



Interoperability with EVPN multi-homing using ESI

Cisco Nexus 9000 switches of second generation (EX model and newer) do not offer full support for EVPN multi-homing.



Note For more information on the EVPN multi-homing functionality, see the [Configuring Multi-Homing](#) chapter.

However, as discussed in the following section, Cisco Nexus 9000 switches can be integrated in the same VXLAN EVPN fabric with switches that fully support the EVPN multi-homing functionality.

This chapter contains these sections:

- [VXLAN EVPN interoperability mechanisms](#) , on page 1
- [EVPN multi-homing supported modes and platform limitations](#), on page 2
- [EVPN ESI multi-homing modes](#), on page 3

VXLAN EVPN interoperability mechanisms

A VXLAN EVPN interoperability mechanism is a network interoperability mechanism that

- enables switches using both reserved and non-reserved ESI values to participate in a VXLAN EVPN fabric
- determines how MAC and IP routes are resolved based on ESI type, and
- supports seamless operation between devices using vPC multi-homing and those using EVPN multi-homing.

Beginning Cisco NX-OS Release 10.2(2)F, EVPN MAC/IP routes (Type 2) with non-reserved and with reserved ESI (0 or MAX-ESI) values are evaluated for forwarding (a functionality usually referred to as "ESI RX"). The definition of the EVPN MAC/IP route resolution is defined in [RFC 7432 Section 9.2.2](#).

- For reserved ESI values (0 or MAX-ESI), MAC/IP route resolution is performed solely using the MAC/IP route (BGP next-hop within Type 2).
- For non-reserved ESI values, MAC/IP route resolution occurs only if an accompanying per-ES Ethernet Auto-Discovery route (Type 1, per-ES EAD) is present.

- MAC/IP route resolution with non-reserved ESI values is supported on Cisco Nexus 9300-EX/FX/FX2/FX3/GX Series switches.

These switches, while continuing to use vPC multi-homing for locally connected devices (as discussed in the previous [Configure vPC Multi-Homing](#) and [Configuring vPC Fabric Peering](#) sections), can coexist in a VXLAN EVPN fabric alongside switches using EVPN multi-homing for local device connectivity. MAC and IP addresses for remote endpoints are learned from remote switches via the EVPN control plane and are assigned multiple next-hop IP addresses (unique VTEP addresses for each switch using EVPN multi-homing).

In a data center network, you can connect some devices locally using vPC multi-homing while other devices rely on EVPN multi-homing. The system learns the MAC and IP addresses of remote devices using control plane messages and resolves their next-hop information according to ESI value and platform capability. This design allows for interoperability between mixed deployments of vPC and EVPN multi-homing.

EVPN multi-homing supported modes and platform limitations

EVPN multi-homing using Ethernet Segment Identifier (ESI) enables high availability and device redundancy in VXLAN fabrics with Cisco Nexus switches. The following table summarizes platform support, modes, and key limitations across NX-OS releases.

Table 1: Supported modes and limitations by platform and NX-OS release

NX-OS Release	Supported Platforms	Supported Mode(s)	Notes & Requirements
up to 10.4(1)F	Cisco Nexus 9300-EX/FX/FX2/FX3/GX/GX2 switches and 9500 switches with 9700-EX/FX/GX line cards	All-active only	Single-active mode not supported; only one active path supported.
10.4(1)F and later	Cisco Nexus 9300-EX/FX/FX2/FX3/GX/GX2 switches and 9500 switches with 9700-EX/FX/GX line cards	All-active and single-active	Coexistence with switches supporting single-active mode introduced.
10.4(2)F and later	Cisco Nexus 9332D-H2R and 93400LD-H1 switches	All-active and single-active	Coexistence with both modes available for these platforms.
10.4(3)F and later	Cisco Nexus 9364C-H1 switches	All-active and single-active	Coexistence with both modes available for these platforms.

Guidelines and requirements

Guidelines and limitations for interoperability with EVPN Multi-Homing using ESI:

- Cisco Nexus-9300 switches do not support EVPN multi-homing connectivity to local devices (both all-active and single-active modes), a functionality referred to as “ESI TX”.
- As remote nodes, Cisco NX-OS devices accept MAC routes from the ESI active node and EAD-ES and EAD-EVI routes from both active and standby ESI nodes. Devices use these routes to calculate primary and backup paths for each endpoint's MAC or IP address. Layer 2 traffic follows the primary path in steady state and switches to the backup path if the primary fails.
- Maintenance mode (GIR) for ESI only supports custom profiles designed to bring down uplinks.

