



## Monitoring and Troubleshooting

This chapter provides an introduction to monitoring and troubleshooting the Cisco Ethernet switches in the Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.0 and Release 1.1, and presents the following major topics:

- [Using Basic CLI Commands to Troubleshoot the Switch, page 5-1](#)
- [Switch in Dhub: Ethernet Troubleshooting Examples, page 5-8](#)
- [More Troubleshooting Tips, page 5-11](#)

### Using Basic CLI Commands to Troubleshoot the Switch

The following commands are addressed:

- [show interface](#) (and variations: gigabitethernet, vlan, port-channel, and tunnel)
- [logging event link-status](#)

#### show interface

The command **show interface**, without a specific interface as an option, retrieves information from every active interface. Below are the counters that are relevant to troubleshooting:

- *input errors*—This is a count of any errors that occurred while the switch is trying to receive packets from the referenced port.

The counter includes both **CRC** (cyclic redundancy check) and **frame errors**, but does not include ignored packets. CRC errors occur when the received packets fail the CRC. Frame errors occur when the receiving frame is not complete. The **ignored** counter counts the number of frames dropped on input because of resource exhaustion in the switch fabric. **Overruns** occur when inter frame gaps (IFGs) are so short that a new Ethernet frame arrives before the previous frame has been completely stored in shared memory.

- *output errors*—This is a count of any errors that occurred while the switch is trying to transmit packets from the referenced port.

**Collisions** shows the number of times a collision occurred while the switch is trying to transmit a packet from the referenced port. This counter should be 0 for a port operating in full-duplex mode. The **interface resets** counter counts the number of times the port reset itself, generally the result of link up or link down transitions. **Underruns** occur when packets are not retrieved quickly enough from shared memory to be transmitted.

- *babbles* and *late collision*—A **babble** is an error caused by the transmission of frames in excess of 1518 bytes in size.

A **late collision** is a collision that occurs outside of the collision window, which is typically caused by a duplex mismatch or a wire length that exceeds the distance limitations (100 meters for 10/100BASE-T ports). The **deferred** counter tabulates the number of times the port had to wait to transmit as a result of traffic on the wire.

- *lost carrier* and *no carrier*—The carrier is an electrical signal that Ethernet devices use to detect whether the wire is currently being used by another transmitting station.

The **lost carrier** counter increases each time a carrier sense loss occurs. This happens when the hardware is transmitting a frame onto the wire and does not see its own carrier wave on the Ethernet. The absence of the carrier signal increments the **no carrier** counter.

## Sample Output for show interfaces on a Headend Switch

The following basic CLI **show interface** commands can provide a great deal of information of assistance in troubleshooting Cisco Catalyst 4500 series switches. The examples below show nominal values for **show interfaces** on a headend switch for a single GE port, a VLAN, and an EtherChannel, respectively:

- [show interfaces gigabitethernet](#)
- [show interfaces vlan](#)
- [show interfaces port-channel](#)



### Note

The VLAN, VLAN50, is in Port-channel2, which contains two GE interfaces—3/5 and 3/6 in this example.

### show interfaces gigabitethernet

```
Headend#show interfaces gigabitethernet 3/5
GigabitEthernet3/5 is up, line protocol is up (connected)
  Hardware is Gigabit Ethernet Port, address is 000c.850e.80bf (bia 000c.850e.80bf)
  Description: Unidirectional Link of EtherChannel 2
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 41/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, link type is force-up, media type is SX
  output flow-control is off, input flow-control is on
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  30 second input rate 0 bits/sec, 0 packets/sec
  30 second output rate 162032000 bits/sec, 14826 packets/sec
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts (0 multicast)
  0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
  0 input packets with dribble condition detected
  835379461 packets output, 1137771157239 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

Headend#show interfaces gigabitethernet 3/6
GigabitEthernet3/6 is up, line protocol is up (connected)
  Hardware is Gigabit Ethernet Port, address is 000c.850e.80bf (bia 000c.850e.80bf)
  Description: Unidirectional Link of EtherChannel 2
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 40/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, link type is force-up, media type is SX
  output flow-control is off, input flow-control is on
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  30 second input rate 0 bits/sec, 0 packets/sec
  30 second output rate 158527000 bits/sec, 14498 packets/sec
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
    0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts (0 multicast)
    0 runs, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
    0 input packets with dribble condition detected
  835489311 packets output, 1137915654913 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

