Topologies

The following Cisco ONS 15454 DWDM network topologies are described in this chapter:

4.1 Overview

There are two main DWDM network types: Metro Core, where the channel power is equalized and dispersion compensation is applied, and Metro Access, where the channels are not equalized and dispersion compensation is not applied. Metro Core networks often include multiple spans and amplifiers, thus making optical signal-to-noise ratio (OSNR) the limiting factor for channel performance. Metro Access networks often include a few spans with very low span loss; therefore, the signal link budget is the limiting factor for performance.

4.2 Hubbed Rings

In the hubbed ring topology (Figure 4-1), a hub node terminates all the DWDM channels. A channel can be provisioned to support protected traffic between the hub node and any node in the ring. Both working and protected traffic use the same wavelength on both sides of the ring. Protected traffic can also be provisioned between any pair of optical add/drop multiplexer (OADM) nodes, except that either the working or the protected path must be regenerated in the hub node.

Protected traffic saturates a channel in a hubbed ring, which means that channel reuse is not available. However, the same channel can be reused in different sections of the ring by provisioning unprotected multihop traffic. From a transmission point of view, this network topology is similar to two bidirectional point-to-point links with OADM nodes.

Figure 4-1 Hubbed Ring

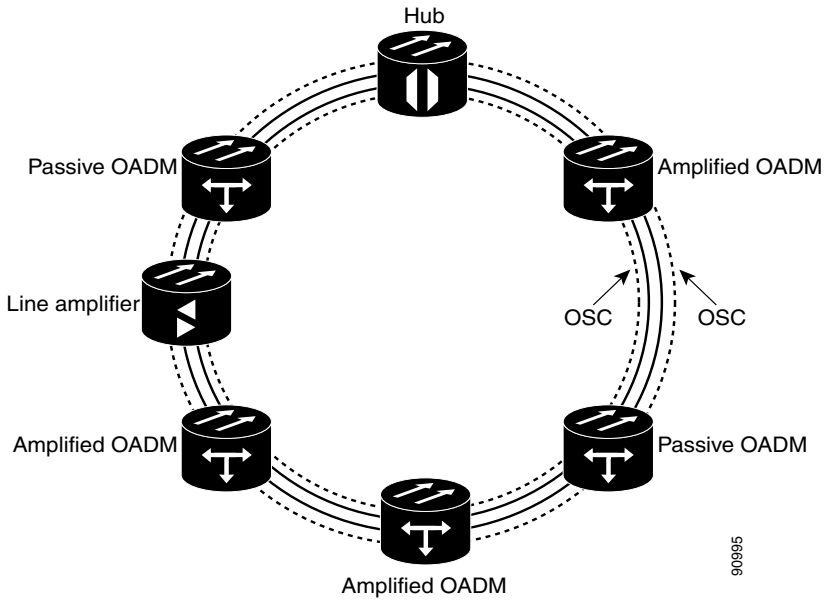

4.3 Multihubbed Rings

A multihubbed ring (Figure 4-2) is based on the hubbed ring topology, except that two or more hub nodes are added to make a multihubbed ring. Protected traffic can only be established between the two hub nodes. Protected traffic can be provisioned between a hub node and any OADM node if the allocated wavelength channel is regenerated through the other hub node. Multihop traffic can be provisioned on this ring. From a transmission point of view, this network topology is similar to two or more point-to-point links with OADM nodes.

