CHAPTER 10

Configuring Link Aggregation on the ML-Series Card

This chapter describes how to configure link aggregation for the ML-Series cards, both EtherChannel and packet-over-SONET (POS) channel. For additional information about the Cisco IOS commands used in this chapter, refer to the *Cisco IOS Command Reference* publication.

This chapter contains the following major sections:

Understanding Link Aggregation

The ML-Series card offers both EtherChannel and POS channel. Traditionally EtherChannel is a trunking technology that groups together multiple full-duplex IEEE 802.3 Ethernet interfaces to provide fault-tolerant high-speed links between switches, routers, and servers. EtherChannel forms a single higher bandwidth routing or bridging endpoint and was designed primarily for host-to-switch connectivity. The ML-Series card extends this link aggregation technology to bridged POS interfaces. POS channel is only supported with LEX encapsulation.

Link aggregation provides the following benefits:

- Logical aggregation of bandwidth
- Load balancing
- Fault tolerance

Port channel is a term for both POS channel and EtherChannel. The port channel interface is treated as a single logical interface although it consists of multiple interfaces. Each port channel interfaces consists of one type of interface, either Fast Ethernet or POS. You must perform all port channel configurations on the port channel (EtherChannel or POS channel) interface rather than on the individual member Ethernet or POS interfaces. You can create the port channel interface by entering the **interface port-channel** interface configuration command.

Port channel connections are fully compatible with IEEE 802.1Q trunking and routing technologies. IEEE 802.1Q trunking can carry multiple VLANs across a port channel.

Each ML100-FX supports up to four FECs plus an additional POS channel, a port channel made up of the two POS ports. A maximum of four Fast Ethernet ports can bundle into one Fast Ethernet Channel (FEC) and provide bandwidth scalability up to 400-Mbps full-duplex Fast Ethernet.

Caution The EtherChannel interface is the Layer 2/Layer 3 interface. Do not enable Layer 3 addresses on the physical interfaces. Do not assign bridge groups on the physical interfaces because doing so creates loops.

Caution Before a physical interface is removed from an EtherChannel (port channel) interface, the physical interface must be disabled. To disable a physical interface, use the **shutdown** command in interface configuration mode.

Note Link aggregation across multiple ML-Series cards is not supported.

Note Policing is not supported on port channel interfaces.

Note The ML-Series does not support the routing of Subnetwork Access Protocol (SNAP) or Inter-Switch Link (ISL) encapsulated frames.

Configuring Link Aggregation

You can configure an FEC or POS channel by creating an EtherChannel interface (port channel) and optionally assigning a network IP address.

Configuring Fast EtherChannel

All interfaces that are members of an FEC should have the same link parameters, such as duplex and speed.

To create an EtherChannel interface, perform the following procedure, beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# interface port-channel` *channel-number*	Creates the EtherChannel interface.
Step 2	`Router(config-if)# ip address` *ip-address subnet-mask*	(Optional) Assigns an IP address and subnet mask to the EtherChannel interface.
Step 3	`Router(config-if)# end`	Exits to privileged EXEC mode.
Step 4	`Router# copy running-config startup-config`	(Optional) Saves configuration changes to NVRAM.

For information on other configuration tasks for the EtherChannel, refer to the *Cisco IOS Configuration Fundamentals Configuration Guide*.

To assign Ethernet interfaces to the EtherChannel, perform the following procedure, beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# interface fastethernet number`	Enters one of the interface configuration modes to configure the Fast Ethernet interface that you want to assign to the EtherChannel.
Step 2	`Router(config-if)# channel-group channel-number`	Assigns the Fast Ethernet interface to the EtherChannel. The channel number must be the same channel number you assigned to the EtherChannel interface.
Step 3	`Router(config-if)# end`	Exits to privileged EXEC mode.
Step 4	`Router# copy running-config startup-config`	(Optional) Saves configuration changes to NVRAM.

EtherChannel Configuration Example

Figure 10-1 shows an example of encapsulation over EtherChannel. The associated commands are provided in Example 10-1 and Example 10-2.

