Protection Schemes and Network Topologies

This chapter describes how protection is implemented on the Cisco ONS 15540. It also describes the supported network topologies and how protection works in these topologies. This chapter contains the following major sections:

- Protection Against Fiber and System Failures, page 2-1
- Supported Topologies, page 2-6

Protection Against Fiber and System Failures

The design of the Cisco ONS 15540 provides for two levels of network protection, facility protection and line card protection. Facility protection provides protection against failures due to fiber cuts or unacceptable signal degradation on the trunk side. Line card protection provides protection against failures both on the fiber and in the transponders, which contain the light emitting and light detecting devices as well as the 3R electronics. Line card protection can also be implemented using redundant client signals. This provides protection against the failure of the client, the transponder, or the fiber.

Splitter Based Facility Protection

To survive a fiber failure, fiber optic networks are designed with both working and protection fibers. In the event of a fiber cut or other facility failure, working traffic is switched to the protection fiber. The Cisco ONS 15540 supports such facility protection using a *splitter* scheme (see Figure 2-1) to send the output of the DWDM transmitter on two trunk side interfaces.

Figure 2-1 Splitter Protection Scheme

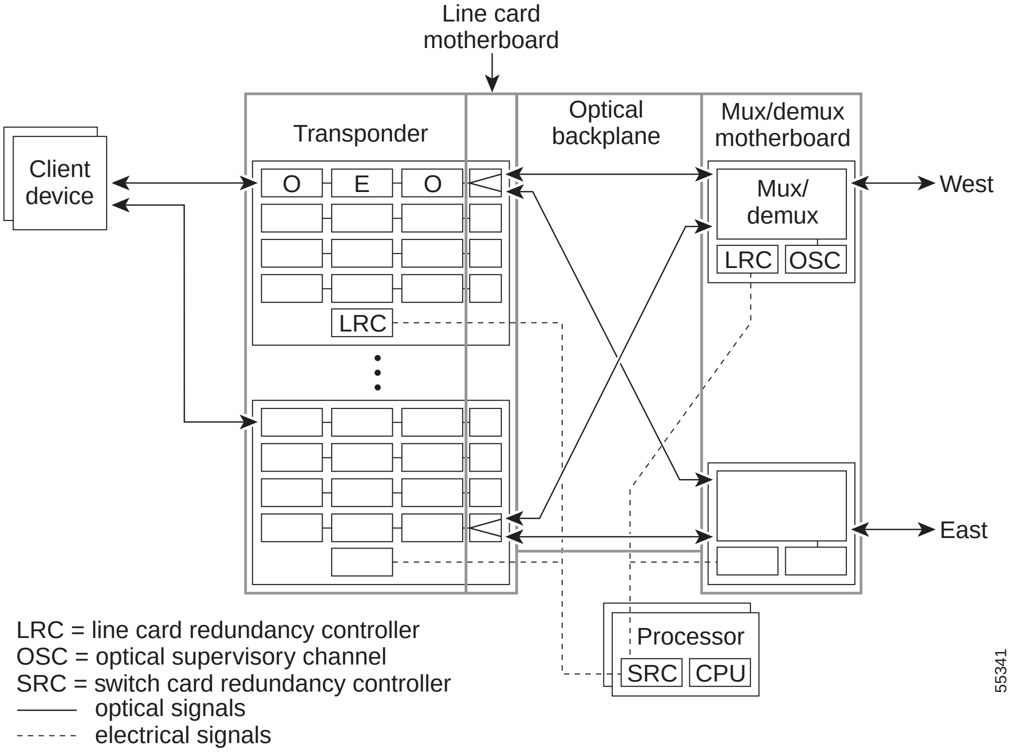

With splitter protection, splitters on the line card motherboard couples each transponder's trunk side interface across the optical backplane to the internal interfaces on the optical mux/demux modules in the east and west slots (0 and 1). On the trunk side, one fiber pair serves as the active connection, while the other pair serves as the standby. The signal is transmitted on both connections, but in the receive direction, an optical switch selects one signal to be the active one. If a failure is detected on the active signal, a switch to the standby signal is made under control of the LRC (line card redundancy controller). Assuming, for example, that the active signal in Figure 2-1 is on the east interface, a failure of the signal on that fiber would result in a switchover, and the signal on the west interface would be selected for the receive signal.