```

### show interfaces vlan

```

Headend#show interfaces vlan50
Vlan50 is up, line protocol is up
  Hardware is Ethernet SVI, address is 000c.850e.80bf (bia 000c.850e.80bf)
  Description: VoD Servers
  Internet address is 192.168.50.2/24
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 155/255
  Encapsulation ARPA, loopback not set
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:06:04, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  30 second input rate 610762000 bits/sec, 56805 packets/sec
  30 second output rate 0 bits/sec, 0 packets/sec
  L3 in Switched: ucast: 3340981987 pkt, 4490274590976 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 4012 pkt, 192576 bytes - mcast: 0 pkt, 0 bytes
    3340981999 packets input, 4490274591912 bytes, 0 no buffer
  Received 12 broadcasts (0 IP multicast)
    0 runs, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
    24215 packets output, 1163224 bytes, 0 underruns
    0 output errors, 0 interface resets
    0 output buffer failures, 0 output buffers swapped out

```

## show interfaces port-channel

```

Headend#show interfaces port-channel 2
Port-channel2 is up, line protocol is up (connected)
  Hardware is EtherChannel, address is 000c.850e.80bf (bia 000c.850e.80bf)
  Description: Unidirectional EtherChannel to DHub_B
  Internet address is 192.168.169.5/30
  MTU 1500 bytes, BW 2000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 40/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Unknown duplex, Unknown Speed, media type is unknown media type
  output flow-control is unsupported, input flow-control is unsupported
  Members in this channel: Gi3/5 Gi3/6
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  30 second input rate 0 bits/sec, 0 packets/sec
  30 second output rate 314028000 bits/sec, 28505 packets/sec
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 1668713589 pkt, 2242751063616 bytes - mcast: 0 pkt, 0 bytes
    0 packets input, 0 bytes, 0 no buffer
  Received 0 broadcasts (0 IP multicast)
    0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
    0 input packets with dribble condition detected
  1668774184 packets output, 2272833746238 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

```

## show interfaces tunnel

This command is illustrated for both headend and Dhub switches in [Troubleshooting Tunnels, page 5-4](#).

## logging event link-status

This command is useful in providing the up/down status of links, sending messages to the console when the status of a link changes. This command can be used in conjunction with management applications such as Cisco Info Center (CIC) (which can pick up console messages and perform a notification process), as there is less latency with this command than there is with MIBs and traps. However, take into account that **logging event link-status** should not be used unless it is necessary, as logging can burden the CPU, especially if traps are also being used.

## Troubleshooting Tunnels

This section covers both working and nonworking unidirectional EtherChannels:

- [Working Unidirectional EtherChannel Between Headend and DHub](#)
- [Nonworking Unidirectional EtherChannel Between Headend and DHub](#)

**Note**Lines of interest are highlighted in **bold**.

## Working Unidirectional EtherChannel Between Headend and DHub

For reference, see [Figure 3-3 on page 3-15](#). The packets on Tunnel2 flow from DHub\_B to Headend. Use the command **show interfaces tunnel2** on both switches to verify that packets are sent by the DHub\_B switch through Tunnel2 and received by the Headend switch. Outputs from the following commands are illustrated:

- [show interfaces tunnel](#)
- [debug tunnel](#)
- [show ip ospf neighbor](#)
- [show arp](#)

### show interfaces tunnel

#### At the Headend

```
Headend#show interfaces Tunnel2
Tunnel2 is up, line protocol is down
  Hardware is Tunnel
  Description: Return Path for Unidirectional EtherChannel to DHub_B
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 1.1.1.5, destination 2.2.2.6
  Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
  Checksumming of packets disabled, fast tunneling enabled
  Tunnel is receive-only UDLR tunnel for send-only interface Port-channel2
  Last input 00:00:01, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (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
    2 packets input, 212 bytes, 0 no buffer
    Received 0 broadcasts (0 IP multicast)
    0 runs, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    0 packets output, 0 bytes, 0 underruns
    0 output errors, 0 collisions, 0 interface resets
    0 output buffer failures, 0 output buffers swapped out
```