Figure 10-1 Encapsulation over EtherChannel Example

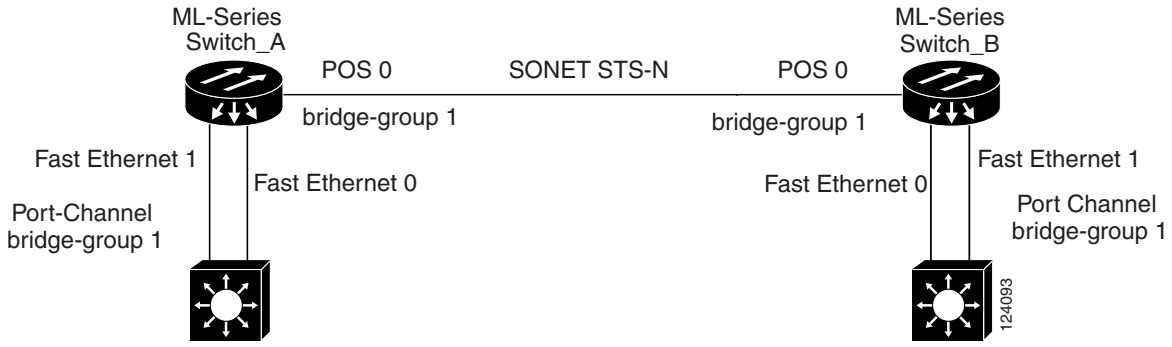

Example 10-1 ML_Series A Configuration

```
hostname Switch A
no ip routing
!
bridge 1 protocol ieee
!
interface Port-channel 1
bridge-group 1
 hold-queue 150 in
!
interface FastEthernet 0
channel-group 1
!
```

```
interface FastEthernet 1
channel-group 1
!
interface POS 0
bridge-group 1
```

Example 10-2 ML_Series B Configuration

```
hostname Switch B
no ip routing
!
bridge 1 protocol ieee
!
interface Port-channel 1
bridge-group 1
 hold-queue 150 in
!
interface FastEthernet 0
channel-group 1
!
interface FastEthernet 1
channel-group 1
!
interface POS 0
bridge-group 1
!
```

Configuring POS Channel

You can configure a POS channel by creating a POS channel interface (port channel) and optionally assigning an IP address. All POS interfaces that are members of a POS channel should have the same port properties and be on the same ML-Series card.

Note POS channel is only supported with G-Series card compatible (LEX) encapsulation.

To create a POS channel interface, perform the following procedure, beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# interface port-channel` *channel-number*	Creates the POS channel interface. You can configure one POS channel on the ML-Series card.
Step 2	`Router(config-if)# ip address` *ip-address subnet-mask*	Assigns an IP address and subnet mask to the POS channel interface (required only for the Layer 3 POS channel).
Step 3	`Router(config-if)# end`	Exits to privileged EXEC mode.
Step 4	`Router# copy running-config startup-config`	(Optional) Saves configuration changes to NVRAM.

Caution The POS channel interface is the routed interface. Do not enable Layer 3 addresses on any physical interfaces. Do not assign bridge groups on any physical interfaces because doing so creates loops.

To assign POS interfaces to the POS channel, perform the following procedure, beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# interface pos` *number*	Enters the interface configuration mode to configure the POS interface that you want to assign to the POS channel.
Step 2	`Router(config-if)# channel-group` *channel-number*	Assigns the POS interface to the POS channel. The channel number must be the same channel number that you assigned to the POS channel interface.
Step 3	`Router(config-if)# end`	Exits to privileged EXEC mode.
Step 4	`Router# copy running-config startup-config`	(Optional) Saves the configuration changes to NVRAM.

POS Channel Configuration Example

Figure 10-2 shows an example of POS channel configuration. The associated code for ML_Series A is provided in Example 10-3 and for ML_Series B in Example 10-4.