A switchover is triggered in hardware by an Loss of Light on the receive signal. Other conditions can be specified to trigger a switchover by configuring thresholds in the APS software.

Splitter Protection Considerations

The following considerations apply when using splitter protection:

- The splitter protected line card motherboard supports splitter protection. Because the signal splitter introduces 4.6 dB of loss in the transmit direction, we recommend using the nonsplitter protected line card motherboards (east or west version) for configurations where splitter protection is not required.
- Switchover after a failure under splitter protection is nonrevertive. After a switchover, manual intervention is required to revert to using the previously failed fiber for the working traffic once the fault has been remedied.
- The OSC plays a crucial role in splitter based protection by allowing the protection fiber to be monitored for a cut or other interruption of service.

For rules on how to configure the shelf for splitter protection, see Chapter 3, "Shelf Configuration Rules." For instructions on configuring the APS software for splitter protection, refer to the *Cisco ONS 15540 ESP Configuration Guide and Command Reference*.

Optical Backplane

The mapping between the internal interfaces on the line card motherboards and on the optical mux/demux modules is shown in Table 2-1. When the splitter protected line card motherboards are used, these cross connections, which are fixed and nonconfigurable, couple the signal from each transponder module to a specific position on the optical mux/demux modules in both slot 0 (west) and slot 1 (east). If west line card motherboards are used, the transponder modules are connected only to the optical mux/demux modules in slot 0; if the east line card motherboards are used, the transponder modules are connected only to the optical mux/demux modules in slot 1.

Table 2-1 Transponder to Optical Mux/Demux Module Mapping

	Optical Mux/Demux Modules Slot/Subslot		
Transponder Module Slot/Subslot	**Splitter**	**West**	**East**
2/0 through 2/3	0/0 and 1/0	0/0	1/0
3/0 through 3/3	0/0 and 1/0	0/0	1/0
4/0 through 4/3	0/1 and 1/1	0/1	1/1
5/0 through 5/3	0/1 and 1/1	0/1	1/1
8/0 through 8/3	0/2 and 1/2	0/2	1/2
9/0 through 9/3	0/2 and 1/2	0/2	1/2
10/0 through 10/3	0/3 and 1/3	0/3	1/3
11/0 through 11/3	0/3 and 1/3	0/3	1/3

Within each 8-channel range, the optical backplane maps the lower four channels (bands A, C, E, and G) to backplane connectors 1–4 on the optical mux/demux modules. The higher four channels (bands B, D, F, and H) are mapped to backplane connectors 5–8.

Figure 2-2 shows this correspondence using slots 4 and 5 as the example slot pair. The transponder modules in slot 4 are mapped to backplane connectors 1–4 for mux/demux subslot 1; the transponder modules in slot 5 are mapped to backplane connectors 5–8 of the same mux/demux subslot. The backplane connections are to either slot 0 or slot 1 (if using west or east line card motherboards) or to both slot 0 and slot 1 (if using splitter protected line card motherboards).