- Configure the **maximum-path** setting under the EVPN address-family on Cisco NX-OS nodes. This enables BGP to select multiple paths for EAD per ES and EAD per EVI routes. Remote VTEPs require this configuration for ESI-RX functionality, supporting multi-homing on ESI-VTEPs.

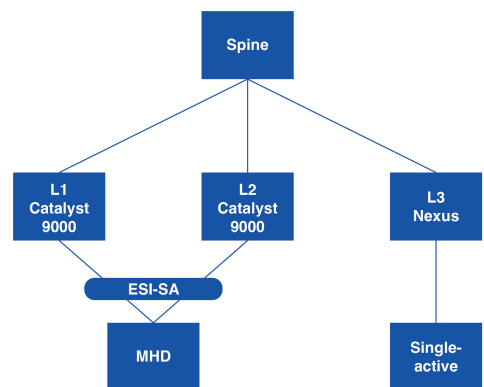
EVPN ESI multi-homing modes

EVPN ESI multi-homing provides redundancy and load sharing in Ethernet VPNs (EVPN) by connecting devices through Ethernet Segment Identifiers (ESIs) that support either single-active or all-active modes. Single-active mode ensures only one node forwards traffic for an ESI, while all-active mode allows multiple nodes to forward simultaneously. You can verify the mode, path details, and MAC address forwarding behavior using show commands on these devices.

In this topology, the Leaf 3 is a Nexus 9000 device which acts as remote VTEP to Catalyst 9000 (Leaf1, Leaf2) devices that support ESI multi-homing connectivity to local devices. It has these capabilities:

- Accepts the MAC, EAD per ES, EAD per EVI routes from ESI-active node and EAD per ES, EAD per EVI routes from ESI-standby node(s).
- Defines whether the ESI is single-active based on flag set in EAD per ES routes.
- Defines whether the ESI single-active is two-way attached or n-way attached based on EAD per ES and EAD per EVI received from how many nodes.

Figure 1: ESI Single-Active Multihoming



Identifying ESI Multi-homing Modes in Operational Outputs

- **ESI Mode Indication:** The field `ESI:1:000000`—a value of 1 means single-active, 0 means all-active.
- **Single-Active Paths:** In MAC address tables, active and standby VTEPs are labeled with `A` (active) and `S` (standby).

For example:

```
nve1(A:192.0.2.11 S:192.0.2.22)
```

- **Path Lists:** The flags (`Si`) identify single-active next-hop paths, and (`A`) identify all-active.
- **Backup Next-Hops:** Path list outputs display backup next-hop addresses for failover scenarios.

Sample operational outputs

• EVPN Route Type Output

```
show bgp l2vpn evpn route-type 1
BGP routing table information for VRF default, address family L2VPN EVPN
Route Distinguisher: 192.0.2.1:3907 (EAD-ES [03de.affe.ed00.0b00.0000 3907])
BGP routing table entry for [1]:[03de.affe.ed00.0b00.0000]:[0xffffffff]/152, version
71
Paths: (1 available, best #1)
Flags: (0x000002) (high32 00000000) on xmit-list, is not in l2rib/evpn

Advertised path-id 1
Path type: local, path is valid, is best path, no labeled nexthop, has esi_gw
AS-Path: NONE, path locally originated
192.0.2.51 (metric 0) from 0.0.0.0 (192.0.2.51)
Origin IGP, MED not set, localpref 100, weight 32768
Received label 0
Extcommunity: RT:12000:1000002 RT:12000:1000003 RT:12000:1000012
RT:12000:1000013 ENCAP:8 ESI:1:000000

Path-id 1 advertised to peers:
192.0.2.2 192.0.2.3
```

ESI:1:000000 -> 1: Indicates single-active.

mode.

• MAC Address Table Output

```
switch# show mac address-table
Legend:
      * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
      age - seconds since last seen, + - primary entry using vPC Peer-Link,
      (T) - True, (F) - False, C - ControlPlane MAC, ~ - vsan,
      (NA) - Not Applicable, A - Active ESI Path, S - Standby ESI Path

VLAN    MAC Address      Type    age    Secure NTFY Ports
-----+-----+-----+-----+-----+-----+-----
C 100    0000.6666.6661    dynamic NA    F      F      nve1 (A:192.0.2.11 S:192.0.2.22)
C 101    0000.6666.6662    dynamic NA    F      F      nve1 (A:192.0.2.11 S:192.0.2.22)
C 101    0000.6666.6663    dynamic NA    F      F      nve1 (A:192.0.2.11 S:192.0.2.22)
C 102    0000.6666.6664    dynamic NA    F      F      nve1 (A:192.0.2.22 S:192.0.2.11)
C 103    0000.6666.6665    dynamic NA    F      F      nve1(192.0.2.33 192.0.2.44)
C 104    0000.6666.6666    dynamic NA    F      F      nve1(192.0.2.33 192.0.2.44)
C 105    0000.6666.6667    dynamic NA    F      F      nve1(192.0.2.33 192.0.2.44)
G -      0091.f3e7.1b08    static  -      F      F      sup-eth1(R)
switch#
```