Figure 4-2 Multihubbed Ring

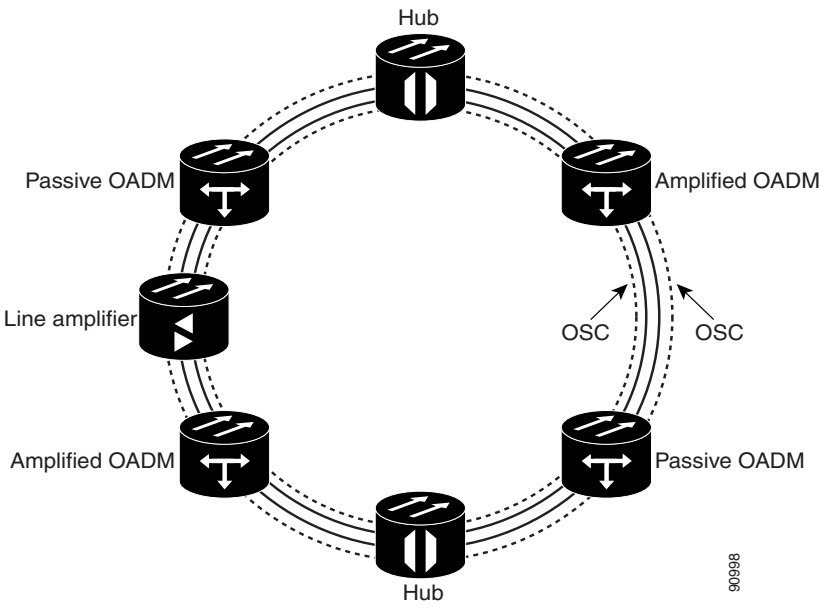

4.4 Any-to-Any Rings

The any-to-any ring topology shown in Figure 4-3 contains only reconfigurable OADM (ROADM) nodes, or ROADM nodes with optical service channel (OSC) regeneration or amplifier nodes. The any-to-any ring topology allows you to route every wavelength from any source to any destination node inside the network.

Figure 4-3 Any-to-Any Ring

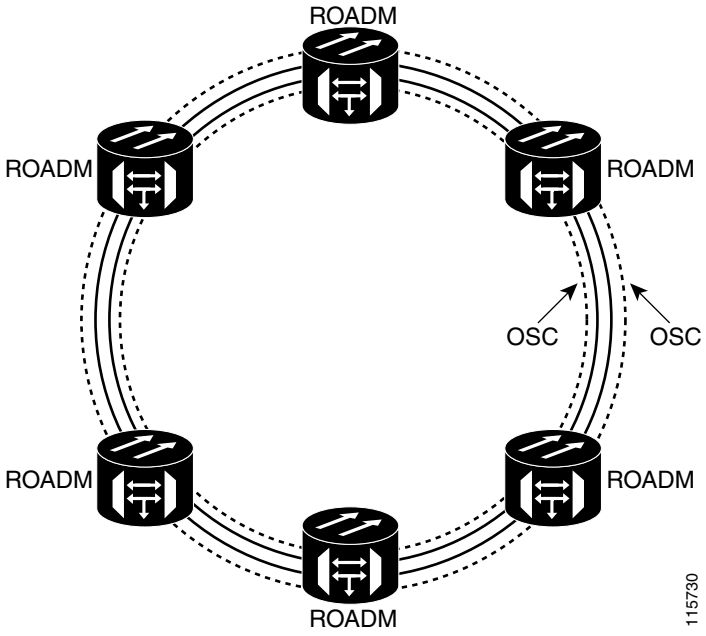

4.5 Meshed Rings

In a meshed ring topology (Figure 4-4), protected traffic can be provisioned between any two nodes; however, the selected channel cannot be reused in the ring. Unprotected multihop traffic can be provisioned in the ring. A meshed ring must be designed to prevent amplified spontaneous emission (ASE) lasing. ASE lasing is prevented by configuring a particular node as an anti-ASE node. An anti-ASE node can be created in two ways:

- Equip an OADM node with 32MUX-O cards and 32DMX-O cards. This solution is adopted when the total number of wavelengths deployed in the ring is higher than ten. OADM nodes equipped with 32MUX-O cards and 32DMX-O cards are called full OADM nodes.
- When the total number of wavelengths deployed in the ring is lower than ten, the anti-ASE node is configured by using an OADM node where all the channels that are not terminated in the node are configured as “optical pass-through.” In other words, no channels in the anti-ASE node can travel through the express path of the OADM node.

Note The example in Figure 4-4 does not use hubbed nodes; only amplified and passive OADM nodes are present.