#### At the Dhub

```
DHub_B#show interfaces Tunnel2
Tunnel2 is up, line protocol is up
  Hardware is Tunnel
  Description: Return Path for Unidirectional EtherChannel from Headend
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 2.2.2.6, destination 1.1.1.5
  Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
```

```

Checksumming of packets disabled, fast tunneling enabled
Tunnel is send-only UDLR tunnel for receive-only interface Port-channel2
Last input never, output 00:00:00, output hang never
Last clearing of "show interface" counters 00:00:17
Input queue: 0/75/0/0 (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, 0 no buffer
  Received 0 broadcasts (0 IP multicast)
  0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
  2 packets output, 184 bytes, 0 underruns
  0 output errors, 0 collisions, 0 interface resets
  0 output buffer failures, 0 output buffers swapped out

```

## debug tunnel

Use the command **debug tunnel** on both switches to see packets travel through the tunnel.

### At the Headend

```

Headend#debug tunnel
Tunnel Interface debugging is on
Headend#
3d18h: Tunnel2: GRE/IP to classify 2.2.2.6->1.1.1.5 (len=92 ttl=253 tos=0x0) ok,
retval=0x0
3d18h: Tunnel2: GRE/IP to decaps 2.2.2.6->1.1.1.5 (len=92 ttl=252)

```

### At the Dhub

```

DHub_B#debug tunnel
Tunnel Interface debugging is on
DHub_B#
2w5d: Tunnel2: GRE/IP encapsulated 2.2.2.6->1.1.1.5 (linktype=7, len=92)

```

## show ip ospf neighbor

If the unidirectional EtherChannel and tunnel are working properly, DHub\_B will be recognized as Headend's OSPF neighbor.

### At the Headend

```

Headend#show ip ospf neighbor Port-channel2

```

Neighbor ID	Pri	State	Dead Time	Address	Interface
2.2.2.6	1	FULL/DR	00:00:30	192.168.169.6	Port-channel2

## show arp

There will also be an ARP entry for DHub\_B's side of Port-channel2 in Headend's ARP table.

### At the Headend

```

Headend#show arp | include Port-channel2

```

Protocol	Address	Age (min)	Hardware Addr	Type	Interface
Internet	192.168.169.5	-	000c.850e.80bf	ARPA	Port-channel2
Internet	<b>192.168.169.6</b>	13	0009.e89a.2dff	ARPA	<b>Port-channel2</b>

## Nonworking Unidirectional EtherChannel Between Headend and DHub

For reference, see [Figure 3-3 on page 3-15](#). When Tunnel2 is down because of a physical connection problem, the line protocol on the DHub\_B side will be “down” and the number of output drops should increment. Outputs from the following commands are illustrated:

- [show interfaces tunnel](#)
- [debug tunnel](#)
- [show ip ospf](#)
- [show arp](#)

### show interfaces tunnel

#### At the Dhub

```
DHub_B#show interfaces Tunnel2
Tunnel2 is up, line protocol is down
  Hardware is Tunnel
  Description: Return Path for Unidirectional EtherChannel from Headend
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 2.2.2.6, destination 1.1.1.5
  Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
  Checksumming of packets disabled, fast tunneling enabled
  Tunnel is send-only UDLR tunnel for receive-only interface Port-channel2
  Last input never, output 00:01:12, output hang never
  Last clearing of "show interface" counters 00:14:49
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 30
  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, 0 no buffer
    Received 0 broadcasts (0 IP multicast)
    0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    135 packets output, 13584 bytes, 0 underruns
    0 output errors, 0 collisions, 0 interface resets
    0 output buffer failures, 0 output buffers swapped out
```

### debug tunnel

If you use **debug tunnel** again, you will see that no packets are being sent through the tunnel.

#### At the Headend

```
Headend#debug tunnel
Tunnel Interface debugging is on
Headend#
```



Note

---

There is no output for **debug tunnel**.

---

**At the Dhub**

```
DHub_B#debug tunnel
Tunnel Interface debugging is on
DHub_B#
```

**Note**


---

There is no output for **debug tunnel**.

---

**show ip ospf**

If you look at Headend's OSPF neighbors, you will no longer see DHub\_B listed as a neighbor.

