

Network Configuration Example

EVPN E-LINE/E-LAN Service on ACX7000 Series



Published
2025-01-29

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Network Configuration Example EVPN E-LINE/E-LAN Service on ACX7000 Series

Copyright © 2023 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

CHAPTER 1 EVPN Overview	4
About This Configuration Example	4
Scope	4
Documentation Feedback	4
Technology Primer: EVPN	4
EVPN Overview	4
EVPN-VPWS E-LINE Service Overview	5
EVPN E-LAN	6
EVPN MAC-VRF	8
Underlay and Customer Equipment (CE)	9
ESI	10
CHAPTER 2 Configure EVPN E-LINE/ E-LAN Service on ACX7000 Series	11
EVPN-MPLS E-LINE/VPWS/FXC	11
EVPN-MPLS VPWS	11
EVPN-MPLS VPWS FXC	12
FXC VLAN-Aware Configuration - Active/Active Mode	13
FXC VLAN-Unaware Configuration - Active/Active Mode	14
EVPN-MPLS E-LAN	14
EVPN-MPLS E-LAN VLAN-Based	14
EVPN-MPLS E-LAN VLAN-BASED (Interop Junos-MX vs Evolved) – No-Normalization	16
EVPN-MPLS E-LAN VLAN-BASED (Interop Junos-MX vs Evolved) – Normalization	17
EVPN-MPLS E-LAN VLAN-BASED (Interop Junos-ACX vs Evolved) – Normalization	18
EVPN-MPLS E-LAN VLAN-Bundle	21
EVPN-VXLAN E-LAN	23
EVPN-VXLAN E-LAN VLAN-BASED	23
EVPN-VXLAN E-LAN VLAN-Bundle	23
EVPN-VXLAN E-LAN VLAN-AWARE	24
CHAPTER 3 Common Configuration	27
Underlay Configuration – MPLS/OSPF	27
Underlay Configuration – SR, ISIS	30
MX Series/ACX Series – CE Configuration in EPL, A/A Mode	32
MX Series – EVPL Configuration Single Tag, A/A mode	33
ACX Series – EVPL Configuration Single Tag, A/A Mode	33
MX Series – CE Configuration in EVPL & Dual-Tag, A/A Mode	34
ACX Series – EVPL Configuration Dual Tag, A/A mode	34

CHAPTER 1 EVPN Overview

About This Configuration Example

Scope

Use this network configuration example (NCE) to configure Ethernet Virtual Private Network (EVPN) Ethernet LAN (E-LAN) on ACX7000 Series.

This example covers:

- Provider Edge (PE) techniques and configurations including scaling of the instances.
- A requirement to move to a more controlled and predictable network for the Service providers and Cloud providers.
- Configuring the EVPN service implementation on the ACX7000 Series with Junos OS Evolved.

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation.

Send your comments to design-center-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Technology Primer: EVPN

EVPN Overview

Ethernet EVPN (EVPN) as a technology has matured with large-scale deployment seen across various use case requirements including Service Provider and Cloud. ACX7000 Series support of a wide range of EVPN services that includes E-LINE, E-LAN services with MPLS/SR-MPLS or VXLAN underlay.

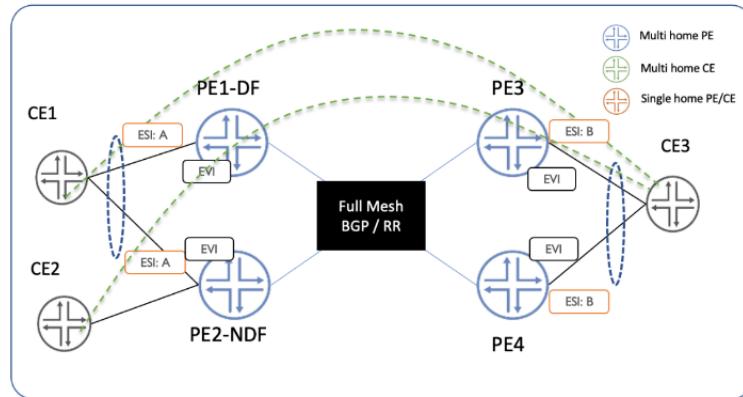
Benefits

EVPN provides advanced L2 deployments with better scale and performance in terms of availability, overcoming underutilization of revenue ports, faster convergence, and advantages over legacy technologies.

Reference Topology

Figure 1 depicts a typical EVPN MPLS architecture.

Figure 1: EVPN Reference Topology – Single Home / Multi Home



EVPN-VPWS E-LINE Service Overview

EVPN-Virtual Private Wired Service (VPWS) is a framework for delivering point-to-point VPN service with the EVPN signaling mechanisms. The endpoints are identified using a service identifier also called VPWS-ID discovered using Border Gateway Protocol (BGP) signaling. The service IDs are 24 bit values encoded in the Ethernet Tag of the Auto Discovery (AD) route /Type 1 per Ethernet Virtual Instance (EVI).

Benefits

- Support for Single homing (SH) and Multi homing (MH) scenario—Advantages in multihoming include load balancing. Both connections are active.
- Support for Inter-AS across the ASBR (Autonomous System Border Router) using the BGP EVPN signaled VPNs.

Private Line Services

EVPN-VPWS delivers the following private line services defined by the Metro Ethernet Forum (MEF):

- Ethernet Private Line (EPL):
 - Offers a point-to-point ethernet virtual connection (EVC) between a dedicated UNI (User Network Interface) with a high degree of transparency.
 - Constructs an EPL-based service based on the ports as per the Metro Ethernet Forum standards. The service is provisioned between a pair of UNI ports. The Ethernet Segment Identifier (ESI) is based on the Ethernet Virtual Instance (EVI) and port. In the multihomed scenario as depicted in Figure 1, the Designate Forwarder (DF) election is Modulus (MOD) based (see [RFC7432](#)). ACX Series supports the preference-based DF election.

- Ethernet Virtual Private Line (EVPL):
 - Offers service multiplexing which means multiple EVCs are paired per UNI.
 - Constructs an EVPL based service based on the IFD as per the Metro Ethernet Forum standards. The service is provisioned between a multiple pair of UNI ports. The ESI is based on the EVI and per IFD. In the multihomed scenario as depicted in [Figure 1](#), the DF election is MOD based.

EVPN E-LAN

Overview

EVPN E-LAN is a framework for delivering multipoint-to-multipoint VPN service with the EVPN signaling mechanisms. E-LAN service allows service providers to offer services that manage the L2 learning very efficiently. In a multihoming scenario, the BUM is handled by the PE-DF, and the learned information is redistributed to other PEs in the network. The multihomed Customer Edge (CE) connects a customer site to two or more PE devices providing redundant services.

From a MEF standard, there are two different services for EVPN E-LAN:

- EP-LAN (Ethernet Private LAN)
 - Offers a multipoint-to-multipoint ethernet virtual connection (EVC) between a dedicated UNI (User Network Interface). EP-LAN is a port-based service.
- EPV-LAN (Ethernet Virtual Private LAN)
 - Offers VLAN based service multiplexing, which means multiple EVCs are paired per UNI.

Instance Types

EVPN E-LAN instance types that are supported across all the Juniper platforms are:

- MAC-VRF—Provides all 3 services (vlan-based, vlan-bundle, vlan-aware-bundle) associated with a single EVI.
- EVPN—Provides vlan-based and vlan-bundle services for an EVI.
- VIRTUAL-SWITCH—Provides vlan-aware-bundle services for an EVI.

ACX Series routers based on Junos OS Evolved support only MAC-VRF instance-type.

EVI Support - MAC-VRF vs EVPN vs VIRTUAL-SWITCH

[Table 1](#) lists different EVIs supported for the MPLS/VXLAN based EVPN implementation on Juniper platforms. This content guides you to identify the Juniper platforms that can interop with ACX Series.

Table 1: EVI Support - MAC-VRF vs EVPN vs VIRTUAL-SWITCH

Instance Type	Platform with OS	VLAN-BASED		VLAN-BUNDLE		VLAN-AWARE	
		MPLS	VXLAN	MPLS	VXLAN	MPLS	VXLAN
MAC-VRF	PTX Series (Junos OS Evolved)	✓	✓	✓	✓	✓	✓
	ACX Series (Junos OS Evolved)	✓	✓	✓	✓	✓	✓

Instance Type	Platform with OS	VLAN-BASED		VLAN-BUNDLE		VLAN-AWARE	
		MPLS	VXLAN	MPLS	VXLAN	MPLS	VXLAN
	MX Series (Junos OS)	×	✓	×	✓	×	✓
EVPN	MX Series (Junos OS)	✓	×	✓	×	×	×
	ACX Series (Junos OS)	✓	×	✓	×	×	×
	PTX Series (Junos OS)	×	×	×	×	×	×
VIRTUAL-SWITCH	MX Series (Junos OS)	×	×	×	×	✓	✓
	ACX Series (Junos OS)	×	×	×	×	×	×
	PTX Series (Junos OS)	×	×	×	×	×	×

EVPN E-LAN Services Behavior Per Instance Type

Table 2 outlines the behavior of each EVPN E-LAN service model for implementation on ACX Series with Junos OS Evolved. This behavior is useful during interop with other Juniper platforms or other Networking Equipment Manufacturers (NEM).