Figure 10-2 POS Channel Example

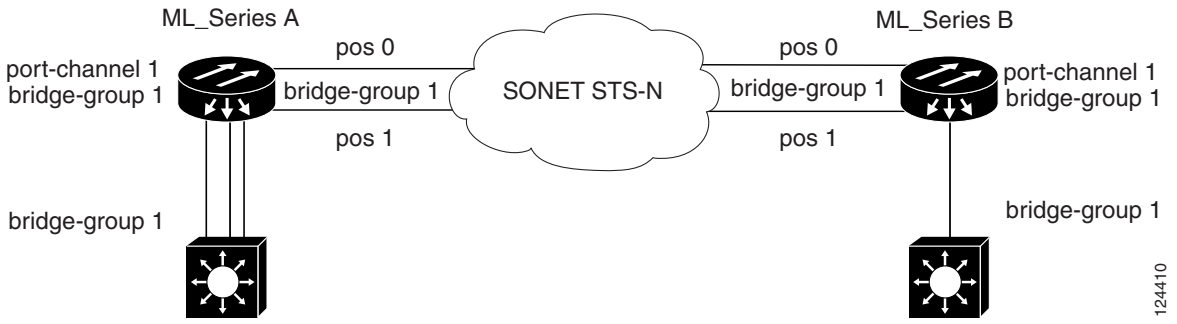

Example 10-3 ML_Series A Configuration

```
no ip routing
bridge 1 protocol ieee
!
!
interface Port-channel1
 no ip address
bridge-group 1
!
interface FastEthernet0
 no ip address
 bridge-group 1
```

```
!
interface POS0
channel-group 1
!
interface POS1
channel-group 1
```

Example 10-4 ML_Series B Configuration

```
bridge irb
bridge 1 protocol ieee
!
!
interface Port-channel1
bridge-group 1
!
interface FastEthernet0
bridge-group 1
!
interface POS0
channel-group 1
!
interface POS1
 no ip address
 channel-group 1
```

Understanding Encapsulation over FEC or POS Channel

When configuring encapsulation over FEC or POS, be sure to configure IEEE 802.1Q on the port-channel interface, not its member ports. However, certain attributes of port channel, such as duplex mode, need to be configured at the member port levels. Also make sure that you do not apply protocol-level configuration (such as an IP address or a bridge group assignment) to the member interfaces. All protocol-level configuration should be on the port channel or on its subinterface. You must configure IEEE 802.1Q encapsulation on the partner system of the EtherChannel as well.

Configuring Encapsulation over EtherChannel or POS Channel

To configure encapsulation over the FEC or POS channel, perform the following procedure, beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# interface port-channel` *channel-number.subinterface-number*	Configures the subinterface on the created port channel.
Step 2	`Router(config-subif)# encapsulation dot1q` *vlan-id*	Assigns the IEEE 802.1Q encapsulation to the subinterface.
Step 3	`Router(config-subif)# bridge-group` *bridge-group-number*	Assigns the subinterface to a bridge group.

	Command	Purpose
Step 4	`Router(config-subif)# end`	Exits to privileged EXEC mode. **Note** Optionally, you can remain in interface configuration mode and enable other supported interface commands to meet your requirements.
Step 5	`Router# copy running-config startup-config`	(Optional) Saves the configuration changes to NVRAM.

Encapsulation over EtherChannel Example

Figure 10-3 shows an example of encapsulation over EtherChannel. The associated code for ML_Series A is provided in Example 10-5 and for ML_Series B in Example 10-6.

Figure 10-3 Encapsulation over EtherChannel Example

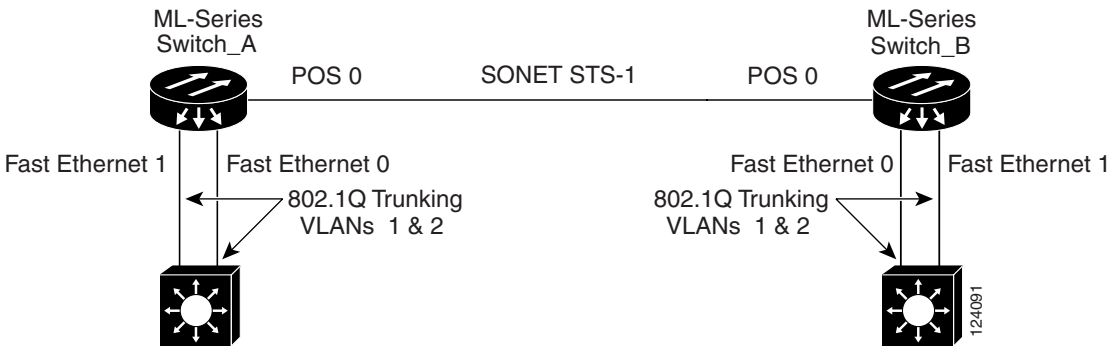

This encapsulation over EtherChannel example shows how to set up two ONS 15310-CL nodes or ONS 15310-MA nodes with ML-Series cards to interoperate with two switches that also support IEEE 802.1Q encapsulation over EtherChannel. To set up this example, use the configurations in the following sections for both Switch A and Switch B.

