



Ethernet Virtual Connections Configuration

An Ethernet Virtual Connection (EVC) is defined by the Metro-Ethernet Forum (MEF) as an association between two or more user network interfaces that identifies a point-to-point or multipoint-to-multipoint path within the service provider network. An EVC is a conceptual *service pipe* within the service provider network. A *bridge domain* is a local broadcast domain that is VLAN-ID-agnostic. An Ethernet flow point (EFP) service instance is a logical interface that connects a bridge domain to a physical port or to an EtherChannel group.

An EVC broadcast domain is determined by a bridge domain and the EFPs that are connected to it. You can connect multiple EFPs to the same bridge domain on the same physical interface, and each EFP can have its own matching criteria and rewrite operation. An incoming frame is matched against EFP matching criteria on the interface, learned on the matching EFP, and forwarded to one or more EFPs in the bridge domain. If there are no matching EFPs, the frame is dropped.

You can use EFPs to configure VLAN translation. For example, if there are two EFPs egressing the same interface, each EFP can have a different VLAN rewrite operation, which is more flexible than the traditional switchport VLAN translation model.

QoS policies on EFPs are supported with ingress rewrite type as push. In the ingress direction with one VLAN tag is pushed and in the egress direction one VLAN tag is popped.

This document describes how to configure EVC features.

For detailed information about the commands, see:

- The Cisco IOS XE 3S Carrier Ethernet Command Reference:
http://www.cisco.com/en/US/docs/ios/cether/command/reference/ce_book.html
- Master Command Index for Cisco IOS XE Release 3S:
http://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_book.html
- [Supported EVC Features, on page 1](#)
- [Restrictions for Ethernet Virtual Connections Configuration, on page 2](#)
- [Configuring EFPs, on page 3](#)
- [Configuring Other Features on EFPs, on page 9](#)
- [Configuring a Static MAC Address, on page 15](#)
- [Monitoring EVC, on page 17](#)

Supported EVC Features

- Service instance—you create, delete, and modify EFP service instances on Ethernet interfaces.

- Encapsulation—you can map traffic to EFPs based on:
 - 802.1Q VLANs (a single VLAN or a list or range of VLANs)
 - 802.1Q tunneling (QinQ) VLANs (a single outer VLAN and a list or range of inner VLANs)
 - Double-tagged frames mapped to EVC based on C-tags (wildcard S-Tags)
- Bridge domains—you can configure EFPs as members of a bridge domain (up to 64 EFPs per bridge domain for bridge domain with BDIs.).
- Rewrite (VLAN translation)
 - Pop symmetric
 - pop 1** removes the outermost tag
 - pop 2** removes the two outermost tags
 - pop symmetric** adds a tag (or 2 tags for **pop 2 symmetric**) on egress for a *push* operation
 - QinQ with rewrite
- EVC forwarding
- MAC address learning and aging
- EVCs on EtherChannels
- Split horizon
- Layer 2 protocol tunneling and QinQ
- Bridging between EFPs
- MSTP (MST on EVC bridge domain)
- EFP statistics (packets and bytes)
- QoS aware EVC/EFP per service instance
- Static MAC Addresses

These Layer 2 port-based features can run with EVC configured on the port:

- LACP
- CDP
- MSTP
- EVC egress filtering

Restrictions for Ethernet Virtual Connections Configuration

- Translate operations are not supported.
- You can create a maximum of 128 EFPs per bridge-domain.

- Only dot1q encapsulation is supported on trunk EFPs.
- Ingress mapping of Differentiated Services Code Point (DSCP) or Class of Service (CoS) to the C-CoS or S-CoS is supported.
- Egress classification and queuing is based on DSCP or CoS.
- 1024 EFPs per port are supported, which fall into the category of second or higher EFP configured under any BDI on that port.
- 64 EFPs per BD is supported for BDs with BDI.
- Egress filtering is not supported for EFP with double tagged encapsulation.
- Frame forwarding of double tagged frames with both outermost and innermost tag as 0x88a8 is not supported.
- Frame forwarding of double tagged frames with outermost tag as 0x8100 and innermost tag as 0x88a8 is not supported.
- Static MAC entry does not get cleared on deleting the encapsulation dot1q VLAN.
- Classification based on multiple encapsulation types in single EFP is not supported.
- MAC ageing happens randomly, it takes around 700 seconds to age out.
- Show MAC address table output takes a maximum of 60 seconds to syncup between the hardware entries.
- Custom Ethertype is not supported.
- Hair pin switching is not supported.
- Encapsulation criteria as Ethertype is not supported.

Configuring EFPs

Default EVC Configuration

No EFPs are configured. No service instances or bridge domains are configured.

Configuration Guidelines

The following guidelines apply when you configure EVCs on the router.

- To configure a service instance on an interface, these commands are prerequisites:

```
Router (config)# interface gigabitethernet0/0/1
Router (config-if)# service instance 22 Ethernet ether
Router (config-if-srv)# encapsulation dot1q 10
Router (config-if-srv)# bridge-domain 10
```

- You must configure encapsulation on a service instance before configuring bridge domain.
- ISL trunk encapsulation is not supported.