Table 2: Service Comparison MAC-VRF vs EVPN vs VIRTUAL-SWITCH

Services	MAC-VRF	EVPN	Virtual-Switch
Control-word enabled by default	✓	✗	✗
Normalization	✓	✓	✗
no-normalization	✓	✓	✗
Default normalization behavior	No-normalization	Normalization	NA
Full RFC7432 Compliance	✓	✓	✓
Core Isolation	✓	✓	✓

EVPN E-LAN Services Behavior with the SP or EP Style of Configuration

Service Provider (SP) style and Enterprise (EP) have different styles for configuring interfaces using Juniper platforms.

Table 3 highlights the supported options available for configuring vlan-id or vlan-id-list.

Table 3: Configuration SP vs EP Style - MAC-VRF

Services vlan-id vs vlan-id-list	VLAN-BASED		VLAN-BUNDLE		VLAN-AWARE	
	SP Style	EP Style	SP Style	EP Style	SP Style	EP Style
VLAN-id	✓	✗	✓	✗	✓	✗
VLAN-id-list	✗	✗	✓	✗	✗	✗

EVPN MAC-VRF

Overview

MAC-VRF unifies EVPN E-LAN services configuration across all Juniper platforms for EVPN-MPLS or EVPN-VXLAN. With this instance, the CLI `service-type` drives the requirements of the service under a single umbrella for the E-LAN services. These are driven by the service provider requirements, RFC compliance, and design choices.

MAC-VRF is not just about simplicity, it also addresses RFC compliance without any additional CLIs. With `instance-type mac-vrf`, the bridge domains are extended into EVPN by default. There is no need to configure `extended-vlan-list` to extend the BDs in the EVPN. For MPLS encapsulation, only SP Style is supported in MAC-VRF instances.

VLAN-Based Service Models

VLAN-based service is a one-to-one mapping of the VLAN ID association with the EVPN Instance (EVI). The MAC-VRF EVI corresponds one VLAN to a single flooding domain and single MAC learning domain to handle all the BUM and Unicast traffic. In some use cases, VLANs on two ends of the EVPN network are different and, in that case, you need to perform VLAN translation.

Table 4 lists different translation models with examples. The ACX Series are flexible to adhere to the compliance partial or full based on the requirement.

Table 4: VLAN-Based Service Models Options

Options	Description	Route Example
No VLAN Translation, Ethernet tag = 0	This is a strict compliance of RFC7432 Section 6.1. The T2 route consists of Ethernet TAG=0	2:22.22.22.22:1::0:00:0 0:aa:00:00:00/304 MAC/IP
No VLAN Translation, Ethernet-tag = non zero	This is a partial compliance of RFC7432 Section 6.1. The T2 route consists of Ethernet TAG=non zero. In this case, the VLANs facing Access & Core are symmetric.	2:22.22.22.22:1::1:00:0 0:aa:00:00:00/304 MAC/IP
VLAN Translation, Ethernet-tag = non zero	This is a partial compliance of RFC7432 Section 6.1. The T2 route consists of Ethernet TAG=non zero. In this case, the VLANs	2:22.22.22.22:1::1:00:0 0:aa:00:00:00/304 MAC/IP

Options	Description	Route Example
	Facing Access & Core are Asymmetric. A VLAN translation needs to be done.	

With MAC-VRF, the control plane advertises routes with ethernet-tag-id = 0 for vlan-based service by default. To interop with other Juniper platforms (Junos OS based), an EVPN instance is configured with `vlan-id none` and `protocol evpn no-normalization` for routing instances with `instance-type evpn`.

At the same time, routing-instance on other platforms configured using `instance-type evpn` in default mode without `vlan-id none` and `protocol evpn no-normalization`. Then, the MAC-VRF instance is configured with `normalize tag`.

EVPN vs MAC-VRF (Normalization)

[Table 5](#) provides guidelines to Interop between different Juniper platforms implementing VLAN-based service using instance-type EVPN or MAC-VRF.

Table 5: Interop Configuration Guidelines

Feature	EVPN	MAC-VRF
normalization	Default behavior	With CLI under <code>protocol evpn set normalization</code> .
no normalization	Using CLIs, set <code>vlan-id none</code> and <code>set no-normalization</code> .	Default Behavior MAC/IP
control-word (default)	Disabled	Enabled

VLAN Bundle

MAC-VRF instance with service-type “vlan-bundle” offers service with multiple bridge-domain/VLANS under it. The MAC-VRF EVPN Instance (EVI) corresponds to one VLAN to a bridge table. Then, the created bridge domain handles all the BUM and Unicast traffic within the Bridge domain. For more information, see [RFC7432 Section 6.2](#).

When maintaining single MAC learning domain (single MAC table, single bridge domain), you can apply VLAN bundle service to multiple broadcast domains (multiple VLANs). This implies that MAC address must be unique across VLANs. There cannot be MAC overlap between VLANs in the VLAN bundle service. Using VLAN bundle, you can consolidate many similar VLANs to be part of a single BD.

Underlay and Customer Equipment (CE)

Configuration

All PEs are configured using the OSPF and MPLS for underlay. The PEs are configured with fully meshed BGP session with `family evpn signaling` enabled to exchange the EVPN routes among each other. For sample configuration, see [Underlay Configuration – MPLS/OSPF](#). For CE configuration, see [CE Configuration in EPL, A/A Mode](#).

ESI

Configuration

Ethernet Segment Identifier (ESI) identifies multihomed Ethernet segments. Information about ESI in EVPN messages is encoded in NRLI and distributed through ES-Import Route Target (RT) communities (*RFC 7432, Section 7.6*) in Type 4 messages. For example:

```
4:11.11.11.11:0::1111111111111111:11.11.11.11/296 ES (1 entry, 1 announced) <===== encoded
ESI: Type 0, Value: 11:11:11:11:11:11:00:00:00

TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0x56375d6c5030 (adv_entry)
  Advertised metrics:
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [64512] I
    Communities: es-import-target:11-11-11-11-11-11 <===== ES-IMPORT RT
    Advertise: 00000003
Path 4:11.11.11.11:0::1111111111111111:11.11.11.11
```

ESI generates the ES-Import RTs. For example:

```
set interfaces ae0 esi 00:11:11:11:11:11:11:11:11
```

ESI has 10 bytes and ES-Import RT has space for six bytes of payload. Therefore, the first 6 most significant bytes from ESI payload are taken to generate ES-Import RT. In ESI, the first byte is ESI Type. For example, Type 0: manually configured. In the above example, 11:11:11:11:11:11 are taken from configured ESI and used to populate advertised ES-Import RT. The ES-Import RT enables all the PEs connected to the same multi-homed site to import the Ethernet Segment (Type 4) routes.

If the following ESIs are configured on some interfaces, then the ES-Import RT is 00:00:00:11:11:11 with ESI configured as 00:00:00:00:11:11:11:11:11 on PE1 and PE2 in [Figure 2](#). See the following configuration:

```
set interfaces ae0 esi 00:00:00:00:11:11:11:11:11
```

The output from the PE1 with ESIs configuration is as follows:

```
4:11.11.11.11:0::111111111111:11.11.11.11/296 ES (1 entry, 1 announced)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0x56375d6c72c0 (adv_entry) <===== encoded
ESI: Type 0, Value: 00:00:00:11:11:11:11:00:00:00
  Advertised metrics:
    Flags: Nexthop Change
    Nexthop: Self
    Localpref: 100
    AS path: [64512] I
    Communities: es-import-target:0-0-0-11-11-11 <===== ES-IMPORT RT
    Advertise: 00000003
Path 4:11.11.11.11:0::111111111111:11.11.11.11
```

As the configuration is same for all ESIs, all PEs will import Type 4 routes. The route is not used if there is no match based on the actual ESI encoded as part of NRLI in these Type 4 routes and the actual ESI configured on some local interface. This EVPN configuration is not optimal. The most optimal configuration for control plane is to differentiate ESIs in the first six useful bytes instead of using the last three bytes. The generated ES-Import RT prevents the routes from being imported on PEs that do not need them.

CHAPTER 2 Configure EVPN E-LINE/E-LAN Service on ACX7000 Series

This chapter explains the following configuration scenarios:

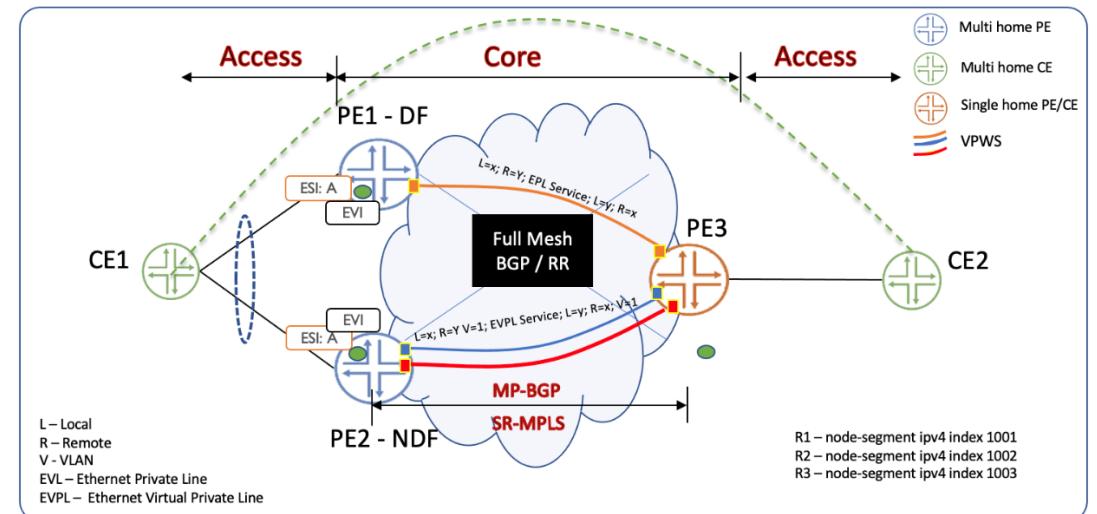
- EVPN-MPLS E-LINE/VPWS/FXC
- EVPN-MPLS E-LAN
- EVPN-VXLAN E-LAN

EVPN-MPLS E-LINE/VPWS/FXC

EVPN-MPLS VPWS

All three PEs in [Figure 1](#) are BGP-EVPN PE. The PEs are set up using the EVPN VPWS instance and the required configurations on all the PEs. The PEs discover each other in the network and establish the EVPN sessions.