Figure 2-2 Transponder to Optical Mux/Demux Module Mapping

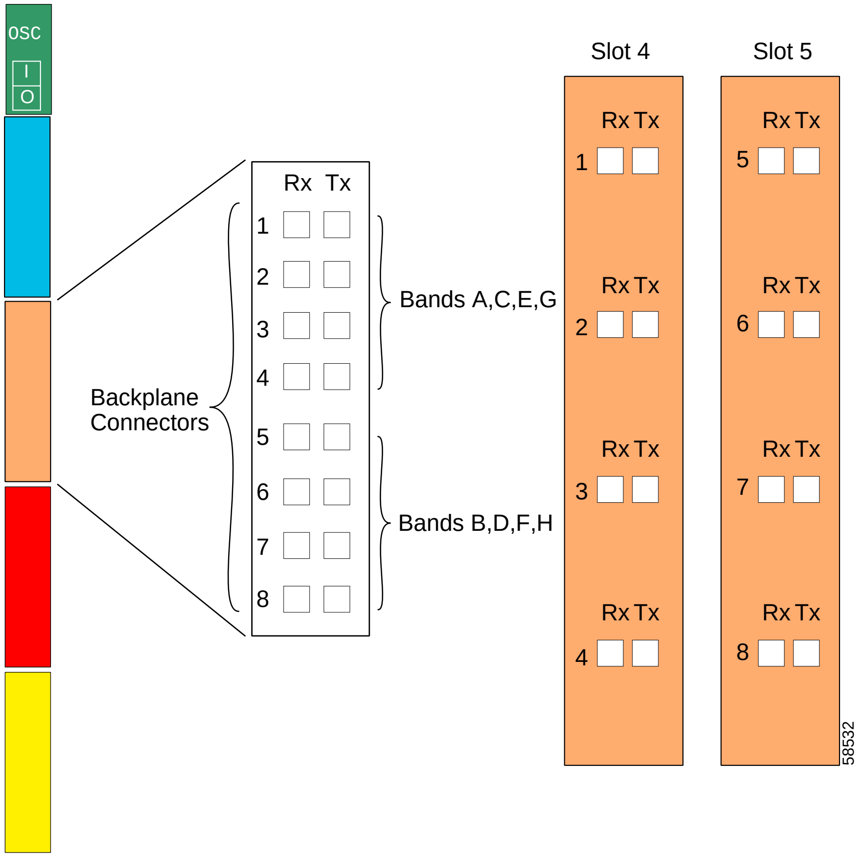

Y-Cable Based Line Card Protection

Line card protection is implemented on the Cisco ONS 15540 using a *y-cable* scheme. Y-cable protection protects against both facility failures and failure of the transponder module. Using an external 2:1 combiner (the y-cable), connections between the client equipment and the transponder interfaces are duplicated so that each input and output client signal is connected to two transponder interfaces. This arrangement is illustrated in Figure 2-3.

Figure 2-3 Y-Cable Protection Scheme

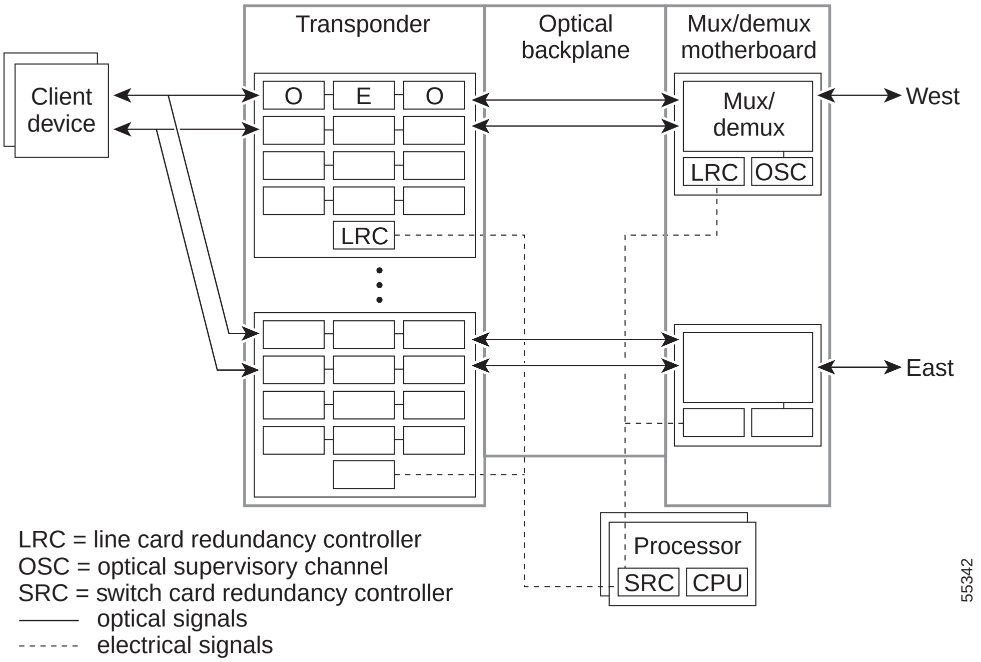

During any interval, one of the transponders is functioning as the active and the other as the standby. On the working transponder, the client side laser is turned on, the trunk side laser is transmitting, and the trunk side receiver is active. On the other transponder, the client side laser is turned off (to avoid corrupting the signal back to the client), the trunk side laser is transmitting, and the trunk side receiver is standing by. The received signal on the trunk side is optically monitored by the LRC. If a loss of light, signal failure, or signal degrade is detected, and an acceptable standby signal is available, a switch to the standby signal is made. The precise conditions that trigger a switchover are configurable in the APS software.