- The router does not support overlapping configurations on the same interface and same bridge domain. If you have configured a VLAN range encapsulation, or encapsulation default, or encapsulation any on service instance 1, you cannot configure any other encapsulations that also match previous encapsulations in the same interface and bridge domain.
- QinQ is not supported on Trunk EFP interfaces.
- Trunk EFPs should be configured with the **rewrite ingress tag pop 1 symmetric** command.
- In MST instance, when you add or remove VLAN from Trunk EFP, the BDI interface goes down which results in loss of packets.
- On an access interface configured with EFP untagged and TEFP, when a tagged packet with encapsulation equal to bridge-domain of untagged EFP passes through the access interface, the packet passes through TEFP and returns to the source through EFP untagged configuration and the packet is not dropped.

Creating Service Instances

Beginning in privileged EXEC mode, follow these steps to create an EFP service instance:

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to attach to the policy map, and enter interface configuration mode. Valid interfaces are physical ports.
Step 3	service instance <i>number</i> ethernet [<i>name</i>]	Configure an EFP (service instance) and enter service instance configuration) mode. <ul style="list-style-type: none"> • The number is the EFP identifier, an integer from 1 to 4000. • (Optional) ethernet name is the name of a previously configured EVC. You do not need to use an EVC name in a service instance.
Step 4	encapsulation { default dot1q priority-tagged untagged }	Configure encapsulation type for the service instance. <ul style="list-style-type: none"> • default—Configure to match all unmatched packets. • dot1q—Configure 802.1Q encapsulation. See for details about options for this keyword. • priority-tagged—Specify priority-tagged frames, VLAN-ID 0 and CoS value of 0 to 7.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • untagged—Map to untagged VLANs. Only one EFP per port can have untagged encapsulation.
Step 5	rewrite ingress tag pop {1 2} symmetric	(Optional) Specify that encapsulation modification to occur on packets at ingress. <ul style="list-style-type: none"> • pop 1—Pop (remove) the outermost tag. • pop 2—Pop (remove) the two outermost tags. • symmetric—Configure the packet to undergo the reverse of the ingress action at egress. If a tag is popped at ingress, it is pushed (added) at egress. This keyword is required for rewrite to function properly.
Step 6	bridge-domain <i>bridge-id</i> [split-horizon group <i>group-id</i>]	Configure the bridge domain ID. The range is from 1 to 4000. You can use the split-horizon keyword to configure the port as a member of a split horizon group. The <i>group-id</i> range is from 0 to 2.
Step 7	end	Return to privileged EXEC mode.
Step 8	show ethernet service instance show bridge-domain [<i>n</i> split-horizon]	Verify your entries.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file. Use the no forms of the commands to remove the service instance, encapsulation type, or bridge domain or to disable the rewrite operation.

Creating a Trunk EFP

Beginning in privileged EXEC mode, follow these steps to create an EFP service instance:



Note Use the **no** forms of the commands to remove the service instance, encapsulation type, or bridge domain or to disable the rewrite operation.



Note Trunk EFPs on port-channel interfaces is supported. Traffic may *not* flow to the TEFP when the port-channel or its member links are in down state.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the port to attach to the policy map, and enter interface configuration mode. Valid interfaces are physical ports.
Step 3	service instance [trunk] <i>number</i> ethernet	<p>Configure an EFP (service instance) and enter service instance configuration) mode.</p> <ul style="list-style-type: none"> • The number is the EFP identifier, an integer from 1 to 4000. • The trunk keyword identifies the trunk ID to which the service instance is assigned. <p>Note Trunk EFP (without port channel) supports encapsulation of up to 1000 VLANs.</p>
Step 4	encapsulation {default dot1q priority-tagged untagged}	<p>Note Only dot1q encapsulation is supported on trunk EFPs.</p> <p>Configure encapsulation type for the service instance.</p> <ul style="list-style-type: none"> • default —Configure to match all unmatched packets. • dot1q —Configure 802.1Q encapsulation. See Table 1 for details about options for this keyword. • priority-tagged —Specify priority-tagged frames, VLAN-ID 0 and CoS value of 0 to 7. • untagged —Map to untagged VLANs. Only one EFP per port can have untagged encapsulation.
Step 5	rewrite ingress tag pop {1 2} symmetric	<p>(Optional) Specify that encapsulation modification to occur on packets at ingress.</p> <ul style="list-style-type: none"> • pop 1 —Pop (remove) the outermost tag.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • pop 2—Pop (remove) the two outermost tags. <p>Caution The pop2 option is not currently supported on Trunk EFPs.</p> <ul style="list-style-type: none"> • symmetric—Configure the packet to undergo the reverse of the ingress action at egress. If a tag is popped at ingress, it is pushed (added) at egress. This keyword is required for rewrite to function properly.
Step 6	bridge-domain <i>bridge-id</i>	Configures the router to derive bridge domains from the encapsulation VLAN list.
Step 7	end	Return to privileged EXEC mode.
Step 8	Use one of the following commands <ul style="list-style-type: none"> • show ethernetservice instance • show bridge-domain [<i>n</i> split-horizon] 	Verify your entries.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuration Examples