Figure 2: EVPN E-LINE/VPWS



VPWS Configuration - Active/Active Mode

The output highlights the EVI and the UNI-facing interface configuration on all three PE nodes shown in [Figure 2](#). This configuration focus on the Active/Active mode.

In Active/Active mode, you must bundle the UNI links from CE to PE for load balancing the traffic. You must make the necessary configuration on the router to achieve the load balance behavior.

Here, you change the configuration to support the Active/Active mode. The configuration shows only to add Active/Active on one of the MH PEs (PE1, PE2). You cannot configure an MH PEs ESI in different modes, for example: PE1 in Active/Standby and PE2 in Active/Active.

```

PE1
r1-re0# show interfaces ae0
ae0 {
    encapsulation ethernet-ccc;
    esi {
        00:11:11:11:11:11:11:11;
        all-active;
    }
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:00:00:00:11:01;
        }
    }
    unit 0 {
        family ccc;
    }
}

r1-re0# show routing-instances
evpn-vpws-evpl {
    instance-type evpn-vpws;
    protocols {
        evpn {
            interface ae0.0 {
                vpws-service-id {
                    local 1;
                    remote 2;
                }
            }
        }
    }
    interface ae0.0;
    route-distinguisher 11.11.11.11:1;
    vrf-target target:11:1;
}

PE2
r2-re0# show interfaces ae0
ae0 {
    encapsulation ethernet-ccc;
    esi {
        00:11:11:11:11:11:11:11;
        all-active;
    }
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:00:00:00:11:01;
        }
    }
    unit 0 {
        family ccc;
    }
}

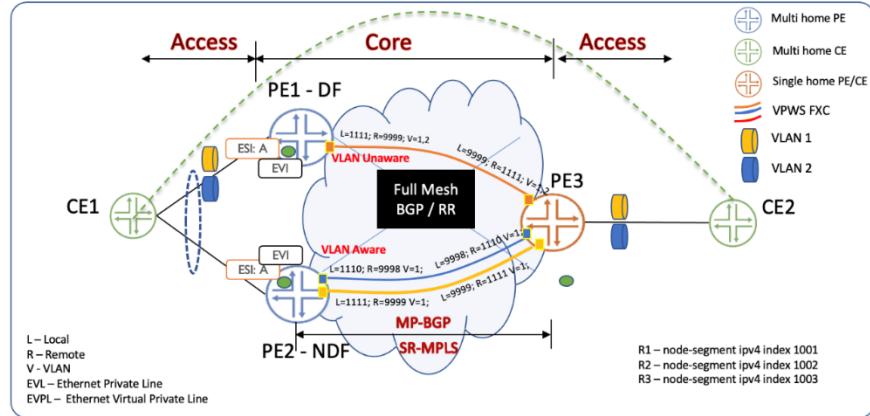
r2-re0# show routing-instances
evpn-vpws-evpl {
    instance-type evpn-vpws;
    protocols {
        evpn {
            interface ae0.0 {
                vpws-service-id {
                    local 1;
                    remote 2;
                }
            }
        }
    }
    interface ae0.0;
    route-distinguisher 22.22.22.22:1;
    vrf-target target:11:1;
}

```

EVPN-MPLS VPWS FXC

All three PEs in [Figure 2](#) are BGP-EVPN PE. The PEs are set up using the EVPN VPWS instance and the required configurations on all the PEs. These PEs discover each other in the network and establish the EVPN sessions.

Figure 3: EVPN VPWS FXC



FXC VLAN-Aware Configuration - Active/Active Mode

The output highlights the EVI and the UNI-facing interface configuration on all three PE nodes (see Figure 3). This configuration focus on the Active/Active mode.

PE1 <pre>r1-re0# show interfaces ae0 ae0 { flexible-vlan-tagging; encapsulation flexible-ethernet-services; esi { 00:11:11:11:11:11:11:11; all-active; } aggregated-ether-options { lACP { active; system-id 00:00:00:00:11:01; } unit 1 { encapsulation vlan-ccc; vlan-id 1; } unit 2 { encapsulation vlan-ccc; vlan-id 2; } } r1-re0# show routing-instances evpn-vpws-evpl { instance-type evpn-vpws; protocols { evpn { interface ae0.1 { vpws-service-id { local 1111; remote 9991; } } interface ae0.2 { vpws-service-id { local 1112; remote 9992; } } flexible-cross-connect-vlan-aware; } } interface ae0.1; interface ae0.2; route-distinguisher 11.11.11.11:1; vrf-target target:11:1; }</pre>	PE3 <pre>r2-re0# show interfaces ae0 ae0 { flexible-vlan-tagging; encapsulation flexible-ethernet-services; aggregated-ether-options { lACP { active; } unit 1 { encapsulation vlan-ccc; vlan-id 1; } unit 2 { encapsulation vlan-ccc; vlan-id 2; } } r2-re0# show routing-instances evpn-vpws-evpl { instance-type evpn-vpws; protocols { evpn { interface ae0.1 { vpws-service-id { local 9991; remote 1111; } } interface ae0.2 { vpws-service-id { local 9992; remote 1112; } } flexible-cross-connect-vlan-aware; } } interface ae0.1; interface ae0.2; route-distinguisher 33.33.33.33:1; vrf-target target:11:1; }</pre>
---	--

FXC VLAN-Unaware Configuration - Active/Active Mode

The output highlights the EVI and the UNI-facing interface configuration on all three PE nodes shown in [Figure 3](#). This configuration focus on the Active/Active mode.

```

PE1
r1-re0# show interfaces ae0
ae0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    esi {
        00:11:11:11:11:11:11:11;
        all-active;
    }
    aggregated-ether-options {
        lACP {
            active;
            system-id 00:00:00:00:11:01;
        }
        unit 1 {
            encapsulation vlan-ccc;
            vlan-id 1;
        }
        unit 2 {
            encapsulation vlan-ccc;
            vlan-id 2;
        }
    }
r1-re0# show routing-instances
evpn-vpws-evpl {
    instance-type evpn-vpws;
    protocols {
        evpn {
            flexible-cross-connect-vlan-unaware;
            group fxc {
                interface ae0.1;
                interface ae0.2;
                service-id {
                    local 9999;
                    remote 1111;
                }
            }
        }
    }
    route-distinguisher 11.11.11.11:1;
    vrf-target target:11:1;
}

PE3
r2-re0# show interfaces ae0
ae0 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lACP {
            active;
        }
        unit 1 {
            encapsulation vlan-ccc;
            vlan-id 1;
        }
        unit 2 {
            encapsulation vlan-ccc;
            vlan-id 2;
        }
    }
r2-re0# show routing-instances
evpn-vpws-evpl {
    instance-type evpn-vpws;
    protocols {
        evpn {
            flexible-cross-connect-vlan-unaware;
            group fxc {
                interface ae0.1;
                interface ae0.2;
                service-id {
                    local 9999;
                    remote 1111;
                }
            }
        }
    }
    route-distinguisher 33.33.33.33:1;
    vrf-target target:11:1;
}

```

EVPN-MPLS E-LAN

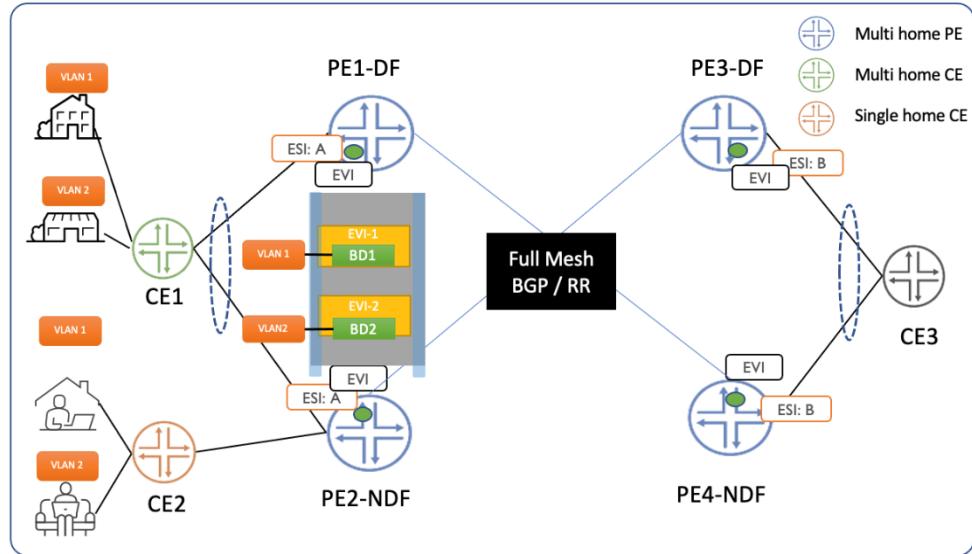
EVPN-MPLS E-LAN VLAN-Based

All four PEs in [Figure 2](#) are BGP-EVPN PE. The PEs are set up using the EVPN MAC-VRF instance and the required configurations on the PEs. These PEs discover each other in the network, establish the EVPN sessions, and exchange the EVPN routes. With VLAN based service, there is one VLAN associated with an EVI. This VLAN has one-to-one mapping of the VLAN-ID and the EVI. So, one bridge table for one VLAN.