Example 10-5 ML_Series A Configuration

```
hostname ML_Series_A
!
bridge irb
bridge 1 protocol ieee
bridge 2 protocol ieee
!
interface Port-channel1
hold-queue 150 in
!
interface Port-channel1.1
 encapsulation dot1Q 1 native
bridge-group 1
!
interface Port-channel1.2
 encapsulation dot1Q 2
bridge-group 2

!
```

```
interface FastEthernet0
channel-group 1
!
interface FastEthernet1
channel-group 1
!
interface POS0
!
interface POS0.1
 encapsulation dot1Q 1 native
 bridge-group 1
!
interface POS0.2
 encapsulation dot1Q 2
 bridge-group 2
```

Example 10-6 ML_Series B Configuration

```
hostname ML_Series_B
!
bridge irb
bridge 1 protocol ieee
bridge 2 protocol ieee
!
interface Port-channel1
hold-queue 150 in
!
interface Port-channel1.1
 encapsulation dot1Q 1 native
bridge-group 1
!
interface Port-channel1.2
 encapsulation dot1Q 2
bridge-group 2
!
interface FastEthernet0
channel-group 1
!
interface FastEthernet1
channel-group 1
!
interface POS0
!
interface POS0.1
 encapsulation dot1Q 1 native
 bridge-group 1
!
interface POS0.2
 encapsulation dot1Q 2
 bridge-group 2
!
```

Monitoring and Verifying EtherChannel and POS

After FEC or POS is configured, you can monitor its status using the **show interfaces port-channel** command.