Example for Configuring a Service Instance

Example for Encapsulation Using a VLAN Range

Configuration Example for Larger String VLAN in Encapsulation

Configuration Example

```

show running config

ethernet service multi-line
!
interface GigabitEthernet0/0/0
service instance 1 ethernet
encapsulation dot1q 10,13,19-21,24,29,32-36,41,46-48,55,61,63-66
encapsulation dot1q add 69-73,78,80,83-86
!
service instance 2 ethernet
encapsulation dot1q 1 second-dot1q 10,13,19-21,24,29,32-36,41
encapsulation dot1q add outer 2-5,7
encapsulation dot1q add inner 46-48,55,61,63-66,69-73,78,80,83-86
encapsulation dot1q add inner 91,95-99,101
!
interface GigabitEthernet0/0/0

```

Example for Two Service Instances Joining the Same Bridge Domain

```

ethernet dot1ad nni
service instance 3 ethernet
encapsulation dot1ad 10,13,19-21,24,29,32-36,41,46-48,55,61,63-66
encapsulation dot1ad add 69-73,78,80,83-86
!
service instance 4 ethernet
encapsulation dot1ad 1 dot1q 10,13,19-21,24,29,32-36,41,46-48,55
encapsulation dot1ad add inner 61,63-66,69-73,78,80,83-86
!
!
```

Example for Two Service Instances Joining the Same Bridge Domain

In this example, service instance 1 on interfaces Gigabit Ethernet 0/0/1 and 0/0/2 can bridge between each other.

Example for Bridge Domains and VLAN Encapsulation

Unlike VLANs, the bridge-domain number does not need to match the VLAN encapsulation number.

However, when encapsulations do not match in the same bridge domain, traffic cannot be forwarded. In this example, the service instances on Gigabit Ethernet 0/0/1 and 0/0/2 can not forward between each other, since the encapsulations don't match (filtering criteria). However, you can use the **rewrite** command to allow communication between these two.

Example for Rewrite

In this example, a packet that matches the encapsulation will have one tag removed (popped off). The **symmetric** keyword allows the reverse direction to have the inverse action: a packet that egresses out this service instance will have the encapsulation (VLAN 10) added (pushed on).

Example for Split Horizon

In this example, service instances 1 and 2 cannot forward and receive packets from each other. Service instance 3 can forward traffic to any service instance in bridge domain 3000 since no other service instance in bridge domain 3000 is in split-horizon group 2. Service instance 4 can forward traffic to any service instance in bridge domain 3000 since it has not joined any split-horizon groups.

Example for Egress Filtering

In EVC switching, egress filtering is performed before the frame is sent on the egress EFP. Egress filtering ensures that when a frame is sent, it conforms to the matching criteria of the service instance applied on the ingress direction. EFP does not require egress filtering if the number of pops is the same as the number of VLANs specified in the **encapsulation** command.

Egress Filtering is not supported on the RSP3 module.



Note

Specifying the **cos** keyword in the encapsulation command is relevant only in the ingress direction. For egress filtering, **cos** is ignored.

For example, consider the following configuration.

If a packet with VLAN tag 10 or 20 is received on Gigabit Ethernet 0/0/3, the ingress logical port would be service instance 3. For the frame to be forwarded on a service instance, the egress frame must match the encapsulation defined on that service instance after the rewrite is done. Service instance 1 checks for outermost VLAN 20; service instance 2 checks for VLAN 30. In this example, the frame with VLAN tags 10 and 20 can be sent to service instance 1 but not to service instance 2.

Configuring Other Features on EFPs

EFPs and EtherChannels

You can configure EFP service instances on EtherChannel port channels, but EtherChannels are not supported on ports configured with service instances. Load-balancing on port channels is based on the MAC address or IP address of the traffic flow on the EtherChannel interface.

This example configures a service instance on an EtherChannel port channel. Configuration on the ports in the port channel are independent from the service instance configuration.

```
Router (config)# interface port-channel 4
Router (config-if)# service instance 1 ethernet
Router (config-if-srv)# encapsulation untagged
Router (config-if-srv)# bridge-domain {any vlan}
Router (config-if-srv)# l2protocol peer {lACP | pagp}
```

Layer 2 Protocol Peering

For Layer 2 protocols (CDP, UDLD, LLDP, MSTP, LACP,) to peer with a neighbor on a port that has an EFP service instance configured, you need to enter the **l2 protocol peer** *protocol* service-instance configuration command on the service instance.

This example shows how to configure CDP to peer with a neighbor on a service instance:

Layer 2 Protocol Software Forwarding

Layer 2 protocol forwarding is based on the bridge domain ID and the destination MAC address.