[Table 1](#) highlights the different VLAN based models, which are mentioned in the RFC 7432. You can take actions in the Ingress router based on your requirements.

For more information on the VLAN-based service, see [VLAN-Based Service for EVPN](#).

Figure 4: EVPN E-LAN VLAN-BASED



E-LAN (VLAN-Based) Configuration - Active/Active Mode

In the Active/Active mode, you need to bundle the UNI links from CE node to the multihomed PE for load balancing the traffic.

```

PE1
r1-RE0-re0# show interfaces ae0
flexible-vlan-tagging;
encapsulation flexible-ethernet-services;
esi {
    00:11:11:11:11:11:11:11;
    all-active;
}
aggregated-ether-options {
    lACP {
        active;
        system-id 00:00:00:00:11:01;
    }
}
unit 1 {
    encapsulation vlan-bridge;
    vlan-id 1;
}

r1-RE0-re0# show routing-instances
evpn-vlan-based {
    instance-type mac-vrf;
    protocols {
        evpn {
            encapsulation mpls;
        }
    }
    service-type vlan-based;
    route-distinguisher 11.11.11.11:1;
    vrf-target target:11:1;
    vlans {
        bd1 {
            vlan-id 1;
            interface ae0.1;
        }
    }
}

PE3
r3-RE0-re0# show interfaces ae0
flexible-vlan-tagging;
encapsulation flexible-ethernet-services;
esi {
    00:22:22:22:22:22:22:22;
    all-active;
}
aggregated-ether-options {
    lACP {
        active;
        system-id 00:00:00:00:22:01;
    }
}
unit 1 {
    encapsulation vlan-bridge;
    vlan-id 1;
}

r2-RE0-re0# show routing-instances
evpn-vlan-based {
    instance-type mac-vrf;
    protocols {
        evpn {
            encapsulation mpls;
        }
    }
    service-type vlan-based;
    route-distinguisher 33.33.33.33:1;
    vrf-target target:11:1;
    vlans {
        bd1 {
            vlan-id 1;
            interface ae0.1;
        }
    }
}

```

EVPN-MPLS E-LAN VLAN-BASED (Interop Junos-MX vs Evolved) – No-Normalization

All the four PEs in [Figure 2](#) are BGP-EVPN PE. The PEs on one side are Junos OS EVO based platforms configured with MAC-VRF instance and on the other side are Junos OS based MX Series configured with EVPN instance type. This configuration is RFC compliant.

The PEs discover each other in the network, establish the EVPN sessions, and exchange the EVPN routes. With VLAN-based service, there is one VLAN associated with an EVI with a one-to-one mapping of the VLAN-ID and the EVI. So, one bridge table for one VLAN.

[Table 1](#) highlights the different VLAN based models, which are mentioned in the RFC 7432. You can take actions in the Ingress router based on your requirements.

E-LAN (VLAN-Based Interop Junos-MX vs Evolved → No-Normalization) Configuration – A/A Mode

In the Active/Active mode, you must bundle the UNI links from CE node to the multihomed PE for load balancing the traffic. In this case, you are provisioning VLAN-based service with no-normalization. The CEs connected on either side of the network are part of same or different VLANs. Ensure that the traffic is forwarded end-to-end. The RFC7432 section 6.1 mandates that the Ingress PE cannot modify the VLAN information coming from the CE (also known as no-normalization). The PE1 or PE3 preserves the Ethernet tag information when the traffic passes through the WAN. The control plane generates the Type2 (MAC/IP) route, and Type3 (IM) route Ethernet tag set to 0. The following configuration output is used for the Interop requirement.

<pre> r1-RE0-re0# show interfaces ae0 flexible-vlan-tagging; encapsulation flexible-ethernet-services; esi { 00:11:11:11:11:11:11:11:11; all-active; } aggregated-ether-options { lacp { active; system-id 00:00:00:00:11:01; } } unit 1 { encapsulation vlan-bridge; vlan-id 1; } r1-RE0-re0# show routing-instances evpn-vlan-based { instance-type mac-vrf; protocols { evpn { no-control-word; encapsulation mpls; } } service-type vlan-based; route-distinguisher 11.11.11.11:1; vrf-target target:11:1; vlans { bdl { vlan-id 1; interface ae0.1; } } } </pre>	<pre> PE3 (MX-Junos) r2-RE0-re0# show interfaces ae0 flexible-vlan-tagging; encapsulation flexible-ethernet-services; esi { 00:22:22:22:22:22:22:22:22; all-active; } aggregated-ether-options { lacp { active; system-id 00:00:00:00:22:01; } } unit 1 { encapsulation vlan-bridge; vlan-id 1; } r2-RE0-re0# show routing-instances evpn-vlan-based { instance-type evpn; protocols { evpn { no-control-word; encapsulation mpls; } } Vlan-id none; no-normalization; route-distinguisher 22.22.22.22:1; vrf-target target:11:1; interface ae0.1; } </pre>
--	--

EVPN-MPLS E-LAN VLAN-BASED (Interop Junos-MX vs Evolved) – Normalization

All four PEs in [Figure 2](#) are BGP-EVPN PE, where two PEs configured on one side are Junos OS Evolved based platforms with MAC-VRF instance and on the other side are Junos OS based MX Series configured with EVPN instance type. This configuration is partially RFC compliant.

The PEs discover each other in the network, establish the EVPN sessions, and exchange the EVPN routes. With VLAN-based service, there is one VLAN associated with an EVI with a one-to-one mapping of the VLAN-ID and the EVI. So, one bridge table for one VLAN.

[Table 1](#) highlights the different VLAN based models, which are mentioned in the RFC 7432. You can take actions in the Ingress router based on your requirements.

E-LAN (VLAN-Based Interop Junos-MX vs Evolved → Normalization) Configuration – A/A Mode

In the Active/Active mode, you must bundle the UNI links from CE node to the multihomed PE for load balancing the traffic.

In this case, you are provisioning VLAN-based service with normalization. The CEs connected on either side of the network are part of different VLANs. Ensure that the traffic is forwarded. This is not RFC compliant because the VLAN information is normalized in the ingress PE. The following configuration output is used for the Interop requirements.

<pre> PE1 r1-RE0-re0# show interfaces ae0 flexible-vlan-tagging; encapsulation flexible-ethernet-services; esi { 00:11:11:11:11:11:11:11:11; all-active; } aggregated-ether-options { lACP { active; system-id 00:00:00:00:11:01; } } unit 1 { encapsulation vlan-bridge; vlan-id 20; } r1-RE0-re0# show routing-instances evpn-vlan-based { instance-type mac-vrf; protocols { evpn { normalization; no-control-word; encapsulation mpls; } } service-type vlan-based; route-distinguisher 11.11.11.11:1; vrf-target target:11:1; vlangs { bdi { vlan-id 20; interface ae0.1; } } } </pre>	<pre> PE3 r2-RE0-re0# show interfaces ae0 flexible-vlan-tagging; encapsulation flexible-ethernet-services; esi { 00:22:22:22:22:22:22:22:22; all-active; } aggregated-ether-options { lACP { active; system-id 00:00:00:00:22:01; } } unit 1 { encapsulation vlan-bridge; vlan-id 10; } r2-RE0-re0# show routing-instances evpn-vlan-based { instance-type evpn; protocols { evpn { no-control-word; encapsulation mpls; } } vlan-id 20; route-distinguisher 22.22.22.22:1; vrf-target target:11:1; interface ae0.1; } </pre>
--	--

EVPN-MPLS E-LAN VLAN-BASED (Interop Junos-ACX vs Evolved) – Normalization

All the four PEs in [Figure 2](#) are BGP-EVPN PE, where two configured PEs on one side are Junos OS EVO based platforms with MAC-VRF instance and on the other side are Junos OS based ACX platforms configured with setup EVPN instance type. This configuration is partially RFC compliant.

The PEs discover each other in the network, establish the EVPN sessions, and exchange the EVPN routes. With the VLAN-based service, one VLAN is associated with an EVI with a one-to-one mapping of the VLAN-ID and the EVI. So, one bridge table for one VLAN.

[Table 1](#) highlights the different VLAN based models which are mentioned in the RFC 7432 where you would like take actions at the Ingress router based on your requirements.

E-LAN (VLAN-Based Interop Junos-ACX vs Evolved → Normalization) Configuration – A/A Mode

In the Active/Active mode, you must bundle the UNI links from CE node to the multihomed PE for load balancing the traffic.

In this case, you are provisioning VLAN-based service with normalization. The CEs connected on either side of the network are part of different VLANs and ensure that the traffic is forwarded. It is not RFC compliant in this case because the VLAN information is normalized in the ingress PE. The following configuration output is used for the Interop requirements.