Example 10-7 show interfaces port-channel Command

```
ML_Series# show int port-channel 9
Port-channel9 is down, line protocol is down
  Hardware is FEChannel, address is 0000.0000.0000 (bia 0000.0000.0000)
  Internet address is 192.26.24.22/25
  MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 04:00:00
    No. of active members in this channel: 0
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/300/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes
     Received 0 broadcasts (0 IP multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 0 multicast
     0 input packets with dribble condition detected
     0 packets output, 0 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

Understanding Link Aggregation Control Protocol

In Software Release 8.0.0 and later, ML100T-12, ML1000-2, ML100T-8, and CE-100T-8 cards can utilize the link aggregation control protocol (LACP) to govern reciprocal peer packet transmission with respect to LACP's detection of flawed packets. The cards' ports transport a signal transparently (that is, without intervention or termination). However, this transparent packet handling is done only if the LACP is not configured for the ML series card.

Passive Mode and Active Mode

Passive or active modes are configured for a port and they differ in how they direct a card to transmit packets: In passive mode, the LACP resident on the node transmits packets only after it receives reciprocal valid packets from the peer node. In active mode, a node transmits packets irrespective of the LACP capability of its peer.

LACP Functions

LACP performs the following functions in the system:

- Maintains configuration information in order to control aggregation
- Exchanges configuration information with other peer devices

- Attaches or detaches ports from the link aggregation group based on the exchanged configuration information
- Enables data flow when both sides of the aggregation group are synchronized

In addition, LACP provides the following benefits:

- Logical aggregation of bandwidth
- Load balancing
- Fault tolerance

LACP Parameters

LACP utilizes the following parameters to control aggregation:

System Identifier—A unique identification assigned to each system. It is the concatenation of the system priority and a globally administered individual MAC address.

Port Identification—A unique identifier for each physical port in the system. It is the concatenation of the port priority and the port number.

Port Capability Identification—An integer, called a key, that identifies one port's capability to aggregate with another port. There are two types of key: administrative and operational. An administrative key is configured by the network administrator, and an operational key is assigned by LACP to a port based on its aggregation capability.

Aggregation Identifier—A unique integer that is assigned to each aggregator and is used for identification within the system.

LACP Usage Scenarios

In Software Release 8.0.0 and later, LACP functions on ML-Series cards in termination mode and on the CE-Series cards in transparent mode.

Termination Mode

In termination mode, the link aggregation bundle terminates or originates at the ML card. To operate in this mode, LACP should be configured on the Ethernet interface. One protect SONET or SDH circuit can carry the aggregated Ethernet traffic of the bundle. The advantage of termination mode over transparent mode is that the network bandwidth is not wasted. However. the disadvantage is that there is no card protection between the CPE and UNI (ONS 15454) because all the links in the ML card bundle belong to the same card.

Figure 10-4 LACP Termination Mode Example

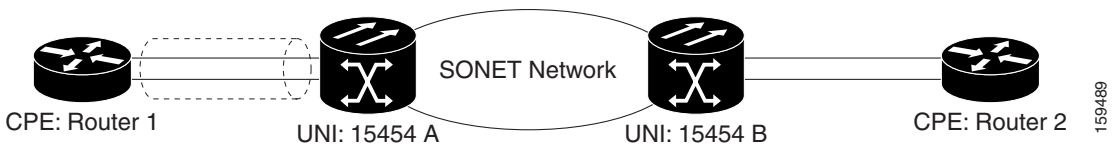

Transparent Mode

In Figure 10-5 the link aggregation bundle originates at router 1 and terminates at router 2. Transparent mode is enabled when the LACP packets are transmitted without any processing on a card. While functioning in this mode, the CE-100T-8 cards pass through LACP packets transparently so that the two CPE devices perform the link aggregation. To operate in this mode, no LACP configuration is required on the CE-100T-8 cards.

Figure 10-5 LACP Transparent Mode Example

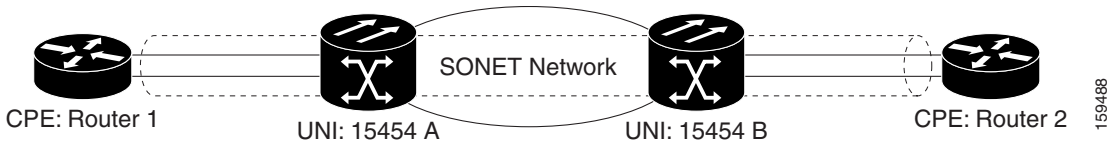

Configuring LACP

To configure LACP over the EtherChannel or POS channel, perform the following procedure, beginning in global configuration mode:

	Command	Purpose
Step 1	`Router(config)# int port( <interface-number>`	Accesses the port interface where you will create the LACP.
Step 2	`Router(config-if)# int fa( <facility-number>`	Access the facility number on the port.
Step 3	`Router(config-if)# channel`	Accesses the channel group of commands.
Step 4	`Router(config-if)# channel-group <channel-number> mode ?`	Queries the current mode of the channel group. Options include active and passive.
Step 5	`Router(config-if)# channel-group <channel-number> mode active`	Places the channel group in active mode.
Step 6	`Router(config-if)# exit`	Exits the channel group configuration.
Step 7	`Router(config-if)# int fa <facility-number>`	Accesses the facility.
Step 8	`Router(config-if)# lacp-port`	Access the link aggregation control protocol commands for the port.
Step 9	`Router(config-if)# lacp port-priority <priority number>`	Sets the LACP port's priority. Range of values is from 1 through 65535. For example, `lacp port-priority 100`
Step 10	`Router(config-if)# exit`	Exits the port's configuration mode.
Step 11	`Router(config)#lacp sys`	Accesses the system LACP settings.

	Command	Purpose
Step 12	`Router(config)#`**`lacp system-priority`** *<system priority>*	Sets the LACP system priority in a range of values from 1 through 65535. For example, `lacp system-priority 100`
Step 13	`Router(config)#` **`exit`**	Exits the global configuration mode.
Step 14	`Router#` **`copy running-config startup-config`**	(Optional) Saves the configuration changes to NVRAM.

In , Example 10-8 the topology includes two nodes with a GEC or FEC transport between them. This example shows one GEC interface on Node 1. (Up to four similar types of links per bundle are supported.)

Example 10-8 LACP Configuration Example

```
ML2-Node1#sh run int gi0
Building configuration...

Current configuration : 150 bytes
!
interface GigabitEthernet0
 no ip address
 no keepalive
 duplex auto
 speed auto
 negotiation auto
 channel-group 1 mode active
 no cdp enable
end

ML2-Node1#
ML2-Node1#sh run int por1
Building configuration...