Y-Cable Protection Considerations

The following considerations apply when using y-cable protection:

- Y-cable protection does not protect against failures of the client equipment. To protect against client failures, protection should be implemented on the client itself.
- Due to their lower optical power loss, we recommend using the nonsplitter protected line card motherboards (east or west version) for configurations with y-cable protection.
- The APS software that supports y-cable protection can be configured as revertive or nonrevertive. After a switchover, the working traffic can be put back on the previously failed fiber, once the fault has been remedied, either automatically (revertive) or through manual intervention (nonrevertive).
- Y-cable protected configurations allow monitoring of the protection fiber without the OSC.

For rules on how to configure the shelf for y-cable protection, see Chapter 3, “Shelf Configuration Rules.” For instructions on configuring the APS software for y-cable protection, refer to the *Cisco ONS 15540 ESP Configuration Guide and Command Reference*.

Client Based Line Card Protection

While y-cable protection protects against failures in the transponders or on the fiber, the client still remains vulnerable. For some applications additional protection of the client equipment may be desirable. As Figure 2-4 shows, the same architecture that supports y-cable protection also supports client protection. The only difference is that rather than using a y-cable to split a single client signal, there are two signals from the client equipment. Operationally, client protection is also different in that signal monitoring and switchover are under control of the client rather than the protection mechanisms on the Cisco ONS 15540.

Figure 2-4 Client Based Line Card Protection Scheme

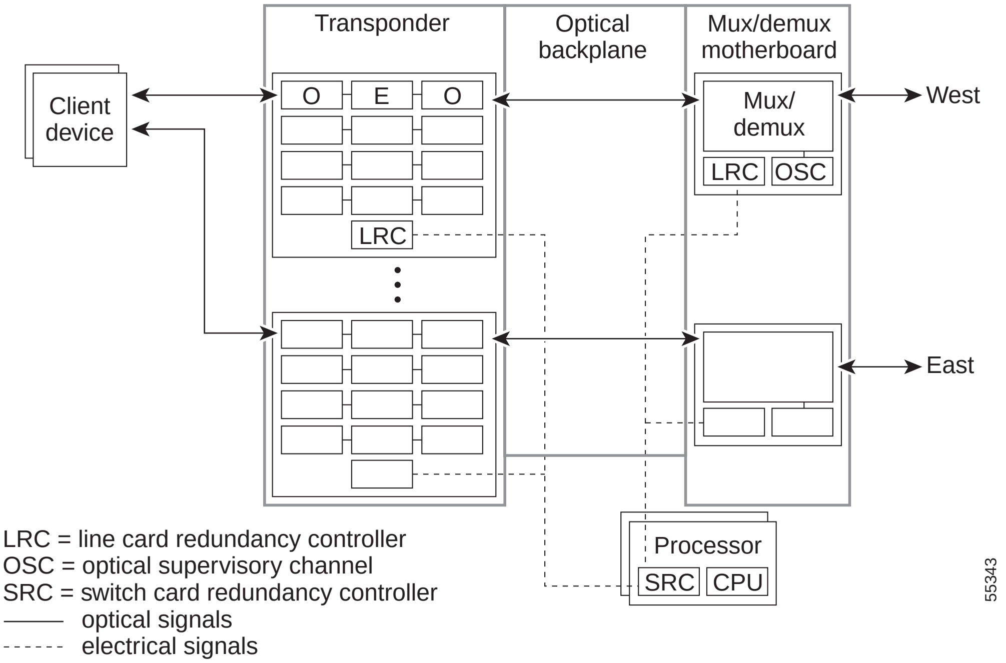

Client Protection Considerations

The following considerations apply when using client protection:

- Client protection uses the same shelf configuration as y-cable based line card protection.
- Due to their lower optical loss, we recommend using the nonsplitter protected line card motherboards (east or west version) for configurations with client protection.
- Client protected configurations allow monitoring of the protection fiber without the OSC.

Supported Topologies

The Cisco ONS 15540 can be used in linear and ring topologies. Linear topologies include protected and unprotected point-to-point and bus. Ring topologies support add/drop nodes and can be hubbed or meshed. The following sections give a brief overview of these topologies.

Linear Topologies

In a pure point-to-point topology all channels terminate on the Cisco ONS 15540 nodes at each end of the trunk. Point-to-point topologies have many common applications, including extending the reach of GE or SONET, and can be configured for unprotected or for protected operation.