Selecting the l2protocol forward option causes the router to flood interfaces in the same VLAN or bridge-domain with untagged or tagged BPDU packets. You can apply the l2protocol forward command to CDP, LACP, LLDP, PAGP, STP, UDLD, and VTP traffic. This is an example how to configure the l2protocol forward option:

Configuring IEEE 802.1Q Tunneling and Layer 2 Protocol Tunneling Using EFPs

Tunneling is a feature used by service providers whose networks carry traffic of multiple customers and who are required to maintain the VLAN and Layer 2 protocol configurations of each customer without impacting the traffic of other customers. The router uses EFPs to support QinQ and Layer 2 protocol tunneling.

802.1Q Tunneling (QinQ)

Service provider customers often have specific requirements for VLAN IDs and the number of VLANs to be supported. The VLAN ranges required by different customers in the same service-provider network might overlap, and traffic of customers through the infrastructure might be mixed. Assigning a unique range of VLAN IDs to each customer would restrict customer configurations and could easily exceed the VLAN limit (4096) of the 802.1Q specification.

Using the EVCs, service providers can encapsulate packets that enter the service-provider network with multiple customer VLAN IDs (C-VLANs) and a single 0x8100 Ethertype VLAN tag with a service provider VLAN (S-VLAN). Within the service provider network, packets are switched based on the S-VLAN. When the packets egress the service provider network onto the customer network, the S-VLAN tag is decapsulated and the original customer packet is restored.

Figure below shows the tag structures of the double-tagged packets.

In figure below, Customer A was assigned VLAN 30, and Customer B was assigned VLAN 40. Packets entering the edge switches with 802.1Q tags are double-tagged when they enter the service-provider network, with the outer tag containing VLAN ID 30 or 40, appropriately, and the inner tag containing the original VLAN number, for example, VLAN 100. Even if both Customers A and B have VLAN 100 in their networks, the traffic remains segregated within the service-provider network because the outer tag is different. Each customer controls its own VLAN numbering space, which is independent of the VLAN numbering space used by other customers and the VLAN numbering space used by the service-provider network. At the outbound port, the original VLAN numbers on the customer's network are recovered.

Method 1

In this example, for Customer A, interface is the customer-facing port, and is a trunk port facing the service provider network. For Customer B, is the customer-facing port, and is the trunk port facing the service provider network.

Customer A

For Customer A, service instance 1 on is configured with the VLAN encapsulations used by the customer: C-VLANs 1–100. These are forwarded on bridge-domain 4000. The service provider facing port is configured with a service instance on the same bridge-domain and with an **encapsulation dot1q** command matching the S-VLAN. The **rewrite ingress pop 1 symmetric** command also implies a push of the configured encapsulation on egress packets. Therefore, the original packets with VLAN tags between 1 and 100 are encapsulated with another S-VLAN (VLAN 30) tag when exiting Gigabit Ethernet port 0/0/2.

Similarly, for double-tagged (S-VLAN = 30, C-VLAN = 1–100) packets coming from the provider network, the **rewrite ingress pop 1 symmetric** command causes the outer S-VLAN tag to be popped and the original C-VLAN tagged frame to be forwarded over bridge-domain 4000 out to .

The same scenario applies to Customer B.

Customer B

Method 2

QinQ is also supported when sending packets between an EFP and a trunk EFP. The same external behavior as Method 1 can be achieved with this configuration:

Customer A

Again, service instance 1 on is configured with the VLAN encapsulations used by the customer. These are forwarded on bridge-domain 30. The service provider facing port is configured as a trunk port. The trunk port pushes a tag matching the bridge-domain that the packet is forwarded on (in this case S-VLAN 30).

For double tagged (S-VLAN = 30, C-VLAN = 1 to 100) packets coming in from the provider network, the trunk port pops the outer S-VLAN (30) and forwards the packet on that bridge-domain.

Customer B

You can also combine the customer A and B configurations, as follows:

Customer A and B

For information about the effect on cost of service (CoS) for different EFT tagging operations, see the .

Layer 2 Protocol Tunneling

Customers at different sites connected across a service-provider network need to use various Layer 2 protocols to scale their topologies to include all remote sites, as well as the local sites. STP must run properly, and every VLAN should build a proper spanning tree that includes the local site and all remote sites across the service-provider network. Cisco Discovery Protocol (CDP) must discover neighboring Cisco devices from local and remote sites.

VLAN Trunking Protocol (VTP) must provide consistent VLAN configuration throughout all sites in the customer network that are participating in VTP. Similarly, DTP, LACP, LLDP, PAgP, and UDLD can also run across the service-provider network.

When protocol tunneling is enabled, edge switches on the inbound side of the service-provider network encapsulate Layer 2 protocol packets with a special MAC address (0100.0CCD.CDD0) and send them across the service-provider network. Core switches in the network do not process these packets but forward them as normal (unknown multicast data) packets. Layer 2 protocol data units (PDUs) for the configured protocols cross the service-provider network and are delivered to customer switches on the outbound side of the service-provider network. Identical packets are received by all customer ports on the same VLANs with these results:

- Users on each of a customer's sites can properly run STP, and every VLAN can build a correct spanning tree based on parameters from all sites and not just from the local site.
- CDP discovers and shows information about the other Cisco devices connected through the service-provider network.
- VTP provides consistent VLAN configuration throughout the customer network, propagating to all switches through the service provider that support VTP.