Figure 4-4 Meshed Ring

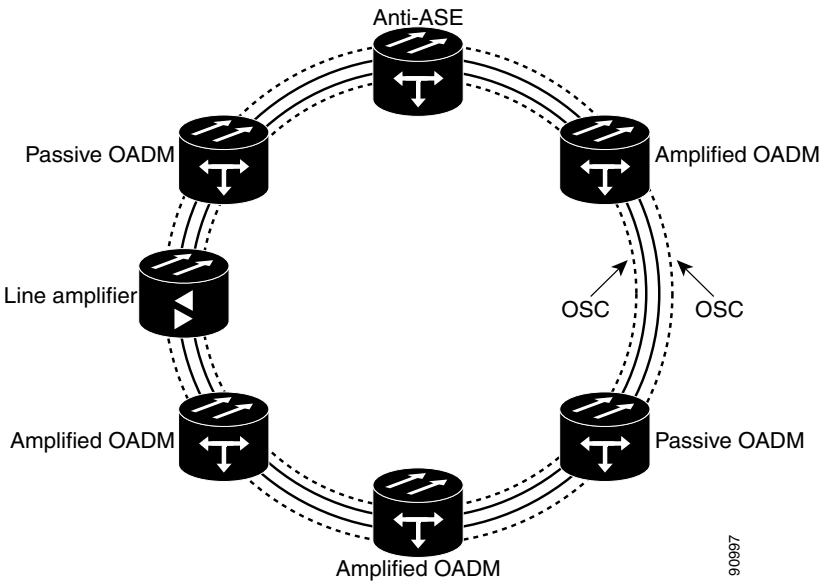

4.6 Linear Configurations

Linear configurations are characterized by the use of two terminal nodes (west and east). The terminal nodes must be equipped with a 32MUX-O card with a 32DMX-O card, or a 32WSS card with a 32DMX or 32DMX-O card. OADM or line amplifier nodes can be installed between the two terminal nodes. Only unprotected traffic can be provisioned in a linear configuration. Figure 4-5 shows five ONS 15454 nodes in a linear configuration with an OADM node.

Figure 4-5 Linear Configuration with an OADM Node

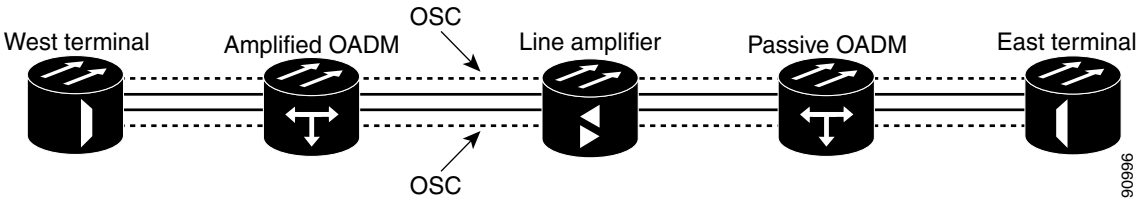

Figure 4-6 shows five ONS 15454 nodes in a linear configuration without an OADM node.

Figure 4-6 Linear Configuration without an OADM Node

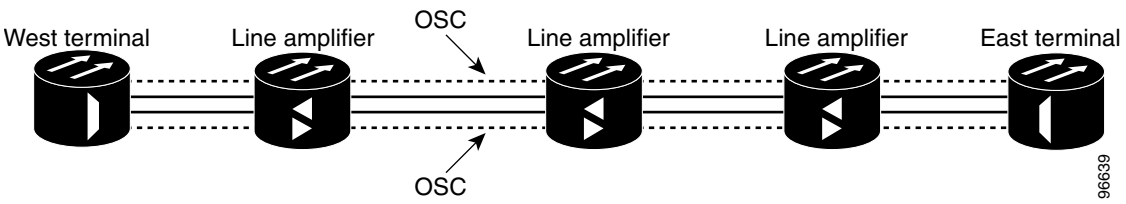

4.7 Single-Span Link

Single-span link is a type of linear configuration characterized by a single-span link with preamplification and post-amplification. A span link is also characterized by the use of two terminal nodes (west and east). The terminal nodes are usually equipped with a 32MUX-O card and a 32DMX-O card. However, a 32WSS card and a 32DMX or a 32DMX-O card can be installed. Single-span links with AD-4C-xx.x cards ar e also supported. Only unprotected traffic can be provisioned on a single-span link.