Unprotected Point-to-Point Topology

Figure 2-5 shows a point-to-point topology without protection. In this configuration only one optical mux/demux slot is used in each of the Cisco ONS 15540 nodes. The west or east trunk side interface (mux/demux slot 0 or 1) of node 1 connects to the corresponding interface on node 2.

Figure 2-5 Unprotected Point-to-Point Topology

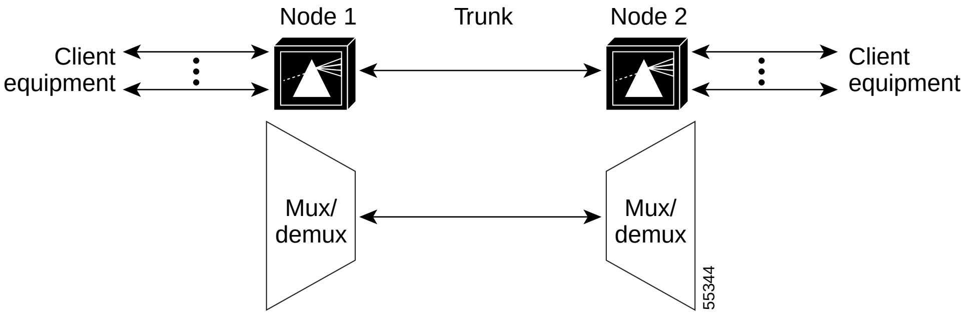

For an example configuration of an unprotected point-to-point topology, see the "Unprotected 32-Channel Point-to-Point Configuration" section on page 5-2.

Protected Point-to-Point Topology

Figure 2-6 shows a protected point-to-point topology. The Cisco ONS 15540 system can be configured for splitter or line card protection. In either case, there are two trunk side interfaces, west and east, connected by two fiber pairs.

Figure 2-6 Protected Point-to-Point Topology

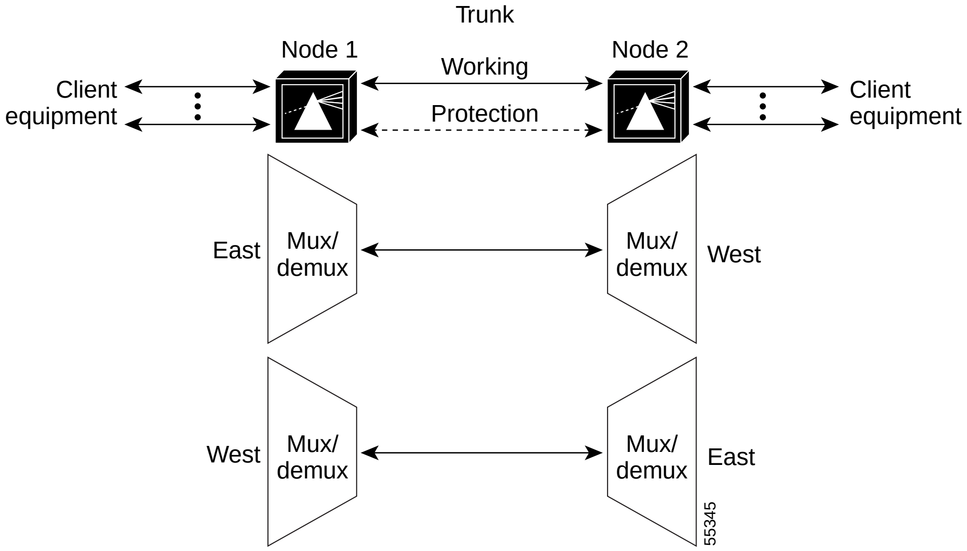

For an example configuration of a protected point-to-point topology, see the "Splitter Protected 32-Channel Point-to-Point Configuration" section on page 5-4 and the "Line Card Protected 16-Channel Point-to-Point Configuration" section on page 5-6.

Bus Topology

In a bus topology, sometimes called *linear add/drop*, there is an intermediate add/drop node between the two terminal nodes. Figure 2-7 shows an example of this type of topology. Bands A and B (channels 1–4 and 5–8) terminate at node 1. Band B is dropped at node 2, which passes band A through. To support this configuration, the add/drop node must have add/drop mux/demux modules in both slots 0 and 1 for west and east directions.