Current configuration : 144 bytes
!
interface Port-channel1
 no ip address
 no negotiation auto
 service instance 30 ethernet
  encapsulation dot1q 301
  bridge-domain 30
 !
end

ML2-Node1#
ML2-Node1#sh lacp int
Flags:  S - Device is requesting Slow LACPDUs
        F - Device is requesting Fast LACPDUs
        A - Device is in Active mode       P - Device is in Passive mode

Channel group 1
                            LACP port     Admin     Oper    Port     Port
Port      Flags   State     Priority      Key       Key     Number   State
Gi0       SA      bndl      32768         0x1       0x1     0x5      0x3D
ML2-Node1#
Configuration  remains same for the ML2-Node2 also.
```

Load Balancing on the ML-Series cards

The load balancing for the Ethernet traffic on the portchannel is performed while sending the frame through a port channel interface based on the source MAC and destination MAC address of the Ethernet frame.

On a 2 port channel interface, the Unicast Ethernet traffic (Learned MAC with unicast SA and DA) is transmitted on either first or second member of the port-channel based on the result of the "Exclusive OR" (XOR) operation applied on the second least significant bits (bit 1) of DA-MAC and SA-MAC. So, if the "XOR" result of the Ethernet frames SA-MAC second least significant bit and DA-MAC second least significant bit is 0 then the frame is sent on the first member and if the result is 1 then the frame is transmitted on the second member port of the port channel.

Table 10-1 MAC Based - 2- Port Channel Interface

Second Least Significant bit of the MAC-DA	Second Least Significant bit of the MAC-SA	XOR Result	Used Member Interface for the Frame Forwarding to the EtherChannel and/or Port Channel
0	0	1	Port 1
0	1	1	Port 2
1	0	1	Port 2
1	1	0	Port 1

Table 10-2 IP Based - 2-Port Channel Interface

Second Least Significant bit of the IP-DA	Second Least Significant bit of the IP-SA	XOR Result	Used Member Interface for the Frame Forwarding to the EtherChannel and/or Port Channel
0	0	1	Port 1
0	1	1	Port 2
1	0	1	Port 2
1	1	0	Port 1

The Flood Ethernet traffic (Unknown MAC, Multicast and Broadcast frames) is transmitted on the first active member of the port-channel.

The routed IP Unicast traffic from the ML-Series towards the port channel ports is transmitted on either interface based on the result of the "Exclusive OR" (XOR) operation applied on the second least significant bits of the source and destination IP address of the IP packet. So if the "XOR" result of the IP packets Source Address least significant bit and Destination Address least significant bit is 0 then the frame is on the first member port and if the result is 1 then the frame is transmitted on the second member port.

On the 4 port EtherChannel the second and third least significant bits are used for load balancing.

Table 10-3 MAC Based - 4-Port Channel Interface

Third Least Significant bit of the MAC-DA	Third Least Significant bit of the MAC-SA	Second Least Significant bit of the MAC-DA	Second Least Significant bit of the MAC-SA	XOR Result	Used Member Interface for the Frame Forwarding to the EtherChannel and/or Port Channel
0	0	0	0	00	First
0	0	0	1	01	Second
0	0	1	0	01	Second
0	0	1	1	00	First
0	1	0	0	10	Third
0	1	0	1	11	Fourth
0	1	1	0	11	Fourth
0	1	1	1	10	Second
1	0	0	0	10	Second
1	0	0	1	11	Third
1	0	1	0	11	Third
1	0	1	1	10	Second
1	1	0	0	00	First
1	1	0	1	01	Second
1	1	1	0	01	Second
1	1	1	1	00	First

Table 10-4 IP Based - 4-Port Channel Interface

Third Least Significant bit of the IP-DA	Third Least Significant bit of the IP-SA	Second Least Significant bit of the IP-DA	Second Least Significant bit of the IP-SA	XOR Result	Used Member Interface for the Frame Forwarding to the EtherChannel and/or Port Channel
0	0	0	0	00	First
0	0	0	1	01	Second
0	0	1	0	01	Second

Table 10-4 IP Based - 4-Port Channel Interface

Third Least Significant bit of the IP-DA	Third Least Significant bit of the IP-SA	Second Least Significant bit of the IP-DA	Second Least Significant bit of the IP-SA	XOR Result	Used Member Interface for the Frame Forwarding to the EtherChannel and/or Port Channel
0	0	1	1	00	First
0	1	0	0	10	Third
0	1	0	1	11	Fourth
0	1	1	0	11	Fourth
0	1	1	1	10	Second
1	0	0	0	10	Second
1	0	0	1	11	Third
1	0	1	0	11	Third
1	0	1	1	10	Second
1	1	0	0	00	First
1	1	0	1	01	Second
1	1	1	0	01	Second
1	1	1	1	00	First

The routed IP Multicast traffic from the ML-Series towards the RPR ring is transmitted on the first active member of the port channel.