<pre> PE1 r1-RE0-re0# show interfaces ae0 flexible-vlan-tagging; encapsulation flexible-ethernet-services; esi { 00:11:11:11:11:11:11:11; all-active; } aggregated-ether-options { lacp { active; system-id 00:00:00:00:11:01; } } unit 1 { encapsulation vlan-bridge; vlan-id 20; } r1-RE0-re0# show routing-instances evpn-vlan-based { instance-type mac-vrf; protocols { evpn { normalization; no-control-word; encapsulation mpls; } } service-type vlan-based; route-distinguisher 11.11.11.11:1; vrf-target target:11:1; vlans { bd1 { vlan-id 20; interface ae0.1; } } } </pre>	<pre> PE3 r2-RE0-re0# show interfaces ae0 flexible-vlan-tagging; encapsulation flexible-ethernet-services; esi { 00:22:22:22:22:22:22:22; all-active; } aggregated-ether-options { lacp { active; system-id 00:00:00:00:22:01; } } unit 1 { encapsulation vlan-bridge; vlan-id 10; } r2-RE0-re0# show routing-instances evpn-vlan-based { instance-type evpn; protocols { evpn; } vlan-id 20; route-distinguisher 22.22.22.22:1; vrf-target target:11:1; interface ae0.1; } </pre>
--	---

E-LAN (VLAN-Based Interop Junos-ACX vs Evolved -> Normalization) Operational – A/A Mode

```

I
MAC IP flags (S - Static, D - Dynamic, L - Local , R - Remote, Lp - Local Proxy,
Rp - Remote Proxy, K - Kernel, RT - Dest Route, (N)AD - (Not) Advt to remote,
RE - Re-ARP/ND, RO - Router, OV - Override, Ur - Unresolved,
RTS - Dest Route Skipped, RGw - Remote Gateway, FU - Fast Update)
Routing instance : evpn-vlan
Bridging domain : bd1
  IP           MAC           Flags           Logical           Active
  address     address     flags           Interface     source
  6.6.6.2     00:00:6e:9e:4e:65  DR,K           ae10.10      00:11:11:11:11:11:11
  6.6.6.1     00:00:6e:9e:4e:69  DL,K,AD
acx-481-06>
acx-481-06> show mac-vrf forwarding mac-table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static, C - Control MAC
SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned
Routing instance : evpn-vlan
  Vlan           MAC           MAC           Age           Logical           NH           RTR
  name         address     flags           interface     Index       ID
  bd1          00:00:6e:9e:4e:65  DC           -           ae10.10      15135      15135
  bd1          00:00:6e:9e:4e:69  D            -           ae10.10      0           0

acx-481-06> show route advertising-protocol bgp 3.3.3.3

bgp.evpn.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
  Prefix      Nexthop      MED      Lclpref      AS path
  1:1.1.1.1:0::22222222222222222222::FFFF:FFFF/192 AD/ESI
*           Self          100      I
  1:1.1.1.1:51::22222222222222222222::10/192 AD/EVI
*           Self          100      I
  2:1.1.1.1:51::10::00:00:6e:9e:4e:69/304 MAC/IP
*           Self          100      I
  2:1.1.1.1:51::10::00:00:6e:9e:4e:69::6.6.6.1/304 MAC/IP
*           Self          100      I
  3:1.1.1.1:51::10::1.1.1.1/248 IM
*           Self          100      I
  4:1.1.1.1:0::2222222222222222:1.1.1.1/296 ES
*           Self          100      I
evpn-vlan.evpn.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix      Nexthop      MED      Lclpref      AS path
  1:1.1.1.1:51::2222222222222222::10/192 AD/EVI
*           Self          100      I
  2:1.1.1.1:51::10::00:00:6e:9e:4e:69/304 MAC/IP
*           Self          100      I
  2:1.1.1.1:51::10::00:00:6e:9e:4e:69::6.6.6.1/304 MAC/IP
*           Self          100      I
  3:1.1.1.1:51::10::1.1.1.1/248 IM
*           Self          100      I
_default_evpn_evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  Prefix      Nexthop      MED      Lclpref      AS path
  1:1.1.1.1:0::2222222222222222::FFFF:FFFF/192 AD/ESI
*           Self          100      I
  4:1.1.1.1:0::2222222222222222:1.1.1.1/296 ES
*           Self          100      I

  Remote-PE      MAC-label      Alias-label      Mode
  5.5.5.5        161           161           all-active
  ESI: 00:22:22:22:22:22:22:22:22:22:22:22
  Status: Resolved by IFL ae10.10
  Local interface: ae10.10, Status: Up/Forwarding
  DF Election Algorithm: MOD based
  Designated forwarder: 1.1.1.1
  Last designated forwarder update: Jan 17 23:38:10
  Advertised split horizon label: 232
  SMET Forwarding: Disabled

acx-481-06> show mac-vrf forwarding mac-ip-table

```



```

Routing instance : evpn-junos
Bridging domain : __evpn-junos__
  IP                                MAC                                Flags          Logical          Active
  address                         address                         Interface      source
  6.6.6.2                          00:00:6e:9e:4e:65    DL,K,AD        ae0.10
  6.6.6.1                          00:00:6e:9e:4e:69    DR,K
                                         00:22:22:22:22:22:22:22:22

acx18> show evpn mac-table

MAC flags (S - static MAC, D - dynamic MAC, L - locally learned, P - Persistent static, C - Control MAC
          SE - statistics enabled, NM - non configured MAC, R - remote PE MAC, O - ovsdb MAC)

Ethernet switching table : 2 entries, 2 learned
Routing instance : evpn-junos
  Vlan          MAC          MAC          Age          Logical          NH          RTR
  name         address      flags        -           interface      Index      ID
  __evpn-junos__ 00:00:6e:9e:4e:65  D          -           ae0.10        0          0
  __evpn-junos__ 00:00:6e:9e:4e:69  DC         -           ae0.10      2097155  2097155

acx18> show route advertising-protocol bgp 3.3.3.3

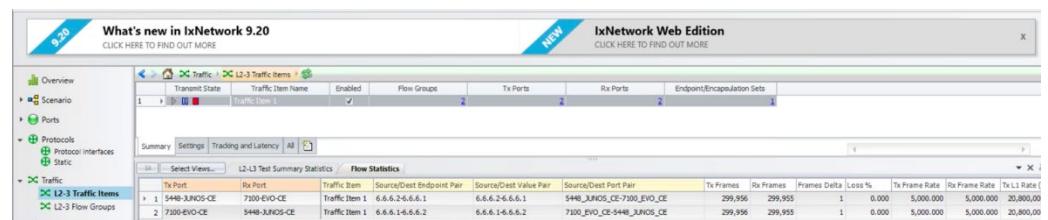
bgp.evpn.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
  Prefix      Nexthop      MED      Lclpref      AS path
  1:5.5.5.0::11111111111111111111111111111111:FFFF:FFFF/192 AD/ESI
*          Self          100      I
  1:5.5.5.5:1801::11111111111111111111111111111111::0/192 AD/EVI
*          Self          100      I
  2:5.5.5.5:1801::10::00:00:6e:9e:4e:65/304 MAC/IP
*          Self          100      I
  2:5.5.5.5:1801::10::00:00:6e:9e:4e:65::6.6.6.2/304 MAC/IP
*          Self          100      I
  3:5.5.5.5:1801::10::5.5.5.5/248 IM
*          Self          100      I
  4:5.5.5.5:0::11111111111111111111111111111111:5.5.5.5/296 ES
*          Self          100      I

evpn-junos.evpn.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
  Prefix      Nexthop      MED      Lclpref      AS path
  1:5.5.5.5:1801::11111111111111111111111111111111::0/192 AD/EVI
*          Self          100      I
  2:5.5.5.5:1801::10::00:00:6e:9e:4e:65/304 MAC/IP
*          Self          100      I
  2:5.5.5.5:1801::10::00:00:6e:9e:4e:65::6.6.6.2/304 MAC/IP
*          Self          100      I
  3:5.5.5.5:1801::10::5.5.5.5/248 IM
*          Self          100      I

_default_evpn_.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
  Prefix      Nexthop      MED      Lclpref      AS path
  1:5.5.5.5:0::11111111111111111111111111111111:FFFF:FFFF/192 AD/ESI
*          Self          100      I
  4:5.5.5.5:0::11111111111111111111111111111111:5.5.5.5/296 ES
*          Self          100      I

```

Figure 5: Traffic Snapshot

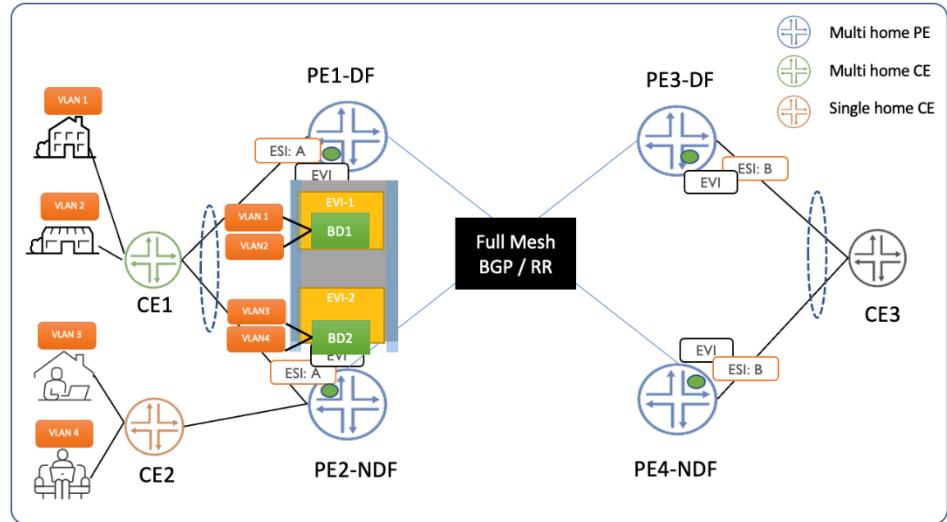


EVPN-MPLS E-LAN VLAN-Bundle

All three PEs in Figure 2 are BGP-EVPN PE. The PEs are set up using the EVPN MAC-VRF instance and the required configurations on all the PEs. These PEs discover each other in the network, establish the EVPN sessions, and exchange EVPN routes.