Figure 2-7 Bus Topology

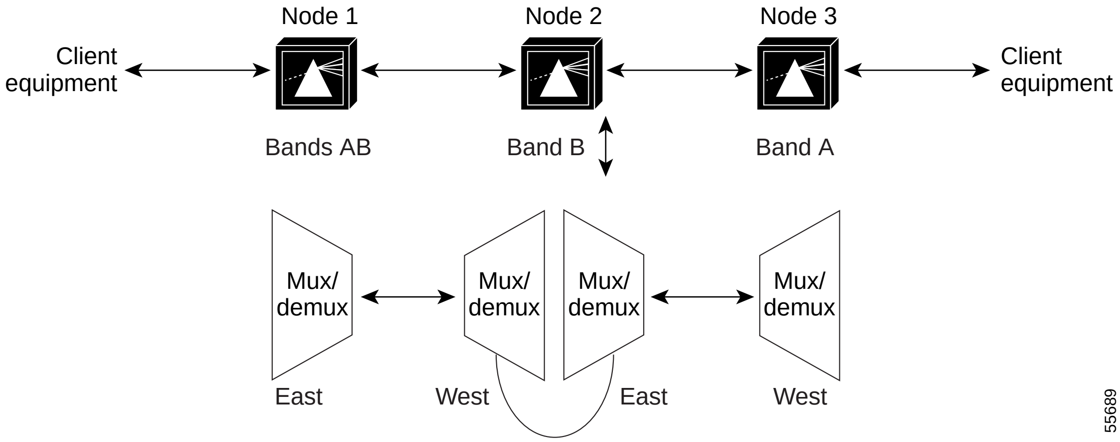

This type of topology offers limited protection. In this example a failure on the link between node 2 and node 3 would result in the loss of band A, but band B would remain operational between nodes 1 and 2. A failure on the link between node 1 and node 2 would result in a loss of both bands.

Note If protection for all bands is required in a topology such as this, where there is a single add/drop node between two terminal nodes, a ring can be used.

Ring Topologies

In a ring topology, client equipment is attached to three or more Cisco ONS 15540 systems, which are interconnected in a closed loop. Channels are dropped and added at two nodes on the ring. Rings have many applications, including providing extended access to SANs (storage area networks) and upgrading existing SONET rings. In the cases where SONET rings are at capacity, the SONET equipment can be moved off the ring and connected to the Cisco ONS 15540 systems. Then the SONET client signals are multiplexed and transported over the DWDM link, thus increasing the capacity of existing fiber.

Hubbed Ring

A hubbed ring is composed of a hub node and two or more add/drop or satellite nodes. All channels on the ring originate and terminate on the hub node. At the add/drop node certain channels are terminated (dropped and added back) while the channels that are not being dropped (express channels) are passed through optically, without being electrically regenerated.

Channels are dropped and added in bands. Figure 2-8 shows a four-node hubbed ring in which bands ABC terminate on node 1. Nodes 1 and 2 communicate using band A, nodes 1 and 3 communicate using band B, and nodes 1 and 4 communicate using band C.

Figure 2-8 Hubbed Ring Topology Example

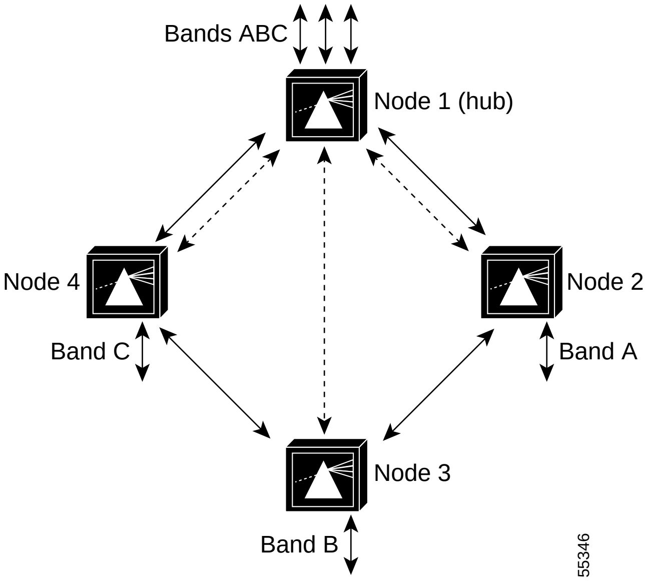

For example configurations of hubbed ring topologies, see the "Hubbed Ring Topologies" section on page 5-11.