**At the Headend**

```
Headend#show ip ospf neighbor Port-channel2
```

**Note**


---

DHub\_B is no longer listed.

---

**show arp**

The ARP entry for DHub\_B's side of Port-channel2 might still be in Headend's ARP table, but it will be aging. If you execute a **clear arp**, the ARP entry will be removed from the ARP table and will not be added until the tunnel is reestablished.

**At the Headend**

```
Headend#show arp | include Port-channel2
Protocol Address Age (min) Hardware Addr Type Interface
Internet 192.168.169.5 - 000c.850e.80bf ARPA Port-channel2
Internet 192.168.169.6 21 0009.e89a.2dff ARPA Port-channel2
```

```
Headend#clear arp
```

```
Headend#show arp | include Port-channel2
Protocol Address Age (min) Hardware Addr Type Interface
Internet 192.168.169.5 - 000c.850e.80bf ARPA Port-channel2
```

## Troubleshooting Cisco Catalyst 4500 Series Switches

For additional tips on troubleshooting Cisco Catalyst 4500 series switches, see “Troubleshooting Tips” at the following URL:

[http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/trbl\\_ja.htm](http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/trbl_ja.htm)

## Switch in Dhub: Ethernet Troubleshooting Examples

The following example illustrates a configuration in which links are broken between a headend and a Dhub switch connected by two asymmetric EtherChannels. This scenario only supports 2 Gbps of video traffic, as Port-channel2 is used for link redundancy.

**Figure 5-1** illustrates the basic topology. The response of the switches to the various failure modes provides useful information about anticipated behavior, as well as clues to troubleshooting should such failures occur.

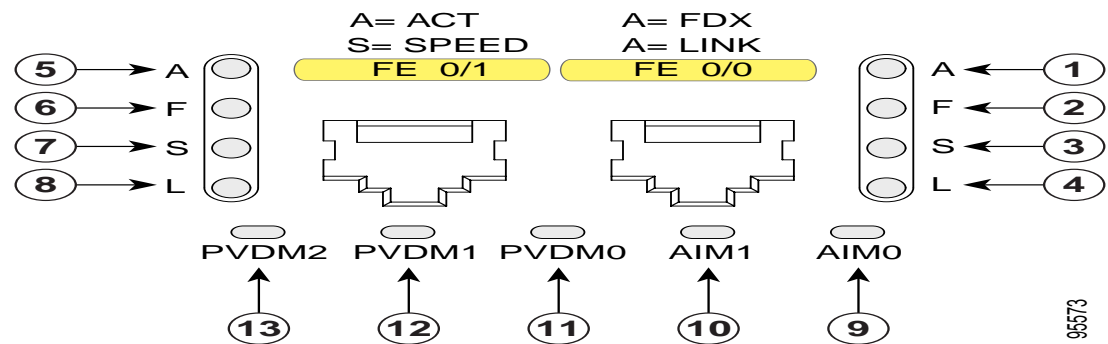




Note

For a discussion of link failure modes, see [Chapter 4, “Providing Redundancy and Reliability.”](#)

Figure 5-1 Basic Ethernet Troubleshooting Topology



## Preconditions

- There are two EtherChannels: Port-channel1 (P1) and Port-channel2 (P2).
- OSPF is enabled, with equal-cost OSPF routes for P1 and P2.
- Load balancing is applied to the destination ports for both EtherChannels with CEF (Cisco Express Forwarding). This causes video flows to be distributed evenly across all four GE ports.

## IP Routes

Route tables for the headend and Dhub switches are presented below.