Customers use Layer 2 protocol tunneling to tunnel BPDUs through a service-provider network without interfering with internal provider network BPDUs.

In figure below, Customer X has four switches in the same VLAN, which are connected through the service-provider network. If the network does not tunnel PDUs, switches on the far ends of the network cannot properly run STP, CDP, and other Layer 2 protocols. For example, STP for a VLAN on a switch in Customer X, Site 1, will build a spanning tree on the switches at that site without considering convergence parameters based on Customer X's switch in Site 2. This could result in the topology shown in figure below.

Figure 1: Layer 2 Protocol Tunneling

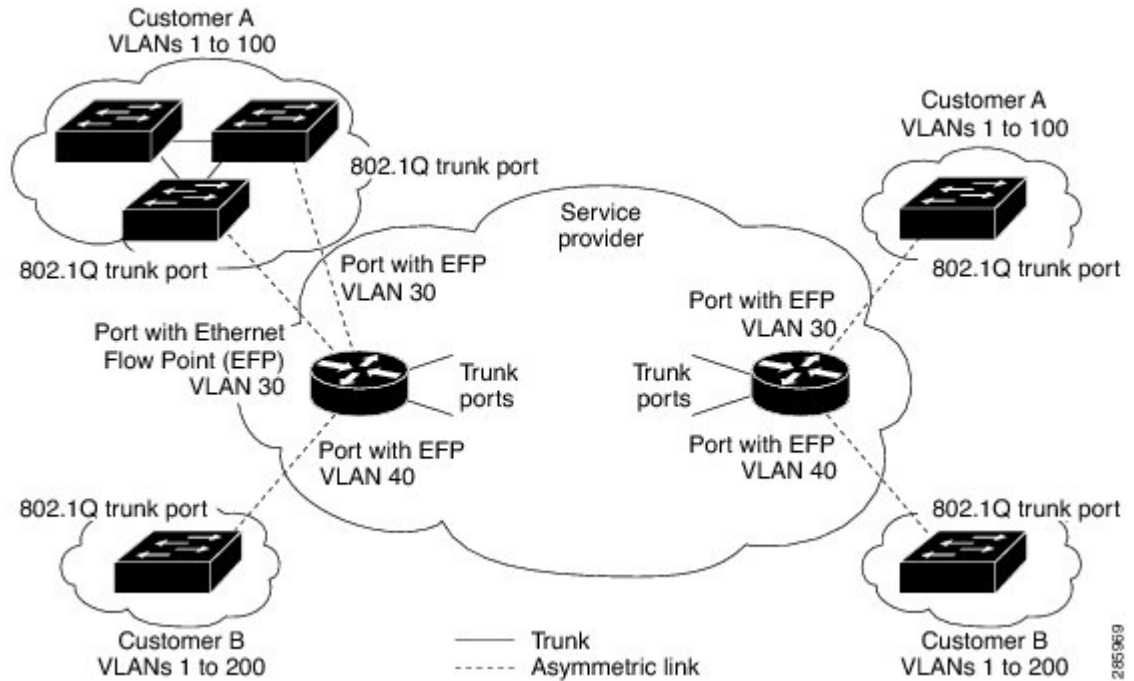
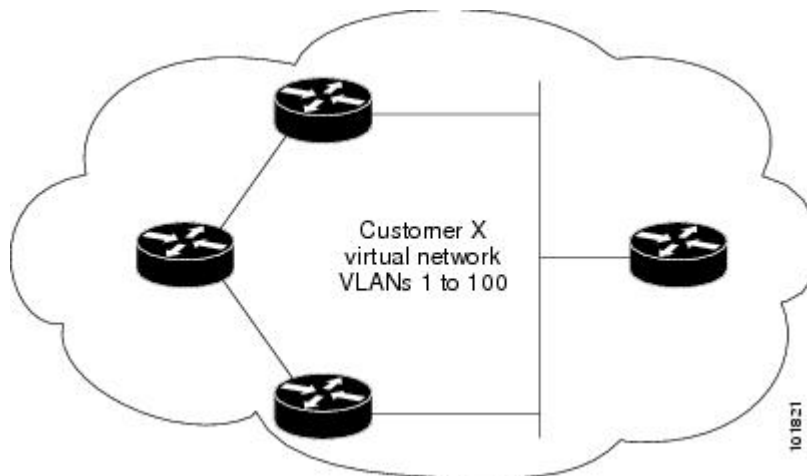


