


[Virtual Leased Lines - Martini](#)

“Martini” Drafts

[draft-martini-l2circuit-encap-mpls-04.txt](#) defines the handling and encapsulation of layer two packets.

[draft-martini-l2circuit-trans-mpls-08.txt](#) defines the signaling and forwarding of traffic across the MPLS network.

In order to provide L2 VPN services across an MPLS network a methodology needed to be defined. The two martini drafts define the different encapsulation techniques accompanied by signaling and transport functions. They combine to form the basis for point-to-point layer two services across an MPLS network, equivalent in concept to ATM or Frame Relay. The point-to-point service is facilitated through a pair of LSPs in opposite directions, which form a single virtual pipe. Label stacking is used to create hierarchies separating the common tunnel LSP and the virtual channels that exist inside. The *VC Label* represents each virtual channel. This label is used by the egress in order to map the channel to the individual *Group ID and Virtual Circuit ID, or VCID*. If present the Group ID is a collection of VC IDs, which identifies the different services within a single VC Label. If the Group ID is not present the VC ID represents the service.



Tunnel LSP Creation

The *Tunnel LSP* is the point-to-point connection within which individual customer virtual circuits will exist. The tunnel LSP scales the network core by aggregating many virtual circuits into a single common tunnel LSP. It is not feasible to have a tunnel per customer, or even worse, a tunnel per customer per service. A model that followed that direction is doomed to run into scalability issues early on in the deployment. The protocol with which to create the tunnel LSP is left to the discretion of the implementer and is largely based on the question, which protocol services the needs better. For example, if traffic engineering, including the signaling explicit paths is required, RSVP-TE provides the solution. However, if it is acceptable to allow the IGP to make the hop-by-hop decisions and end-to-end path significance is not required LDP may be acceptable. The RS platform supports LDP over LDP and LDP over RSVP-TE, so either LDP or RSVP-TE can be used for the tunnel LSP.

Signaling the VC Label

A peering session must exist for LDP information to be exchanged between LDP capable nodes. Since the edges of the network represent remote LDP peers the configuration must explicitly instruct nodes to form a peering relationship using extended discovery. Once the peering session has been established Virtual Circuit Labels, *VC Labels*, can be exchanged to identify the components of the individual L2-FEC. The signaling of the VC Label is performed using LDP extensions to the *Label Mapping* message. When an edge router is configured with a new L2-FEC, a local LDP label is selected from the database and using the extensions to

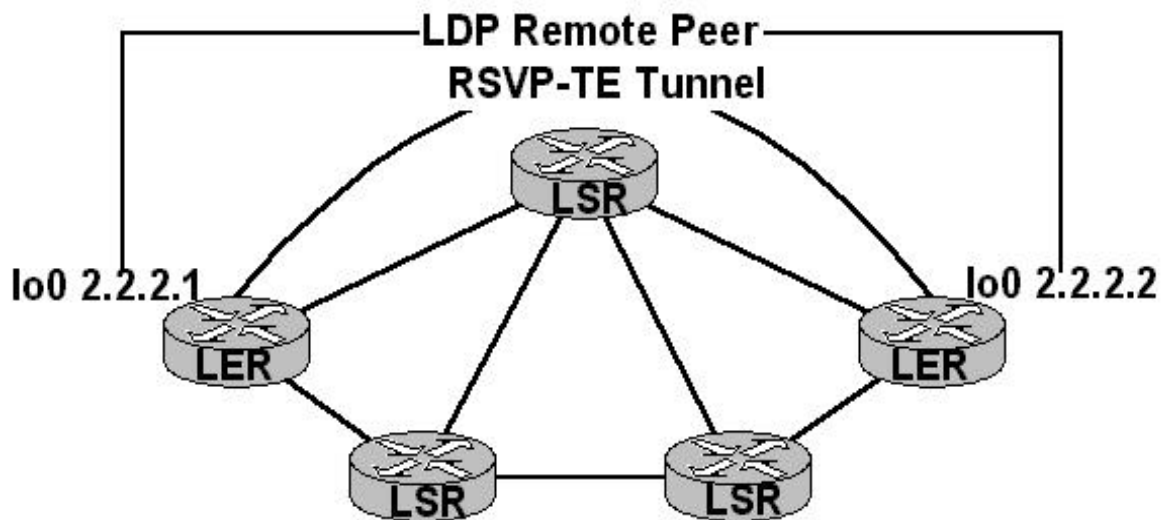
the Label Mapping Message, all the information relating to the L2-FEC (Group ID, VC ID, Interface and other) is forwarded to the remote peer.