### Headend Route Table

```
C 192.168.151.0/24 is directly connected, Vlan151
O 192.168.150.0/24 [110/2] via 192.168.170.2, 00:00:32, Port-channel2
    [110/2] via 192.168.169.2, 00:00:32, Port-channel1
192.168.170.0/30 is subnetted, 1 subnets
C    192.168.170.0 is directly connected, Port-channel2
192.168.169.0/30 is subnetted, 1 subnets
C    192.168.169.0 is directly connected, Port-channel1
```

### Dhub RouteTable

```
O 192.168.151.0/24 [110/2] via 192.168.170.1, 00:01:44, Port-channel2
    [110/2] via 192.168.169.1, 00:01:44, Port-channel1
C 192.168.150.0/24 is directly connected, Vlan150
192.168.170.0/30 is subnetted, 1 subnets
C    192.168.170.0 is directly connected, Port-channel2
192.168.169.0/30 is subnetted, 1 subnets
C    192.168.169.0 is directly connected, Port-channel1
```

## Actions and Responses

- Pull the fiber on the send-only GE of P1 (GE 2/2 on the Dhub and GE 3/1 on the headend).
  - Sessions on that port are lost and never recover. EtherChannel load balancing does not rebalance on a loss of a send-only link.
  - GE 2/2 shows down/down on the Dhub:
 

```
2d21h: %EC-5-UNBUNDLE: Interface GigabitEthernet2/2 left the port-channel
Port-channell
```
  - P1 on the Dhub now only shows GE 2/1 as a member.
  - No event is shown on the headend. The disabled send-only link is still reported as up/up.
  - No alarms are noted on the QAM device, although the missing sessions are removed from the output extraction page.
- Replace the fiber on the send-only GE of P1.
  - Sessions recover within 2 seconds.
- Pull the fiber on the bidirectional GE of P1 (GE 2/1 on the Dhub and GE 3/1 on the headend).
  - Video is momentarily rebalanced to the unidirectional link (evidenced by “blinking”), but after a few seconds the entire EtherChannel goes down as OSPF rerouting takes over.
  - No sessions are lost.
  - On the Dhub, GE 2/1 is down/down; P1 is up/up and shows GE 2/2 as member.
  - On the headend, GE 3/1 down/down; P1 is up/up and shows G3/2 as member
  - The Dhub log reads as follows:
 

```
2d21h: %EC-5-UNBUNDLE: Interface GigabitEthernet2/1 left the port-channel
Port-channell
2d21h: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet2/1, changed
state to down
2d21h: %LINK-3-UPDOWN: Interface GigabitEthernet2/1, changed state to down
```
  - The headend log reads as follows:
 

```
Jun 13 05:19:52.174: %EC-5-UNBUNDLE: Interface GigabitEthernet3/1 left the
port-channel Port-channell
*Jun 13 05:19:53.174: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet3/1, changed state to down
*Jun 13 05:19:54.174: %LINK-3-UPDOWN: Interface GigabitEthernet3/1, changed state
to down
*Jun 13 05:19:54.302: %OSPF-5-ADJCHG: Process 100, Nbr 192.168.170.2 on
Port-channell from FULL to DOWN, Neighbor Down: Dead timer d
*Jun
```
  - P1 is no longer in the **sh ip route** for 192.168.150.0.
  - No alarms are on the QAM device.
  - All traffic now flows through P2.



**Note** A **sh int g2/2** on the Dhub shows a large output rate for this receive-only link. This is erroneous, as there is no output, as shown below:

```
30 second output rate 4084883554000 bits/sec, 1651326426 packets/sec
```

## 4. Pull the P1 DWDM fiber.

- Recovery time is 7 seconds
- P1 is down/down in the Dhub, but is up/up in the headend.
- The IP route now has P2 as the only link for 192.168.150.0.
- The Dhub log reads as follows:

```
00:09:45: %EC-5-UNBUNDLE: Interface GigabitEthernet2/1 left the port-channel
Port-channell
00:09:45: %EC-5-UNBUNDLE: Interface GigabitEthernet2/2 left the port-channel
Port-channell
00:09:46: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet2/1,
changed state to down
00:09:46: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channell, changed
state to down
00:09:47: %OSPF-5-ADJCHG: Process 100, Nbr 192.168.170.1 on Port-channell from
FULL to DOWN, Neighbor Down: Dead timer expired
00:09:47: %LINK-3-UPDOWN: Interface GigabitEthernet2/1, changed state to down
00:09:47: %LINK-3-UPDOWN: Interface Port-channell, changed state to down
```

- The headend log reads as follows:

```
*Jun 13 05:51:29.331: %EC-5-UNBUNDLE: Interface GigabitEthernet3/1 left the
port-channel Port-channell
*Jun 13 05:51:30.335: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet3/1, changed state to down
*Jun 13 05:51:31.331: %LINK-3-UPDOWN: Interface GigabitEthernet3/1, changed state
to down
*Jun 13 05:51:31.955: %OSPF-5-ADJCHG: Process 100, Nbr 192.168.170.2 on
Port-channell from FULL to DOWN, Neighbor Down: Dead timer d
```

## 5. Pull the P2 DWDM fiber.

- The results are similar to those for pulling the P1 DWDM fiber.