The current configuration is with the Active-Active mode for the multihomed PEs.

Figure 6: EVPN E-LAN VLAN-Bundle



E-LAN (VLAN-Bundle) Configuration - Active/Active Mode

In the Active/Active mode, you must bundle the UNI links from CE node to the multihomed PE for load balancing the traffic. You have two VLAN-IDs part that are of the same EVI. In the multihomed scenario, the ESI must be defined individually for each IFLs.

For a single/active scenario, the common ESI can be defined for both vlan-id in an EVI.

```

PE1
r1-RE0-re0# show routing-instances evpn-vlan-
bundle
instance-type mac-vrf;
protocols {
    evpn {
        encapsulation mpls;
    }
}
service-type vlan-bundle;
route-distinguisher 11.11.11.11:1;
vrf-target target:11:1;
vlans {
    bd1 {
        interface ae0.1;
        interface ae0.2;
    }
}
r1-RE0-re0# show interfaces ae0
flexible-vlan-tagging;
encapsulation flexible-ethernet-services;
aggregated-ether-options {
    lacp {
        active;
        system-id 00:00:00:00:11:01;
    }
}
unit 1 {
    encapsulation vlan-bridge;
    vlan-id 1;
    esi {
        00:11:11:11:11:11:11:11:11;
        all-active;
    }
}
unit 2 {
    encapsulation vlan-bridge;
    vlan-id 2;
    esi {
        00:11:11:11:11:11:11:11:12;
        all-active;
    }
}

PE3
r2-RE0-re0# show routing-instances evpn-vlan-
bundle
instance-type mac-vrf;
protocols {
    evpn {
        encapsulation mpls;
    }
}
service-type vlan-bundle;
route-distinguisher 22.22.22.22:1;
vrf-target target:11:1;
vlans {
    bd1 {
        interface ae0.1;
        interface ae0.2;
    }
}
r2-RE0-re0# show interfaces ae0
flexible-vlan-tagging;
encapsulation flexible-ethernet-services;
aggregated-ether-options {
    lacp {
        active;
        system-id 00:00:00:00:22:01;
    }
}
unit 1 {
    encapsulation vlan-bridge;
    vlan-id 1;
    esi {
        00:22:22:22:22:22:22:22:21;
        all-active;
    }
}
unit 2 {
    encapsulation vlan-bridge;
    vlan-id 2;
    esi {
        00:22:22:22:22:22:22:22:22;
        all-active;
    }
}

```

EVPN-VXLAN E-LAN

EVPN-VXLAN E-LAN VLAN-BASED

All four PEs in [Figure 4](#) are BGP-EVPN PE. The PEs are set up using the EVPN MAC-VRF instance and the required configurations on all the PEs. The PEs discover each other in the network, establish the EVPN sessions, and exchange the EVPN routes.

E-LAN (VLAN-Based) Configuration - Active/Active Mode

In the Active/Active mode, you must bundle the UNI links from CE node to the multihomed PE for load balancing the traffic.

```

PE1
r1-RE0-re0# show interfaces ae0
flexible-vlan-tagging;
encapsulation flexible-ethernet-services;
esi {
    00:11:11:11:11:11:11:11;
    all-active;
}
aggregated-ether-options {
    lACP {
        active;
        system-id 00:00:00:00:11:01;
    }
}
unit 1 {
    encapsulation vlan-bridge;
    vlan-id 1;
}

{master}[edit]
r1-RE0-re0# show routing-instances
evpn-vlan-based {
    instance-type mac-vrf;
    protocols {
        evpn;
        encapsulation vxlan;
    }
    service-type vlan-based;
    vtep-source-interface lo0.0
    route-distinguisher 11.11.11.11:1;
    vrf-target target:11:1;
    vlans {
        bd1 {
            vlan-id 1;
            interface ae0.1;
            vxlan vni 1
        }
    }
}

PE3
r2-RE0-re0# show interfaces ae0
flexible-vlan-tagging;
encapsulation flexible-ethernet-services;
esi {
    00:22:22:22:22:22:22:22:21;
    all-active;
}
aggregated-ether-options {
    lACP {
        active;
        system-id 00:00:00:00:22:01;
    }
}
unit 1 {
    encapsulation vlan-bridge;
    vlan-id 1;
}

{master}[edit]
r2-RE0-re0# show routing-instances
evpn-vlan-based {
    instance-type mac-vrf;
    protocols {
        evpn;
        encapsulation vxlan;
    }
    service-type vlan-based;
    vtep-source-interface lo0.0
    route-distinguisher 33.33.33.33:1;
    vrf-target target:11:1;
    vlans {
        bd1 {
            vlan-id 1;
            interface ae0.1;
            vxlan vni 1
        }
    }
}

```

EVPN-VXLAN E-LAN VLAN-Bundle

All four PEs in [Figure 7](#) topology are BGP-EVPN PE. The PEs are set up using the EVPN MAC-VRF instance and the required configurations on all the PEs. The PEs discover each other in the network, establish the EVPN sessions, and exchange EVPN routes.

The current configuration is with the Active-Active mode for the multihomed PEs.

E-LAN (VLAN-Bundle) Configuration - Active/Active Mode

In the Active/Active mode, you must bundle the UNI links from CE node to the multihomed PE for load balancing the traffic. You have two VLAN-IDs part that are of the same EVI. In the multihomed scenario, the ESI must be defined individually for each IFBs.

For a single/active scenario, the common ESI can be defined for both vlan-id in an EVI.

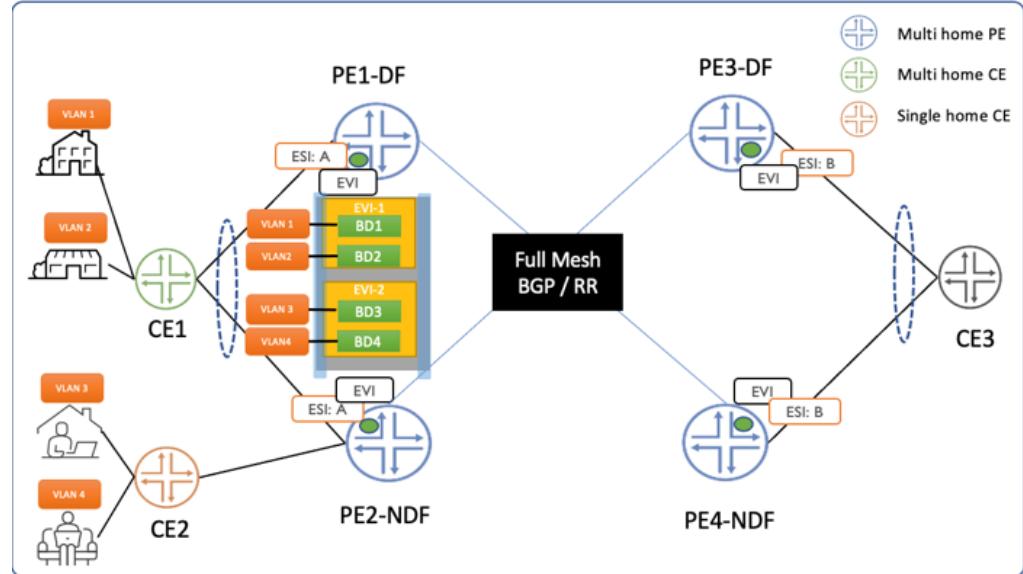
<pre> PE1 acx-481-09# show routing-instances evpn-vlan- bundle instance-type mac-vrf; protocols { evpn { encapsulation vxlan; } } vtep-source-interface lo0.0; service-type vlan-bundle; route-distinguisher 11.11.11.11:1; vrf-target target:11:1; vlans { bdl { interface ae0.1; interface ae0.2; vxlan { vni 1; } } } [edit] acx-481-09# show interfaces ae3 flexible-vlan-tagging; encapsulation flexible-ethernet-services; unit 1 { encapsulation vlan-bridge; vlan-id 1; esi { 00:11:11:11:11:11:11:11:11:11; all-active; } } unit 2 { encapsulation vlan-bridge; vlan-id 2; esi { 00:11:11:11:11:11:11:11:11:12; all-active; } } </pre>	<pre> PE3 acx-481-09# show routing-instances evpn-vlan- bundle instance-type mac-vrf; protocols { evpn { encapsulation vxlan; } } vtep-source-interface lo0.0; service-type vlan-bundle; route-distinguisher 22.22.22.22:1; vrf-target target:11:1; vlans { bdl { interface ae0.1; interface ae0.2; vxlan { vni 1; } } } [edit] acx-481-09# show interfaces ae0 flexible-vlan-tagging; encapsulation flexible-ethernet-services; unit 1 { encapsulation vlan-bridge; vlan-id 1; esi { 00:22:22:22:22:22:22:22:22:21; all-active; } } unit 2 { encapsulation vlan-bridge; vlan-id 2; esi { 00:11:11:11:11:11:11:11:11:12; all-active; } } </pre>
---	---

EVPN-VXLAN E-LAN VLAN-AWARE

All four PEs in the [Figure 8](#) topology are BGP-EVPN PE. The PEs are set up using the EVPN MAC-VRF instance and the required configurations on all the PEs. The PEs discover each other in the network, establish the EVPN sessions, and exchange EVPN routes.