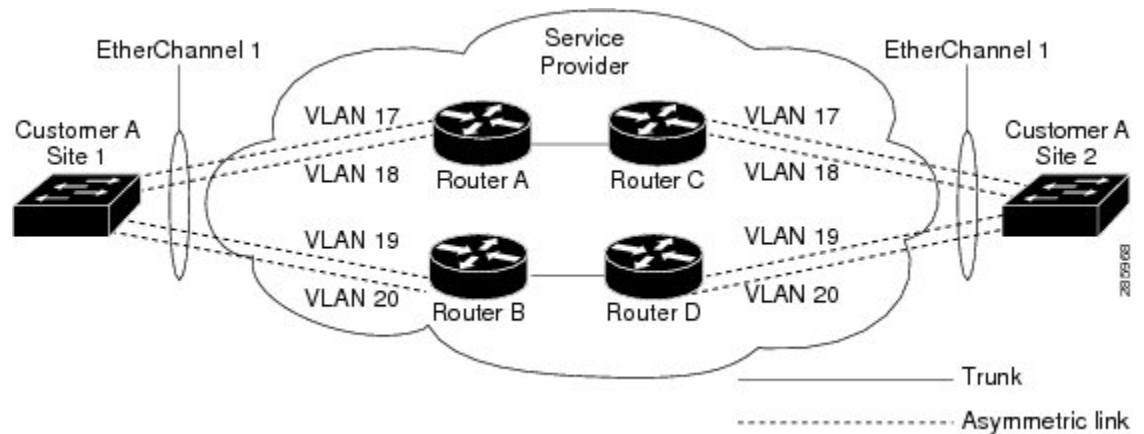
Figure 2: Layer 2 Network Topology without Proper Convergence



In a service-provider network, you can use Layer 2 protocol tunneling to enhance the creation of EtherChannels by emulating a point-to-point network topology. When you enable protocol tunneling (PAGP or LACP) on the service-provider switch, remote customer switches receive the PDUs and can negotiate the automatic creation of EtherChannels.

For example, in figure below, Customer A has two switches in the same VLAN that are connected through the SP network. When the network tunnels PDUs, switches on the far ends of the network can negotiate the automatic creation of EtherChannels without needing dedicated lines

Figure 3: Layer 2 Protocol Tunneling for EtherChannels



Use the **`l2protocol tunnel protocol`** service-instance configuration command to enable Layer 2 protocol tunneling on a service instance

Valid protocols include CDP, LACP, LLDP, PAgP, STP, UDLD, and VTP. If a protocol is not specified for a service instance, the protocol frame is dropped at the interface.

This is an example of Layer 2 protocol tunneling configuration:

```
Router (config)# interface gigabitethernet0/2
Router (config-if)# service instance 10 Ethernet
Router (config-if-srv)# encapsulation untagged , dot1q 200 second-dot1q 300
Router (config-if-srv)# l2protocol tunnel cdp stp vtp pagp lacp
Router (config-if-srv)# bridge-domain 10
```



Note To enable tunneling of most Layer 2 protocol, you must configure **encapsulation untagged** because Layer 2 protocol PDUs are usually untagged.

Bridge Domain Routing

The switch supports IP routing and multicast routing for bridge domains, including Layer 3 and Layer 2 VPNs, using the BDI model. There are the limitations:

- You must configure BDIs for bridge-domain routing.
- The bridge domain must be in the range of 1 to 4094 to match the supported VLAN range.
- You can use bridge domain routing with only native packets.

Bridge domain routing only works if proper tag popping is configured on the corresponding EFP BD. For example, if an EFP is configured with a single tag then **rewrite** should be **pop 1 symmetric**. If the EFP is configured with double tag then **rewrite** should be **pop 2 symmetric**. For double tag EFP, **pop 1 symmetric** and routing on the BDI is not supported.



Note Traffic engineering is not supported for BDI Routing.

This is an example of configuring bridge-domain routing with a single tag EFP:

```
Router (config)# interface gigabitethernet0/2
Router (config-if)# service instance 1 Ethernet
Router (config-if-srv)# encapsulation dot1q 10
Router (config-if-srv)# rewrite ingress tag pop 1 symmetric
Router (config-if-srv)# bridge-domain 100

Router (config)# interface bdi 100
Router (config-if)# ip address 20.1.1.1 255.255.255.255
```

This is an example of configuring bridge-domain routing with two tags:

```
Router (config)# interface gigabitethernet0/2
Router (config-if)# service instance 1 Ethernet
Router (config-if-srv)# encapsulation dot1q 10 second-dot1q 20
Router (config-if-srv)# rewrite ingress tag pop 2 symmetric
Router (config-if-srv)# bridge-domain 100

Router (config)# interface bdi 100
Router (config-if)# ip address 20.1.1.1 255.255.255.255
```

EFPs and Trunk Port MAC Addresses

Because forwarding can occur between EFPs and trunk ports, MAC address movement can occur on learned addresses. Addresses learned on EFPs will have the format of interface + EFP ID, for example gigabitethernet 0/0/1 + EFP 1. When an address moves between a non-secured EFP and a trunk port, the behavior is similar to that of moving between trunk ports.

To see MAC address information for bridge domains, use the **show mac-address-table bdomain domain** command.