## 6. Remove the receive-only GBIC or pull the associated fiber (same effect).

- All video streams on the failed port are lost.
- There is no recovery from streams lost on the failed port.
- The Dhub log reads as follows:

```
2w1d: %EC-5-UNBUNDLE: Interface GigabitEthernet2/2 left the port-channel
Port-channell
```

## More Troubleshooting Tips

The following discussions can be of assistance in troubleshooting:

- [Example Link Utilization and Related Traffic Statistics](#)
- [Understanding the Load-Balancing Behavior of Asymmetric EtherChannels](#)
- [Preventing Lost Traffic: Adding Send-Only GE Links to Operational EtherChannels](#)
- [Understanding ARP and MAC Address Table Timeouts](#)
- [Understanding Buffers on Cisco Catalyst 4500 Series Switches](#)

## Example Link Utilization and Related Traffic Statistics

As a benchmark for troubleshooting links that support video streams, it can be useful to see some nominal traffic statistics for various degrees of link utilization. As the number of video streams increases, so does link utilization—and also latency. Table 5-1 lists example link utilization and related traffic statistics for a Concurrent VoD server. Numbers for other servers may vary.

Table 5-1 Example Link Utilization and Related Traffic Statistics for a Concurrent VoD Server

Link Utilization, percent	Number of Video Streams	Arrival Time, msec			Latency, microsec		Packets	
		Minimum	Average	Maximum	Minimum	Maximum	Out of Order	Dropped
88	240	0.84	2.82	4.41	21.8	968.5	0	0
90	245	0.84	2.82	4.41	21.8	1369.7	0	0
92	250	0.84	2.82	4.41	21.8	1068.0	0	0
93	255	0.84	2.82	4.41	—	—	—	—
95	260	—	—	—	—	—	—	—

## Understanding the Load-Balancing Behavior of Asymmetric EtherChannels

The load balancing behavior of an EtherChannel with both bidirectional and send-only links is as follows:

- If a bidirectional link is lost, the flows on that link are rebalanced to the other links. If this link is the only bidirectional link in the EtherChannel, then it is possible that OSPF will lose this route and force the traffic to a new route.
- If a send-only link is lost, the flows on that link are not rebalanced to the other links.

## Preventing Lost Traffic: Adding Send-Only GE Links to Operational EtherChannels

When adding send-only links to EtherChannels that are already carrying video traffic, be sure to configure the receive-side (Dhub) switch and all physical connections before the configuring the send-side (headend) switch. If the send-side switch is configured first, then the send-only link is forced up (even though the receive-side switch is not ready to receive), and traffic will be load-balanced onto that link. This causes that traffic to be lost.

## Understanding ARP and MAC Address Table Timeouts

Because the video data flow in VoD is unidirectional, the MAC address table for each GE port connected to an edge QAM device eventually times out, causing data bound for that device to be flooded to all ports. To prevent this, set the ARP timeout for that port (or VLAN) to a value less than the value for mac-address-table aging-time. Because the default value of **aging-time** is 300 (seconds), the ARP timeout can be set to 150, as follows: **arp timeout 150**.

## Understanding Buffers on Cisco Catalyst 4500 Series Switches

The shared-memory switching buffer for Cisco Catalyst 4500 series switches running Supervisor Engines 2, 3, and 4 is 16 Mbytes. Unlike with many Cisco platforms, there is no command to change or tune the buffers on Cisco Catalyst 4500 series switches. The various transmit-queue buffer sizes are shown in [Table 5-2](#).

*Table 5-2 Transmit-Queue Buffer Sizes for Cisco Catalyst 4500 Series Switches*

Paquets per Queue	Packets per Queue on a Subport	Packets on a Nonblocking GE Port	
		Per Queue	Per Port
240	960	1920	7680