The current configuration is with the Active-Active mode for the multihomed PEs.

Figure 7: EVPN E-LAN VLAN-Bundle



E-LAN (VLAN-Aware Bundle) Configuration - Active/Active Mode

In the Active/Active mode, you must bundle the UNI links from CE node to the multihomed PE for load balancing the traffic. You have two VLAN-IDs part that are of the same EVI. In the multihomed scenario, the ESI must be defined individually for each IFLs.

For a single/active scenario, the common ESI can be defined for both vlan-id in an EVI.

```

PE1
r1-RE0-re0# show routing-instances evpn-vlan-aware
instance-type mac-vrf;
protocols {
    evpn;
    encapsulation vxlan;
}
service-type vlan-aware;
route-distinguisher 11.11.11.11:1;
vrf-target target:11:1;
vtep-source-interface lo0.0
vlans {
    bd1 {
        interface ae0.1;
        vxlan vni 1
    }
    bd2 {
        interface ae0.2;
        vxlan vni 2
    }
}
r1-RE0-re0# show interfaces ae0
flexible-vlan-tagging;
encapsulation flexible-ethernet-services;
aggregated-ether-options {
    lacp {
        active;
        system-id 00:00:00:00:11:01;
    }
}
unit 1 {
    encapsulation vlan-bridge;
    vlan-id 1;
    esi {
        00:11:11:11:11:11:11:11:11;
        all-active;
    }
}

PE3
r2-RE0-re0# show routing-instances evpn-vlan-aware
instance-type mac-vrf;
protocols {
    evpn;
    encapsulation vxlan;
}
service-type vlan-aware;
route-distinguisher 22.22.22.22:1;
vrf-target target:11:1;
vtep-source-interface lo0.0
vlans {
    bd1 {
        interface ae0.1;
        vxlan vni 1
    }
    bd2 {
        interface ae0.2;
        vxlan vni 2
    }
}
r2-RE0-re0# show interfaces ae0
flexible-vlan-tagging;
encapsulation flexible-ethernet-services;
aggregated-ether-options {
    lacp {
        active;
        system-id 00:00:00:00:22:01;
    }
}
unit 1 {
    encapsulation vlan-bridge;
    vlan-id 1;
    esi {
        00:22:22:22:22:22:22:22:21;
        all-active;
    }
}

```

```
unit 2 {  
    encapsulation vlan-bridge;  
    vlan-id 2;  
    esi {  
        00:11:11:11:11:11:11:11:11:12;  
        all-active;  
    }  
}  
  
unit 2 {  
    encapsulation vlan-bridge;  
    vlan-id 3;  
    esi {  
        00:22:22:22:22:22:22:22:22:22;  
        all-active;  
    }  
}
```

CHAPTER 3 Common Configuration

For any EVPN services (ELINE/ELAN/ETREE), the below mentioned common configurations are the underlay configurations for the four Multihomed PEs that remain same. Here, we have IGP as ISIS and Segment Routing (SR) as transport. Each PE has a unique System ID, Loopback ID, and Node segment ipv4-index. These configurations are operator/Network engineer's choice and can be different than the example given below.

Underlay Configuration – MPLS/OSPF

```

PE1
r1-re0# show interfaces
et-0/0/1 {
    unit 0 {
        description to-r4;
        family inet {
            address 10.1.1.1/30;
        }
        family mpls;
    }
}
et-0/0/3 {
    unit 0 {
        description to-r2;
        family inet {
            address 10.1.2.1/30;
        }
        family mpls;
    }
}
et-0/0/5 {
    unit 0 {
        description to-r3;
        family inet {
            address 10.1.3.1/30;
        }
        family mpls;
    }
}
et-0/0/7 {
    description "To R5";
    ether-options {
        802.3ad ae0;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.1/32;
        }
    }
}
r1-re0# show protocols
bgp {
    group ibgp {
        type internal;
        local-address 192.168.0.1;
        family evpn {
            signaling;
        }
        neighbor 192.168.0.2;
        neighbor 192.168.0.3;
        neighbor 192.168.0.4;
    }
}
ldp {
    interface all;
}

PE2
r2-re0# show interfaces
et-0/0/1 {
    unit 0 {
        description to-r3;
        family inet {
            address 10.1.23.1/30;
        }
        family mpls;
    }
}
et-0/0/3 {
    unit 0 {
        description to-r1;
        family inet {
            address 10.1.2.2/30;
        }
        family mpls;
    }
}
et-0/0/5 {
    unit 0 {
        description to-r4;
        family inet {
            address 10.1.24.1/30;
        }
        family mpls;
    }
}
et-0/0/9 {
    description "To R5";
    ether-options {
        802.3ad ae0;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.2/32;
        }
    }
}
r2-re0# show protocols
bgp {
    group ibgp {
        type internal;
        local-address 192.168.0.2;
        family evpn {
            signaling;
        }
        neighbor 192.168.0.1;
        neighbor 192.168.0.3;
        neighbor 192.168.0.4;
    }
}
ldp {
    interface all;
}

```

```

        interface fxp0.0 {
            disable;
        }
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface et-0/0/1.0;
            interface et-0/0/3.0;
            interface et-0/0/5.0;
            interface lo0.0 {
                passive;
            }
        }
    }
    lldp {
        interface all;
    }

r1-re0# show routing-options
router-id 192.168.0.1;
autonomous-system 65000;

i      interface fxp0.0 {
            disable;
        }
    mpls {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    ospf {
        area 0.0.0.0 {
            interface et-0/0/1.0;
            interface et-0/0/3.0;
            interface et-0/0/5.0;
            interface lo0.0 {
                passive;
            }
        }
    }
    lldp {
        interface all;
    }

r2-re0# show routing-options
router-id 192.168.0.2;
autonomous-system 65000;

```

<pre> PE3 r3-re0# show interfaces et-0/0/1 { unit 0 { description to-r2; family inet { address 10.1.23.2/30; } family mpls; family iso; } } et-0/0/3 { unit 0 { description to-r1; family inet { address 10.1.3.2/30; } family mpls; family iso; } } et-0/0/5 { unit 0 { description to-r4; family inet { address 10.1.34.1/30; } family mpls; family iso; } } et-0/0/7 { description "To R7"; ether-options { 802.3ad ae0; } } lo0 { unit 0 { family inet { address 192.168.0.3/32; } family iso { address 49.0001.3333.3333.00 } } } </pre>	<pre> PE4 r4-re0# show interfaces et-0/0/1 { unit 0 { description to-r1; family inet { address 10.1.1.2/30; } family mpls; family iso; } } et-0/0/3 { unit 0 { description to-r2; family inet { address 10.1.24.2/30; } family mpls; family iso; } } et-0/0/5 { unit 0 { description to-r3; family inet { address 10.1.34.2/30; } family mpls; family iso; } } et-0/0/7 { description "To R7"; ether-options { 802.3ad ae0; } } lo0 { unit 0 { family inet { address 192.168.0.4/32; } family iso { address 49.0001.4444.4444.00 } } } </pre>
---	---

```
r3-re0# show protocols
bgp {
  group ibgp {
    type internal;
    local-address 192.168.0.3;
    family evpn {
      signaling;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
    neighbor 192.168.0.4;
  }
  isis {
    interface all {
      point-to-point;
    }
    interface fxp0.0 {
      disable;
    }
    source-packet-routing {
      node-segment ipv4-index 1003;
    }
    level 2 wide-metrics-only;
    level 1 disable;
  }
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  lldp {
    interface all;
  }
}
r3-re0# show routing-options
router-id 192.168.0.3;
autonomous-system 65000;

r4-re0# show protocols
bgp {
  group ibgp {
    type internal;
    local-address 192.168.0.4;
    family evpn {
      signaling;
    }
    neighbor 192.168.0.1;
    neighbor 192.168.0.2;
    neighbor 192.168.0.3;
  }
  isis {
    interface all {
      point-to-point;
    }
    interface fxp0.0 {
      disable;
    }
    source-packet-routing {
      node-segment ipv4-index 1004;
    }
    level 2 wide-metrics-only;
    level 1 disable;
  }
  mpls {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  lldp {
    interface all;
  }
}
r4-re0# show routing-options
router-id 192.168.0.4;
autonomous-system 65000;
```