When an EFP property changes (bridge domain, rewrite, encapsulation, split-horizon, secured or unsecured, or a state change), the old dynamic MAC addresses are flushed from their existing tables. This is to prevent old invalid entries from lingering.

EFPs and MSTP

EFP bridge domains are supported by the Multiple Spanning Tree Protocol (MSTP). These restrictions apply when running STP with bridge domains.

- EVC supports only MSTP.
- All incoming VLANs (outer-most or single) mapped to a bridge domain must belong to the same MST instance or loops could occur.
- For all EFPs that are mapped to the same MST instance, you must configure backup EFPs on every redundant path to prevent loss of connectivity due to STP blocking a port.
- When STP mode is PVST+ or PVRST, EFP information is not passed to the protocol. EVC only supports only MSTP.
- Changing STP mode from MST to PVRST is not allowed.

MAC Address Forwarding, Learning and Aging on EFPs

- Layer 2 forwarding is based on the bridge domain ID and the destination MAC address. The frame is forwarded to an EFP if the binding between the bridge domain, destination MAC address, and EFP is known. Otherwise, the frame is flooded to all the EFPs or ports in the bridge domain.
- MAC address learning is based on bridge domain ID, source MAC addresses, and logical port number. MAC addresses are managed per bridge domain when the incoming packet is examined and matched against the EFPs configured on the interface. If there is no EFP configured, the bridge domain ID equal to the outer-most VLAN tag is used as forwarding and learning look-up key.

If there is no matching entry in the Layer 2 forwarding table for the ingress frame, the frame is flooded to all the ports within the bridge domain. Flooding within the bridge domain occurs for unknown unicast, unknown multicast, and broadcast.

- Dynamic addresses are addresses learned from the source MAC address when the frame enters the router. All unknown source MAC addresses are sent to the CPU along with ingress logical port number and bridge domain ID for learning. Once the MAC address is learned, the subsequent frame with the destination MAC address is forwarded to the learned port. When a MAC address moves to a different port, the Layer 2 forwarding entry is updated with the corresponding port.



Note The router does not currently support the **no mac address-table** learning bridge-domain *bridge-id* global configuration command.

- Dynamic addresses are aged out if there is no frame from the host with the MAC address. If the aged-out frame is received by the switch, it is flooded to the EFPs in the bridge domain and the Layer 2 forwarding entry is created again. The default for aging dynamic addresses is 5 minutes. However, when MST undergoes a topology change, the aging time is reduced to the *forward-delay* time configured by the spanning tree. The aging time reverts back to the last configured value when the topology change expires.

You can configure a dynamic address aging time per bridge domain using the **mac aging-time time** command. The range is in seconds and valid values are 10-600. The default value is 300. An aging time of 0 means that the address aging is disabled.

- MAC address movement is detected when the host moves from one port to another. If a host moves to another port or EFP, the learning lookup for the installed entry fails because the ingress logical port number does not match and a new learning cache entry is created. The detection of MAC address movement is disabled for static MAC addresses where the forwarding behavior is configured by the user.

Configuring a Static MAC Address

This section describes how to configure a static MAC address on the router.

Static MAC Addresses

The router supports multicast static MAC addresses, which allow you to enable multicast at the layer 2 level. You can use multicast static MAC addresses to forward multicast packets to specific EFPs on a network.

Limitations

The following limitations apply when configuring static MAC addresses:

- Static MAC addresses are supported only on egress ports.
- You can configure up to 1024 multicast static MAC addresses
- You can assign up to 24 EFPs to a bridge domain configured with a multicast static MAC address.
- MAC entries configured across different bridge-domains are represented as separate entries in the router MAC table.
- Multicast static MAC addresses apply only to layer 2 traffic; layer 3 multicast traffic is not affected by a static MAC configuration and is forwarded to all EFPs in a bridge domain.

Configuring a Static MAC Address

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: Router# configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i> Example: Router(config)# interface gigabitethernet 0/0/0	Specify the port to attach to the policy map, and enter interface configuration mode. Valid interfaces are physical ports.
Step 3	no ip address Example: Router(config-if)# no ip address	To set a primary or secondary IP address for an interface, use the ip address interface configuration command. To remove an IP address or disable IP processing, use the no form of this command.
Step 4	no negotiation auto Example: Router(config-if)# no negotiation auto	Disables autonegotiation on Gigabit Ethernet interfaces.
Step 5	service instance <i>number</i> ethernet [<i>name</i>] Example: Rotuer(config-if)# service instance 1 ethernet	Configure an EFP (service instance) and enter service instance configuration) mode. <ul style="list-style-type: none"> • The number is the EFP identifier, an integer from 1 to 4000. • (Optional) ethernet name is the name of a previously configured EVC. You do not need to use an EVC name in a service instance.