The RS supports the following types of layer two FECs.

RS L2-FEC	Description	Group ID	VC ID
VLAN	802.1Q VLAN		Value of VLAN ID
Customer-Id	Physical Port		Value of Customer-id configured
Customer-Id, VLAN	Physical Port, 802.1Q VLAN	Value of Customer-id configured	Value of VLAN ID

Monitoring the Signaling Process

Starting with the base network having only the tunnel LSP and the remote LDP peering session established, we examine the effects on the LDP database as we start to add L2-FECs to the router configurations. No L2-Fecs have been created at this time.



With just the base network configured, the LDP databases should receive (Input label database) the remote peers loopback address and send (Output label database) the remote peer the local loopback address.

lo0 2.2.2.1	lo0 2.2.2.2
LER1# ldp show database	LER2# ldp show database
Input label database, 2.2.2.1:0-2.2.2.2:0	Input label database, 2.2.2.2:0-2.2.2.1:0
Label Prefix	Label Prefix
3 2.2.2.2/32	3 2.2.2.1/32
Output label database, 2.2.2.1:0-2.2.2.2:0	Output label database, 2.2.2.2:0-2.2.2.1:0
Label Prefix	Label Prefix
3 2.2.2.1/32	3 2.2.2.2/32

If the following L2-FEC is added to the configuration on router 2.2.2.1...

```
ldp map ports et.3.(5-6) customer-id 2001
ldp add l2-fec customer-id 2001 to-peer 2.2.2.2
```

A *Label Mapping* message will be sent to peer 2.2.2.2 binding label to FEC, in this case the VC ID, "customer-id 2001".

```
LDP Send Label Mapping msg on Interface lo, Remote Neighbor
Our LDP Id: 2.2.2.1:0 Peer LDP Id: 2.2.2.2:0
Label: 2048
FEC:
Customer-id 2001
```

The LDP database will reflect the addition of the FEC. The L2-FEC must be configured on both routers to allow for traffic to reach the ultimate destination. If the L2-FEC is only created on one of the LDP peers the packets will reach the destination with the appropriate VC Label however there are no associated output ports possible for the received L2-FEC. The packet would be dropped at the egress node.

lo0 2.2.2.1	lo0 2.2.2.2
LER1# ldp show database	LER2# ldp show database
Input label database, 2.2.2.1:0-2.2.2.2:0	Input label database, 2.2.2.2:0-2.2.2.1:0
Label Prefix	Label Prefix
3 2.2.2.2/32	2048 Customer ID 2001
	3 2.2.2.1/32
Output label database, 2.2.2.1:0-2.2.2.2:0	Output label database, 2.2.2.2:0-2.2.2.1:0
Label Prefix	Label Prefix
2048 Customer ID 2001	3 2.2.2.2/32
3 2.2.2.1/32	

Adding the complimentary configuration on router 2.2.2.2 will allow it to properly forward traffic out a specific set of ports to the ultimate destination.

```
ldp map ports et.5.1 customer-id 2001
ldp add l2-fec customer-id 2001 to-peer 2.2.2.1
```

lo0 2.2.2.1	lo0 2.2.2.2
LER1# ldp show database	LER2# ldp show database
Input label database, 2.2.2.1:0-2.2.2.2:0	Input label database, 2.2.2.2:0-2.2.2.1:0
Label Prefix	Label Prefix
2048 Customer ID 2001	2048 Customer ID 2001
3 2.2.2.2/32	3 2.2.2.1/32
Output label database, 2.2.2.1:0-2.2.2.2:0	Output label database, 2.2.2.2:0-2.2.2.1:0
Label Prefix	Label Prefix
2048 Customer ID 2001	2048 Customer ID 2001
3 2.2.2.1/32	3 2.2.2.2/32

Forwarding Packets

Ingress Node: The ingress node is responsible for classifying inbound traffic based on the configured L2-FEC information. Once the traffic has been classified the ingress router strips the preamble and FCS and pushes the two labels found in the output tag table. The bottom of stack label represents the VC Label and the top level label represents the tunnel LSP.