Underlay Configuration – SR, ISIS

<pre> PE1 r1-re0# show interfaces et-0/0/1 { unit 0 { description to-r4; family inet { address 10.1.1.1/30; } family mpls; } } et-0/0/3 { unit 0 { description to-r2; family inet { address 10.1.2.1/30; } family mpls; } } et-0/0/5 { unit 0 { description to-r3; family inet { address 10.1.3.1/30; } family mpls; } } et-0/0/7 { description "To R5"; ether-options { 802.3ad ae0; } } lo0 { unit 0 { family inet { address 192.168.0.1/32; } } } r1-re0# show protocols bgp { group ibgp { type internal; local-address 192.168.0.1; family evpn { signaling; } neighbor 192.168.0.2; neighbor 192.168.0.3; neighbor 192.168.0.4; } } isis { interface all { point-to-point; } interface fxp0.0 { disable; } source-packet-routing { node-segment ipv4-index 1001; } level 2 wide-metrics-only; level 1 disable; } mpls { interface all; interface fxp0.0 { disable; } } lldp { interface all; } r1-re0# show routing-options router-id 192.168.0.1; autonomous-system 64512; </pre>	<pre> PE2 r2-re0# show interfaces et-0/0/1 { unit 0 { description to-r3; family inet { address 10.1.23.1/30; } family mpls; } } et-0/0/3 { unit 0 { description to-r1; family inet { address 10.1.2.2/30; } family mpls; } } et-0/0/5 { unit 0 { description to-r4; family inet { address 10.1.24.1/30; } family mpls; } } et-0/0/9 { description "To R5"; ether-options { 802.3ad ae0; } } lo0 { unit 0 { family inet { address 192.168.0.2/32; } } } r2-re0# show protocols bgp { group ibgp { type internal; local-address 192.168.0.2; family evpn { signaling; } neighbor 192.168.0.1; neighbor 192.168.0.3; neighbor 192.168.0.4; } } isis { interface all { point-to-point; } interface fxp0.0 { disable; } source-packet-routing { node-segment ipv4-index 1001; } level 2 wide-metrics-only; level 1 disable; } mpls { interface all; interface fxp0.0 { disable; } } lldp { interface all; } r2-re0# show routing-options router-id 192.168.0.2; autonomous-system 64512; </pre>
--	--

```

PE3
r3-re0# show interfaces
et-0/0/1 {
    unit 0 {
        description to-r2;
        family inet {
            address 10.1.23.2/30;
        }
        family mpls;
        family iso;
    }
}
et-0/0/3 {
    unit 0 {
        description to-r1;
        family inet {
            address 10.1.3.2/30;
        }
        family mpls;
        family iso;
    }
}
et-0/0/5 {
    unit 0 {
        description to-r4;
        family inet {
            address 10.1.34.1/30;
        }
        family mpls;
        family iso;
    }
}
et-0/0/7 {
    description "To R7";
    ether-options {
        802.3ad ae0;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.3/32;
        }
        family iso {
            address 49.0001.3333.3333.00
        }
    }
}
r3-re0# show protocols
bgp {
    group ibgp {
        type internal;
        local-address 192.168.0.3;
        family evpn {
            signaling;
        }
        neighbor 192.168.0.1;
        neighbor 192.168.0.2;
        neighbor 192.168.0.4;
    }
}
isis {
    interface all {
        point-to-point;
    }
    interface fxp0.0 {
        disable;
    }
    source-packet-routing {
        node-segment ipv4-index 1001;
    }
    level 2 wide-metrics-only;
    level 1 disable;
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
lldp {
    interface all;
}
r3-re0# show routing-options
router-id 192.168.0.3;
autonomous-system 65000;

PE4
r4-re0# show interfaces
et-0/0/1 {
    unit 0 {
        description to-r1;
        family inet {
            address 10.1.1.2/30;
        }
        family mpls;
        family iso;
    }
}
et-0/0/3 {
    unit 0 {
        description to-r2;
        family inet {
            address 10.1.24.2/30;
        }
        family mpls;
        family iso;
    }
}
et-0/0/5 {
    unit 0 {
        description to-r3;
        family inet {
            address 10.1.34.2/30;
        }
        family mpls;
        family iso;
    }
}
et-0/0/7 {
    description "To R7";
    ether-options {
        802.3ad ae0;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.4/32;
        }
        family iso {
            address 49.0001.4444.4444.00
        }
    }
}
r4-re0# show protocols
bgp {
    group ibgp {
        type internal;
        local-address 192.168.0.4;
        family evpn {
            signaling;
        }
        neighbor 192.168.0.1;
        neighbor 192.168.0.2;
        neighbor 192.168.0.3;
    }
}
isis {
    interface all {
        point-to-point;
    }
    interface fxp0.0 {
        disable;
    }
    source-packet-routing {
        node-segment ipv4-index 1001;
    }
    level 2 wide-metrics-only;
    level 1 disable;
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
lldp {
    interface all;
}
r4-re0# show routing-options
router-id 192.168.0.4;
autonomous-system 65000;

```

The below CE configuration covers when the MX Series or ACX Series are used as only CE devices in different models that includes single vlan-tag and dual vlan-tag.

MX Series/ACX Series – CE Configuration in EPL, A/A Mode

```

CE1
r5_re0# show interfaces
ge-0/0/0 {
    description "To R2";
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/0/1 {
    description "To R1";
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/0/2 {
    description "To IXIA";
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation ethernet-bridge;
    }
}
ae0 {
    description "Connected to PE1";
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:00:00:00:11:02;
        }
        unit 1 {
            encapsulation ethernet-bridge;
        }
    }
}

CE2
r7_re0# show interfaces
ge-0/0/0 {
    description "To IXIA";
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation ethernet-bridge;
    }
}
ge-0/0/1 {
    description "To R3";
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/0/2 {
    description "To R4";
    gigether-options {
        802.3ad ae0;
    }
}
ae0 {
    description "Connected to PE3-4";
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
        }
        unit 0 {
            encapsulation ethernet-bridge;
        }
    }
}

```

MX Series – EVPL Configuration Single Tag, A/A mode

<pre>CE1 R5-RE0> show configuration interfaces ge-0/0/0 { description "To R2"; gigether-options { 802.3ad ae0; } } ge-0/0/1 { description "To R1"; gigether-options { 802.3ad ae0; } } ge-0/0/2 { description "To IXIA"; flexible-vlan-tagging; encapsulation flexible-ethernet-services; unit 1 { encapsulation vlan-bridge; vlan-id 1; } } ae0 { description "Connected to PE1"; flexible-vlan-tagging; encapsulation flexible-ethernet-services; aggregated-ether-options { lACP { active; system-id 00:00:00:00:11:02; } } unit 1 { encapsulation vlan-bridge; vlan-id 1; } } R5-RE0> show configuration bridge-domains bd1 { vlan-id 1; interface ae0.1; interface ge-0/0/2.1; }</pre>	<pre>CE2 R7-RE0> show configuration interfaces ge-0/0/0 { description "To IXIA"; flexible-vlan-tagging; encapsulation flexible-ethernet-services; unit 1 { encapsulation vlan-bridge; vlan-id 1; } } ge-0/0/1 { description "To R3"; gigether-options { 802.3ad ae0; } } ge-0/0/2 { description "To R4"; gigether-options { 802.3ad ae0; } } ae0 { description "Connected to PE1"; flexible-vlan-tagging; encapsulation flexible-ethernet-services; aggregated-ether-options { lACP { active; system-id 00:00:00:00:22:02; } } unit 1 { encapsulation vlan-bridge; vlan-id 1; } } R7-RE0> show configuration bridge-domains bd1 { vlan-id 1; interface ae0.1; interface ge-0/0/0.1; }</pre>
---	---

ACX Series – EVPL Configuration Single Tag, A/A Mode

The configuration remains same as MX Series - EVPL Configuration Single Tag, A/A mode except the following configuration:

<pre>CE1 R5-RE0> show configuration vlans bd1 { vlan-id 1; interface ae0.1; interface ge-0/0/2.1; }</pre>	<pre>CE2 R7-RE0> show configuration vlans bd1 { vlan-id 1; interface ae0.1; interface ge-0/0/0.1; }</pre>
--	--

MX Series – CE Configuration in EVPL & Dual-Tag, A/A Mode

```

CE1
R5-RE0> show configuration interfaces
ge-0/0/0 {
    description "To R2";
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/0/1 {
    description "To R1";
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/0/2 {
    description "To IXIA";
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
        encapsulation vlan-bridge;
        vlan-tags outer 1 inner 2;
    }
}
ae0 {
    description "Connected to PE1";
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:00:00:00:11:02;
        }
        unit 1 {
            encapsulation vlan-bridge;
            vlan-tags outer 1 inner 2;
        }
    }
}

R5-RE0> show configuration bridge-domains
bd1 {
    vlan-tags outer 1 inner 2;
    interface ae0.1;
    interface ge-0/0/2.1;
}

CE2
R7-RE0> show configuration interfaces
ge-0/0/0 {
    description "To IXXIA";
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
        encapsulation vlan-bridge;
        vlan-tags outer 1 inner 2;
    }
}
ge-0/0/1 {
    description "To R3";
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/0/2 {
    description "To R4";
    gigether-options {
        802.3ad ae0;
    }
}
ae0 {
    description "Connected to PE1";
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    aggregated-ether-options {
        lacp {
            active;
            system-id 00:00:00:00:22:02;
        }
        unit 1 {
            encapsulation vlan-bridge;
            vlan-tags outer 1 inner 2;
        }
    }
}

R7-RE0> show configuration bridge-domains
bd1 {
    vlan-tags outer 1 inner 2;
    interface ae0.1;
    interface ge-0/0/0.1;
}

```

ACX Series – EVPL Configuration Dual Tag, A/A mode

The configuration remains same as MX Series - CE Configuration in EVPL & Dual-Tag with A/A Mode except the following configuration:

```

CE1
R5-RE0> show configuration vlans
bd1 {
    vlan-tags outer 1 inner 2;
    interface ae0.1;
    interface ge-0/0/2.0;
}

CE2
R7-RE0> show configuration vlans
bd1 {
    vlan-id outer 1 inner 2;
    interface ae0.1;
    interface ge-0/0/0.0;
}

```