	Command or Action	Purpose
Step 6	<p>encapsulation {default dot1q priority-tagged untagged}</p> <p>Example:</p> <pre>Router(config-if-srv)# encapsulation dot1q 100</pre>	<p>Configure encapsulation type for the service instance.</p> <ul style="list-style-type: none"> • default—Configure to match all unmatched packets. • dot1q—Configure 802.1Q encapsulation. See for details about options for this keyword. • priority-tagged—Specify priority-tagged frames, VLAN-ID 0 and CoS value of 0 to 7. • untagged—Map to untagged VLANs. Only one EFP per port can have untagged encapsulation.
Step 7	<p>bridge-domain <i>bridge-id</i> [split-horizon group <i>group-id</i>]</p> <p>Example:</p> <pre>Router(config-if-srv)# bridge-domain 100</pre>	<p>Configure the bridge domain ID. The range is from 1 to 4000.</p> <p>You can use the split-horizon keyword to configure the port as a member of a split horizon group. The <i>group-id</i> range is from 0 to 2.</p>
Step 8	<p>mac static address <i>address</i></p> <p>Example:</p> <pre>Router(config-if-srv)# mac static address 0000.bbbb.cccc</pre>	<p>Specifies the multicast MAC address.</p>
Step 9	<p>end</p> <p>Example:</p> <pre>Router(config-if-srv)# end</pre>	<p>Return to privileged EXEC mode.</p>

Monitoring EVC

Table 1: Supported show Commands

	Description
<p>show ethernet service evc [id <i>evc-id</i> interface <i>interface-id</i>] [detail]</p>	<p>Displays information about all EVCs, or a specific EVC when you enter an EVC ID, or all EVCs on an interface when you enter an interface ID. The detail option provides additional information about the EVC.</p>

	Description
show ethernet service instance [id <i>instance-id</i> interface <i>interface-id</i> interface <i>interface-id</i>] { [detail] [stats]}	Displays information about one or more service instance (EFPs). If you specify an EFP ID and interface, only data pertaining to that particular EFP is displayed. If you specify only an interface ID, data is displayed for all EFPs on the interface.
show bridge-domain [<i>n</i>]	When you enter <i>n</i> , this command displays all the members of the specified bridge-domain, if a bridge-domain with the specified number exists. If you do not enter <i>n</i> , the command displays all the members of all bridge-domains in the system.
show bridge-domain <i>n</i> split-horizon [group { <i>group_id</i> all }]	When you do not specify a group <i>group_id</i> , this command displays all the members of bridge-domain <i>n</i> that belong to split horizon group 0. If you specify a numerical <i>group_id</i> , this command displays all the members of the specified group id. When you enter group all , the command displays all members of any split horizon group.
show ethernet service instance detail	This command displays detailed service instance information, including Layer 2 protocol information. This is an example of the output: Router# show ethernet service instance detail Service Instance ID: 1 Associated Interface: Ethernet0/0 Associated EVC: L2protocol tunnel pagp CE-Vlans: State: Up EFP Statistics: Pkts In Bytes In Pkts Out Bytes Out 0 0 0 0
show mac address-table	This command displays dynamically learned or statically configured MAC security addresses.
show mac address-table bridge-domain <i>bridge-domain id</i>	This command displays MAC address table information for the specified bridge domain.
show mac address-table count bridge-domain <i>bridge-domain id</i>	This command displays the number of addresses present for the specified bridge domain.
show mac address-table learning bridge-domain <i>bridge-domain id</i>	This command displays the learning status for the specified bridge domain.

This is an example of output from the **show ethernet service instance detail** command:

```

Router#
Service Instance ID: 1
Associated Interface:
Associated EVC: EVC_P2P_10
L2protocol drop
CE-Vlans:
Encapsulation: dot1q 10 vlan protocol type 0x8100
Interface Dot1q Tunnel Ethertype: 0x8100
State: Up
EFP Statistics:
  Pkts In   Bytes In   Pkts Out   Bytes Out
    214     15408     97150     6994800
EFP Microblocks:
*****
Microblock type: Bridge-domain
Bridge-domain: 10

```

This is an example of output from the **show bridge-domain** command:

```

Router# show bridge-domain 100
Bridge-domain 100 (1 ports in all)
State: UP Mac learning: Enabled
Aging-Timer: 300 second(s)
Maximum address limit: 256000
GigabitEthernet0/0/0 service instance 1

```

Nile Mac Address Entries

```

BD mac addr type ports
-----
100 0000.bbbb.cccc STATIC Gi0/0/0.Efp1

```

```
sh mac-address-table bdomain 100
```

Nile Mac Address Entries

```

BD mac addr type ports
-----
100 0000.bbbb.cccc STATIC Gi0/0/0.Efp1

```

This is an example of output from the **show ethernet service instance** statistics command:

```

Router#
Service Instance 1, Interface
Pkts In   Bytes In   Pkts Out   Bytes Out
    214     15408     97150     6994800

```

This is an example of output from the **show mac-address table count** command:

```

Router# show mac address-table count bdomain 10

Mac Entries for BD 10:
-----
Dynamic Address Count : 20
Static Address Count  : 0
Total Mac Addresses   : 20

```