Meshed Ring

A meshed ring is a physical ring that has the logical characteristics of a mesh. While traffic travels on a physical ring, the logical connections between individual nodes are meshed. An example of this type of configuration, which is sometimes called a *logical mesh*, is shown in Figure 2-9. Nodes 1 and 2 communicate using band A, nodes 1 and 3 communicate using band B, and nodes 1 and 4 communicate using band C, as in the previous example (Figure 2-8). In this example, however, the fact that nodes 3 and 4 communicate independently using band D makes it a meshed ring.

Figure 2-9 Meshed Ring Topology Example

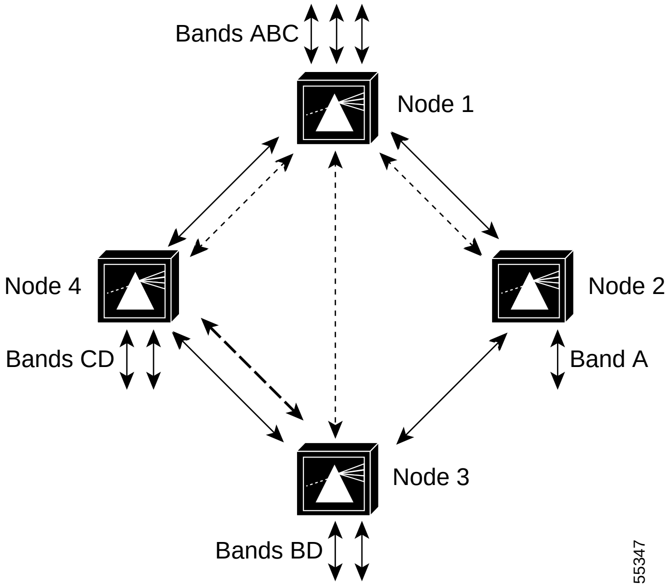

For example configurations of meshed ring topologies, see the "Meshed Ring Topologies" section on page 5-33.

Protection in Ring Topologies

Protection in the Cisco ONS 15540 is supported using per-channel unidirectional path switching or bidirectional path switching. Protection mechanisms are implemented in both the hardware and the APS software.

Unidirectional Path Switching

Unidirectional path switching is based on a variant of the SONET UPSR (Unidirectional Path Switched Ring). For each channel on a ring, traffic is transmitted in both directions from each node, using one of the fiber pair for each direction; on each node, traffic is received from one direction. In the event of a failure of the receive signal at a node, the node switches over to receive the signal from the other direction. Switching decisions are made on a node-by-node basis, and some channels can be received from one direction while others are received from the other direction. Protection in rings can be implemented on the Cisco ONS 15540 using either a splitter or y-cable configuration.

Figure 2-10 shows a three-node hubbed ring. In the example, node 1 is sending traffic from both east and west mux/demux interfaces, but is actively receiving only on the west side.