Transit Routers: The transit routers are oblivious to any tunneling that is occurring beneath the top level label. They perform their label operation based on the tunnel label and never see the VC Label inside.

Egress Router: The egress router may receive a packet that has already had the tunnel label popped by the penultimate router or receive the complete stack if penultimate hop popping is disabled. Regardless, the egress router uses the VC Label to point to the Group Id and VC ID and determines which local ports are included as possible output ports. Once that is determined, regular L2 functions apply, learning and bridging etc.

Creating and Signaling Group ID & VC ID

The RS platform creates and signals the Group ID & VC ID when an L2-FEC is created locally and then added to a remote peer. Basically the process can be broken down into three steps.

- Create the L2-FEC
- Signal the L2-FEC
- Optionally set any transport level preferences

The three different approaches will be detailed below. Before starting the discussion on configuring the three L2-FEC types it is important to realize all the MPLS ports connecting the edge LER to the core MPLS network should be configured as 802.1Q trunks. This allows the MPLS ports to be part of multiple customer VLANs. There are two options when defining these ports as 802.1Q trunk-ports. The packets sent out into the MPLS core can either carry a 802.1Q tag as part of the layer two information that precedes the MPLS label if layer two switches are deployed between the LER and LSR, or the “untagged” option which does not insert an 802.1Q tag between the layer two information and the MPLS label. The untagged option is deployed where the next hop is an MPLS aware device, switching on MPLS label information.

To set the MPLS ports to 802.1Q trunk ports...

```
RS(config)# vlan make trunk-port <port>
```

To send the packet out to the MPLS cloud without the outer VLAN header information...

```
RS(config)# vlan make trunk-port <port> untagged
```

1. L2-FEC Physical Port: The ability to map a port or group of ports to unique identifier, isolating those ports from the rest of the ports.

Using the example below, packets that arrive on port et.3.5, not destined to a local host, will be encapsulated, less the FCS and preamble, inside another layer two header followed by a stack of two MPLS labels. When using port based solutions it is important to note that port may receive 802.1Q tagged or untagged packets. By default the RS does not accept both 802.1Q tagged and untagged packets on the same port. In order to allow a port to act as a hybrid *Native VLAN* must be configured.

Case I - Traditional Access Port (Non-802.1Q tagged packets accepted, tagged packets dropped)

```
vlan make trunk-port gi.2.(1-2) untagged (set MPLS ports on LER to 802.1q capable but do not send 802.1Q tag)
```

```

ldp map ports et.3.(5-6) customer-id 2001          (map the ports to
an identifier)
ldp add l2-fec customer-id 2001 to-peer 2.2.2.2    (signal VC Label
vcid=customer-id 2001} to the remote peer)

```

Case II - Traditional Trunk-Port (802.1Q tagged packets accepted, untagged packets dropped)

```

vlan make trunk-port gi.2.(1-2) untagged          (set MPLS ports on
LER to 802.1q capable but do not send 802.1Q tag)
vlan make trunk-port et.3.(5-6)                  (accept 802.1Q tag
on the customer facing ports)
ldp map ports et.3.(5-6) customer-id 2001        (map the ports to
an identifier)
ldp add l2-fec customer-id 2001 to-peer 2.2.2.1  (signal VC Label
and {vcid customer-id 2001}to the remote peer)

```

Case III – Hybrid Ports (Both 802.1Q tagged and untagged accepted)

This case requires *Native VLAN* support. When the Native VLAN configuration is deployed it means the provider does not have to know whether the customer traffic arriving on the port is or is not 802.1Q tagged, nor does the provider have to be aware of the customer VLAN IDs. This will be completely transparent to both the provider and the customer. Traffic that arrives will ultimately arrive at the destination with the exact same disposition it entered with. The Native VLAN acts as a catch all for those packets that arrive without an 802.1Q tags, inferring a VLAN for it.