A is active ESI path; s is standby ESI path.

• L2 Route Path List Output

The following example shows sample output from Leaf 3 device for the **show l2route evpn path-list all detail** command which is enhanced to capture Single-Active mode flag and backup next-hop details as highlighted below:

```
switch# show l2route evpn path-list all detail
(R) = Remote Global EAD NH Peerid resolved,
(UR) = Remote Global EAD NH Peerid unresolved
Flags - (A):All-Active (Si):Single-Active

Topology ID  Prod   ESI                               ECMP Label Flags  Client Ctx  MACs  NFN
-----
1162         None   aaaa.aaaa.aaaa.aaaa.99aa  1           Si      0          1      8
```

```

CP Next-Hops:
Gbl EAD Next-Hops: 192.0.2.11 (11,R), 192.0.2.22 (22,R)
Res Next-Hops: 192.0.2.22
Bkp Next-Hops: 192.0.2.11
Res Next-Hops from UFDM: 192.0.2.22
Bkp Next-Hops from UFDM: 192.0.2.11
1162          UFDM  aaaa.aaaa.aaaa.aaaa.99aa  1          -          1493172225  0          2
CP Next-Hops:
Gbl EAD Next-Hops:
Res Next-Hops: 192.0.2.22
Bkp Next-Hops: 192.0.2.11

switch# show l2route evpn path-list all detail
(R) = Remote Global EAD NH Peerid resolved,
(UR) = Remote Global EAD NH Peerid unresolved
Flags - (A):All-Active (Si):Single-Active

NH Flags: Asy = Asymmetric VNI
Topology ID  Prod  ESI                      ECMP Label          Flags  Client Ctx  MACs
Sent To
-----
100          None  aaaa.aaaa.aaaa.aaaa.99aa  0          A          0          1
          UFDM
CP Next-Hops: 192.0.2.1, 198.51.100.1
Gbl EAD Next-Hops: 192.0.2.1 (192,R)
                  198.51.100.1 (198,R)
Res Next-Hops: <b>192.0.2.1 (Label: 20000) (Flags: Asy)
                198.51.100.1 (Label: 10000) (Flags: Asy)</b>
Bkp Next-Hops:
Res Next-Hops from UFDM: <b>192.0.2.1 (Label: 20000) (Flags: Asy)
                        198.51.100.1 (Label: 10000) (Flags: Asy)</b>

```

- Flags (Si): Single-active path.
- Bkp Next-Hops: Backup next-hop if active fails.

• L2 Route EAD Output

The following example shows sample output for the **show l2route evpn ead all detail** command which is enhanced to capture Single-Active mode flag and backup next-hop details as highlighted below :

```

switch# show l2route evpn ead all detail

Flags - (A):All-Active (Si):Single-Active (V):Virtual ESI (D):Del Pending(S):Stale

Topology ID  Prod  ESI                      NFN Bitmap  Num PLs  Flags
-----
1162          BGP  aaaa.aaaa.aaaa.aaaa.99aa  0          1        -
          Next-Hops: 192.0.2.11, 192.0.2.22
4294967294   BGP  aaaa.aaaa.aaaa.aaaa.99aa  0          1        Si
          Next-Hops: 192.0.2.11, 192.0.2.22

switch# show l2route evpn ead all detail
Flags - (A):All-Active (Si):Single-Active (V):Virtual ESI (D):Del Pending
(S):Stale
Topology ID  Prod  ESI                      Sent To      Num PLs  Flags
-----
100          BGP  aaaa.aaaa.aaaa.aaaa.99aa  -          1        A
          Next-Hops: <b>192.0.2.1 (Label: 20000) (Flags: Asy)</b>
4294967294   BGP  aaaa.aaaa.aaaa.aaaa.99aa  -          1        A
          Next-Hops: 192.0.2.1
                  198.51.100.1

```

Flags Si: Indicates single-active mode.