Figure 2-10 Per-Channel Unidirectional Path Switching in Normal State

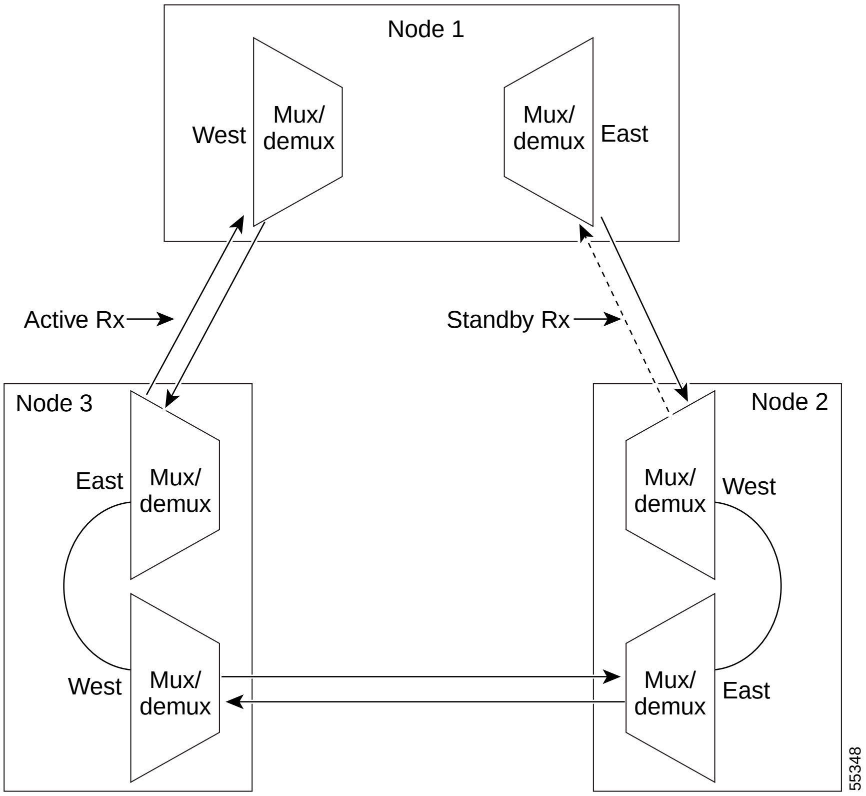

Figure 2-11 shows a failure on the receive signal from node 3 to node 1. In this event, node 1 switches to the receive signal on its east mux/demux interface for the failed channel(s). Assuming that the failure is only in one direction between node 1 and node 3, node 3 would not be required to switch receive directions. If, however, both directions were affected (for example, a fiber cut) and node 3 had been receiving on its east side, it would switch to the west side for its active receive signal.

Figure 2-11 Per-Channel Unidirectional Path Switching after Protection Switch

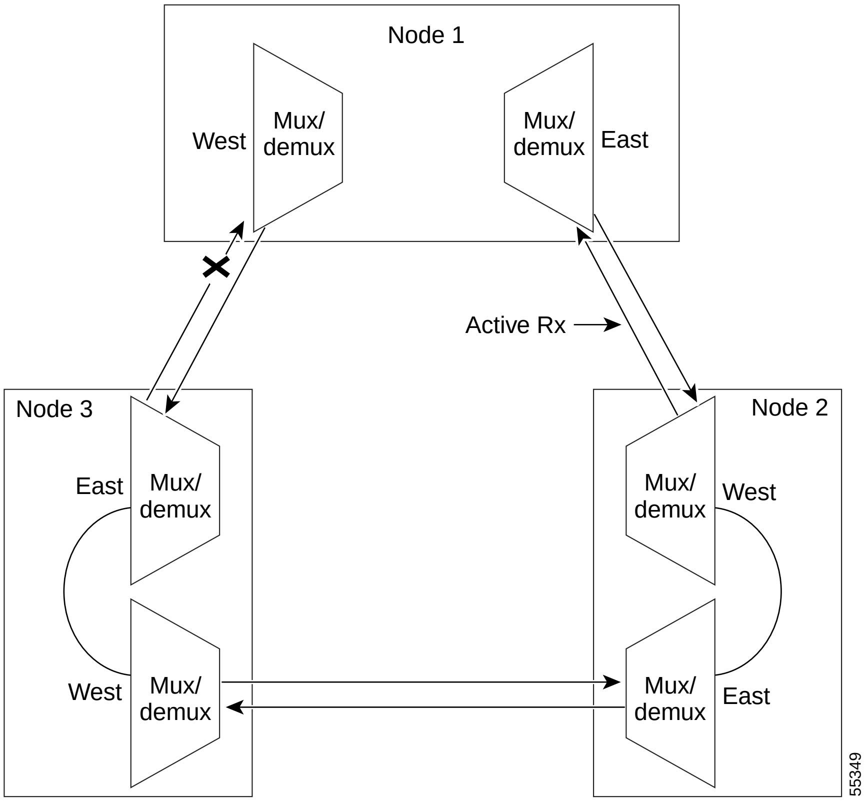

Bidirectional Path Switching

The Cisco ONS 15540 also supports bidirectional switching through the APS software. When configured for bidirectional switching, a node that detects a fault sends a signal over the OSC to the source node to also switch its receive direction. Assume, for example, that the channels configured between node 1 and node 3 in Figure 2-11 were communicating over the link that fails. When node 1 switches to receive those channels from the east side (over node 2), node 3 would also switch to receive those channels from its west side. This ensures that the distance between the two nodes remains the same for those channels. This option is supports protocols that are distance sensitive.