To configure the native VLAN...

```

RS(config)# vlan set native-vlan <ports> <protocol|all> <VLAN name>

```

The following maps all packets that arrive on ports et.3.(1-8) without a tag to the DEFAULT VLAN (1). Notice, for the ports that are to be configured for hybrid, they must be configured as trunk ports first. **WARNING:** VLAN names are case sensitive! When coding a VLAN name, DEFAULT is not the same as default.

```

vlan make trunk-port et.3.(1-8)
vlan set native-vlan et.3.(1-8) all DEFAULT

```

It is not required to use the DEFAULT VLAN as the native VLAN. Any VLAN can be specified as the native VLAN. However, a conflict will arise if an 802.1Q labeled packet arrives on a native VLAN port with the a VLAN tag in the packet that is the same as the configured Native VLAN.

```

vlan make trunk-port et.3.(1-8)
vlan create Native port-based id 4093
vlan add ports et.3.(1-8) to Native
vlan set native-vlan et.3.(1-8) all Native

```

Applying the hybrid approach to the L2 VPN service, the following allows customers to send 802.1Q tagged packets or untagged packets.

```

vlan make trunk-port gi.2.(1-2) untagged          (set MPLS ports on LER to
802.1q capable sending packets untagged)
vlan make trunk-port et.3.(1-8)                  (accept 802.1Q tag)
vlan set native-vlan et.3.(1-8) all DEFAULT      (set the ports to apply
native DEFAULT VLAN when 802.1Q tag not present)
ldp map ports et.3.(5-6) customer-id 2001        (create the L2-FEC)
ldp add l2-fec customer-id 2001 to-peer 2.2.2.1 (signal VC Label and
{vcid customer-id 2001}to peer)
ldp map ports et.3.8 customer-id 4000           (create the L2-FEC)

```

```
ldp add l2-fec customer-d 4000 to-peer 2.2.2.1 (signal VC Label and
{vcid customer-id 4000} to peer)
```

2. L2-FEC VLAN ID: When this is used as the classification mechanism the VLAN ID value represents the VCID. Typically, this approach is used when the provider assigns a VLAN-ID on a per customer basis. For this case, a VLAN must be created for the customer and the trunk port that connects to this customer and the MPLS enabled ports facing the core must be added to the customer defined VLAN. The connection can either be direct to the customer or a connection to a metro network that supports many subscribers of that provider.

In this example the trunk-port connects to a metro network using the et.3.3 interface. The VLAN ID 1001 and 1002 are extended to the remote peer 2.2.2.2.

```
vlan make trunk-port gi.2.(1-2) untagged (set MPLS ports on
LER to 802.1q capable sending packets untagged)
vlan make trunk-port et.3.3 (customer facing -
accept 802.1Q tag)
vlan create Cust1001 port-based id 1001 (create customer
VLAN)
vlan add port et.3.3 to Cust1001 (add the customer
facing ports to customer VLAN)
vlan add ports gi.2.(1-2) (add the MPLS ports
to customer VLAN)
vlan create Cust1002 port-based id 1002 (create customer
VLAN)
vlan add port et.3.3 to Cust1002 (add the customer
facing ports to customer VLAN)
vlan add ports gi.2.(1-2) (add the MPLS ports
to customer VLAN)
ldp add l2-fec vlan 1001 to-peer 2.2.2.2 (signal VC Label
and {vcid customer-id 1001}to peer)
ldp add l2-fec vlan 1002 to-peer 2.2.2.2 (signal VC Label
and {vcid cusomter-id 1002}to peer)
```

3. L2-FEC Physical Port and VLAN: Finally, there is the a physical port and a VLAN ID to al L2-FEC. In this case the VC ID is the VLAN ID and the Group-ID is the customer-id. This approach allows a single customer to use multiple VLAN IDs without depleting the provider VLAN space. The Group-id (customer-id) engulfs all the VC ID (VLAN) information within it. The provider can offer site-specific VLAN significance within a customer-id. The port that connects to an individual customer or shared metro is configured as a trunk, with the port assigned an encapsulating group-ID. This example distributes the traffic to specific sites based n VLAN information within the customer-id.

```
vlan make trunk-port et.3.8 (customer
facing - accept 802.1Q tag)
ldp map ports et.3.8 customer-id 3000 (create the
L2-FEC)
ldp add l2-fec vlan 100 customer-id 3000 to-peer 2.2.2.2 (signal VC 1
to peer 2.2.2.2)
ldp add l2-fec vlan 200 customer-id 3000 to peer 2.2.2.3 (signal VC
Label and {GID=customer-id 3000:VCID=Vlan 200}
to peer 2.2.2.3)
```

