats

- Active-Gateway IPv6 address should be a Link-Local Address if hosts learn the default Gateway via RA messages.
- “**no ipv6 nd suppress-ra**” command must be used for host IPv6 auto-configuration (SLAAC) and for DHCPv6 client that depends on RA messages for DHCPv6 setting (Windows host for instance).
- At least one IPv6 address (GUA, ULA or LLA) must be configured on SVI enabled for IPv6 to be enabled on the VLAN interface (aside the auto-config mode on SVI which is out of scope here).
- “ipv6 nd router-preference high” command is not used in the context of VSX or VTEP Active-Gateway.
- When VLAN interface (SVI) and active-gateway are configured with the same IPv6 address, **Duplicate Address Detection** is disabled on the given SVI. Message will pop-up in CLI if DAD attempts is set. This is a restriction that applies on VTEP and also on non-VTEP VSX.
- ActGW_IPv6=SVI_IPv6 can be IPv6 **Global Unicast Address / Unique Local Address and/or IPv6 Link-Local Address**. (see next slide for options)

VTEP - Anycast SVI IPv6 Gateway Configuration

Recommended IPv6 Option for VTEP and host IPv6 configuration

SVI=GUA/ULA + ActGW=LLA

- VTEP SVI configuration

```
interface vlan 10
  vrf attach VRF1
  ip address 10.1.10.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address 2001:db8:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 fe80:0:1:10::1
  no ipv6 nd suppress-ra
```

VTEP1

```
interface vlan 10
  vrf attach VRF1
  ip address 10.1.10.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address 2001:db8:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 fe80:0:1:10::1
  no ipv6 nd suppress-ra
```

VTEP2

- May work for SLAAC, DHCPv6.
Not for static IPv6 host configuration with GUA gateway (even with static GUA GW in host configuration, most hosts use a LLA-GW).
- If same SVI_IP, duplicate IPv6 address with different MACs.



Not supported.

SVI=ActGW=GUA/ULA

- VTEP SVI configuration

```
interface vlan 10
  vrf attach VRF1
  ip address 10.1.10.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address 2001:db8:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

VTEP1

```
interface vlan 10
  vrf attach VRF1
  ip address 10.1.10.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address 2001:db8:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

VTEP2

- For static IPv6 host configuration and also for SLAAC and DHCPv6.
- Warning:** some host OS may still require LLA to exit subnet even if GUA is configured. As LLA is not configured, each physical VTEP switch uses its own address, **different from other VTEPs** of same subnet.
Not supported: Host connectivity may fail.

recommended

SVI=ActGW=GUA/ULA+LLA

- VTEP SVI configuration

```
interface vlan 10
  vrf attach VRF1
  ip address 10.1.10.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address link-local fe80:0:1:10::1/64
  ipv6 address 2001:db8:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 fe80:0:1:10::1
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

VTEP1

Note*

```
interface vlan 10
  vrf attach VRF1
  ip address 10.1.10.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address link-local fe80:0:1:10::1/64
  ipv6 address 2001:db8:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 fe80:0:1:10::1
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

VTEP2

Note*

- Support for any IPv6 host configuration type.
- * Although SVI auto-generated LLA is enough and SVI static LLA is optional in 10.10 implementation, **it is recommended to configure SVI LLA.**

VTEP - Anycast SVI IPv6 Gateway Configuration

Other invalid options for VTEPs in production (educative purpose for support)

SVI=ActGW=LLA

- VTEP SVI configuration

```
interface vlan 10
  vrf attach VRF1
  ip address 10.1.10.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address link-local fe80:0:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 fe80:0:1:10::1
  no ipv6 nd suppress-ra
```

- Host traffic will stay within subnet and won't get return-routed as there is no IPv6 GUA/ULA configured on SVI.
- **Not supported**

SVI=LLA + ActGW=GUA/ULA

- VTEP SVI configuration

```
interface vlan 10
  vrf attach VRF1
  ip address 10.1.10.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address link-local fe80:0:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

VTEP1

```
interface vlan 10
  vrf attach VRF1
  ip address 10.1.10.1/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address link-local fe80:0:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

VTEP2

- VTEP nodes will source LLA with 2 different MAC addresses.
- **Not supported**

Non-VTEP VSX - Unicast SVI IPv6 and Active-Gateway IPv6 Configuration

Recommended IPv6 Option for VSX and host IPv6 configuration

Anycast SVI

- VSX (non-VTEP) SVI configuration

```
interface vlan 10
  vsx-sync active-gateways
  ip address 10.1.10.2/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address link-local fe80:0:1:10::1/64
  ipv6 address 2001:db8:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 fe80:0:1:10::1
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

primary

Note*

```
interface vlan 10
  vsx-sync active-gateways
  ip address 10.1.10.3/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address link-local fe80:0:1:10::1/64
  ipv6 address 2001:db8:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 fe80:0:1:10::1
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

secondary

Note*

- Supported. Plan for troubleshooting and sourcing ping from a loopback.
- * Although SVI auto-generated LLA is enough and SVI static LLA is optional in 10.10 implementation, **it is recommended to configure SVI LLA.**

Unicast SVI

- VSX (non-VTEP) SVI configuration

```
interface vlan 10
  vsx-sync active-gateways
  ip address 10.1.10.2/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address link-local fe80:0:1:10::2/64
  ipv6 address 2001:db8:1:10::2/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 fe80:0:1:10::1
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

primary

Note*

```
interface vlan 10
  vsx-sync active-gateways
  ip address 10.1.10.3/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  ipv6 address link-local fe80:0:1:10::3/64
  ipv6 address 2001:db8:1:10::3/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  active-gateway ipv6 fe80:0:1:10::1
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```

secondary

Note*

- Best practice for non-VTEP VSX.
- * Although SVI auto-generated LLA is enough and SVI static LLA is optional in 10.10 implementation, **it is recommended to configure SVI LLA.**

recommended

Anycast SVI Gateway Configuration

Duplicate Address Detection Caveat

- The recommended configuration order for Active-gateway IPv6 address same as SVI IPv6 address on both VSX Peers or VTEPs is:
 1. IPv6 active gateway configuration
 2. SVI IPv6 address configuration*i.e. IPv6 active-gateway should be configured first followed by SVI IPv6 address on each VSX peer.*
- If the configuration is applied in a different order it may result in DAD status to be DUPLICATE.
- To remove the DUPLICATE status of the SVI IP address, perform shutdown and no shutdown.

```
interface vlan 10
  vsx-sync active-gateways
  ip address 10.1.10.2/24
  active-gateway ip mac 12:00:00:00:01:00
  active-gateway ip 10.1.10.1
  2 ipv6 address link-local fe80:0:1:10::1/64
  ipv6 address 2001:db8:1:10::1/64
  active-gateway ipv6 mac 12:00:00:00:01:00
  1 active-gateway ipv6 fe80:0:1:10::1
  active-gateway ipv6 2001:db8:1:10::1
  no ipv6 nd suppress-ra
```


Thank you

vincent.giles@hpe.com