L2-Fec Transport LSP

By default, adding an L2-fec to a remote peer, will automatically select an LSP over which to signal the VC Label and Group-ID and VC-ID. The selection process depends on the number of label switched paths that exist between the remote LDP peers. If only a single LSP exists, the selection process is simple it uses that

one. If more than one exists LSP selection is based on preference. The lowest numerical preference the most preferred LSP. Should the preference be equal, the process checks to determine if a corresponding VC ID has been received from the remote LDP peer. If the VC Label mapping to the same Group-ID and VC ID has been received from a remote peer the same LSP will be selected.

In most cases it is beneficial, when more than one LSP exists, to map the traffic to a preferred LSP. This allows the provider to explicitly engineer which are the preferred paths in the network. This does not mean if the LSP fails it represents a single point of failure; alternates can be accepted should a failure occur.

To set specific l2-fec specific transport parameters...

RS(config)# ldp set l2-fec ?

```

ler1(config)# ldp set l2-fec ?
alternate-acceptable  - An alternate LSP (RSVP or LDP) is acceptable in case
                        the transport LSP is not active. If the transport
                        LSP comes up later, it will override the alternate
                        LSP selected.

customer-id           - Sets the transport LSP to be used for this
                        customer-id FEC. The vlan option if specified
                        additionally will select the transport LSP for this
                        customer-id, vlan combination FEC.

no-switchback         - Do not switch back from non-preferred LSP to
                        preferred LSP if preferred LSP comes up

to-peer               - Sets the transport LSP to this peer

transport-lsp         - Name of the RSVP LSP which should be used as a
                        transport LSP

vlan                  - Sets the transport LSP to be used for this vlan FEC.
                        The customer-id option if specified additionally
                        will select the transport LSP for this customer-id,
                        vlan combination FEC.
  
```

Specific to selecting the preferred LSP...

Option	Example	Action
Only one LSP specified in configuration	<i>ldp set l2-fec customer-id 2001 to-peer 2.2.2.2 transport-lsp TunLSP</i>	Use specified LSP. In event of LSP failure service interruption will occur.
LSP is specified but alternates are specified in the configuration	<i>ldp set l2-fec customer-id 2001 to-peer 2.2.2.2 transport-lsp TunLSP alternate-acceptable</i>	Use specified LSP. In the event of failure select any other LSP that terminates on the required remote peer.

Consider the following, no VC Labels have been exchanged between remote peers and the local router has decided the most preferred label switched path is a local decision. After the label exchange has been completed between two remote LDP peers with an established session the data starts to flow. By looking at the L2-FEC and the associated output tag table it is noticed all traffic for these L2-FEC are using a specific Transport LSP (TunLSP2) out a specific interface (To-LSR2).

```
LER1# ldp show l2-fec
FEC: Forward Equivalence class, in-lbl: Label received, out-lbl: Label sent
```

```
Remote neighbor 2.2.2.2:0
```

FEC name/label	in-lbl	out-lbl	Transport LSP
VLAN ID 1002	2048	2049	TunLSP2/17
VLAN ID 1001	2049	2048	TunLSP2/17
Customer ID 2001	-	2050	

```
LER1# mpls show ott-table
```

Interface	OTT	RefCount	HW-OTT	RefCount	NextHop	Vlan
-----------	-----	----------	--------	----------	---------	------

```
Labels
```

```
-----
--
lo          1    1      0    0      192.168.1.5  3    [17]
            2    1      0    0      0.0.0.0      1002
[2048]
            3    1      1    1      192.168.1.5  1002
[17|2048]
            4    1      0    0      0.0.0.0      1001
[2049]
            5    1      2    1      192.168.1.5  1001
[17|2049]
            6    1      0    0      0.0.0.0      0    [3]
```