Figure 4-7 shows ONS 15454 nodes in a single-span link. Eight channels are carried on one span. Single-span link losses apply to OC-192 LR ITU cards. The optical performance values are valid if that the sum of the insertion losses and span losses for the passive OADM nodes does not exceed 35 dB.

Figure 4-7 Single-Span Link

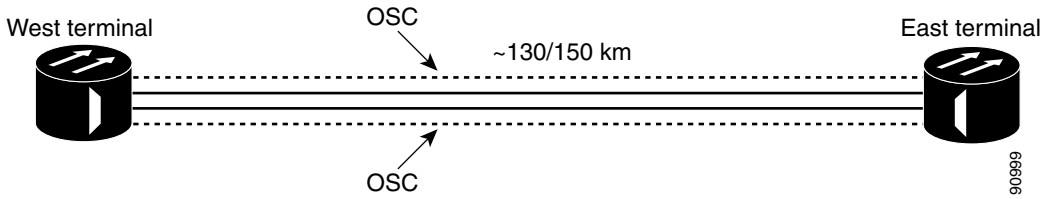

4.8 1+1 Protected Single-Span Link

A 1+1 protected, single-span link configuration uses a single hub or OADM node connected directly to the far-end hub or OADM node through four-fiber links. This node configuration is used in a ring configured with two point-to-point links. The advantage of the 1+1 protected flexible terminal node configuration is that it provides path redundancy for 1+1 protected time division multiplexing (TDM) networks (two transmit paths and two receive paths) using half of the DWDM equipment that is usually required. In the example shown in Figure 4-8, one node transmits traffic to the other node on both the east side and the west side of the ring for protection purposes. If the fiber is damaged on one side of the ring, traffic still arrives safely through fiber on the other side of the ring.

Figure 4-8 Double Terminal Protection Configuration

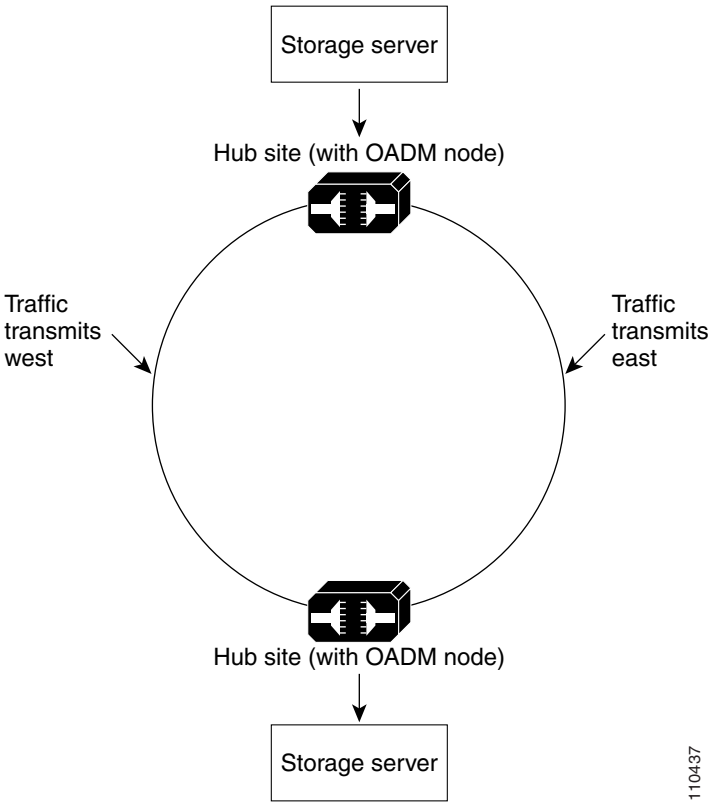

Figure 4-9 shows a functional block diagram of a 1+1 protected, single-span link with hub nodes. A 1+1 protected, single-span link with hub nodes cannot be used in a hybrid configurations.

Figure 4-9 1+1 Protected Single-Span Link with Hub Nodes

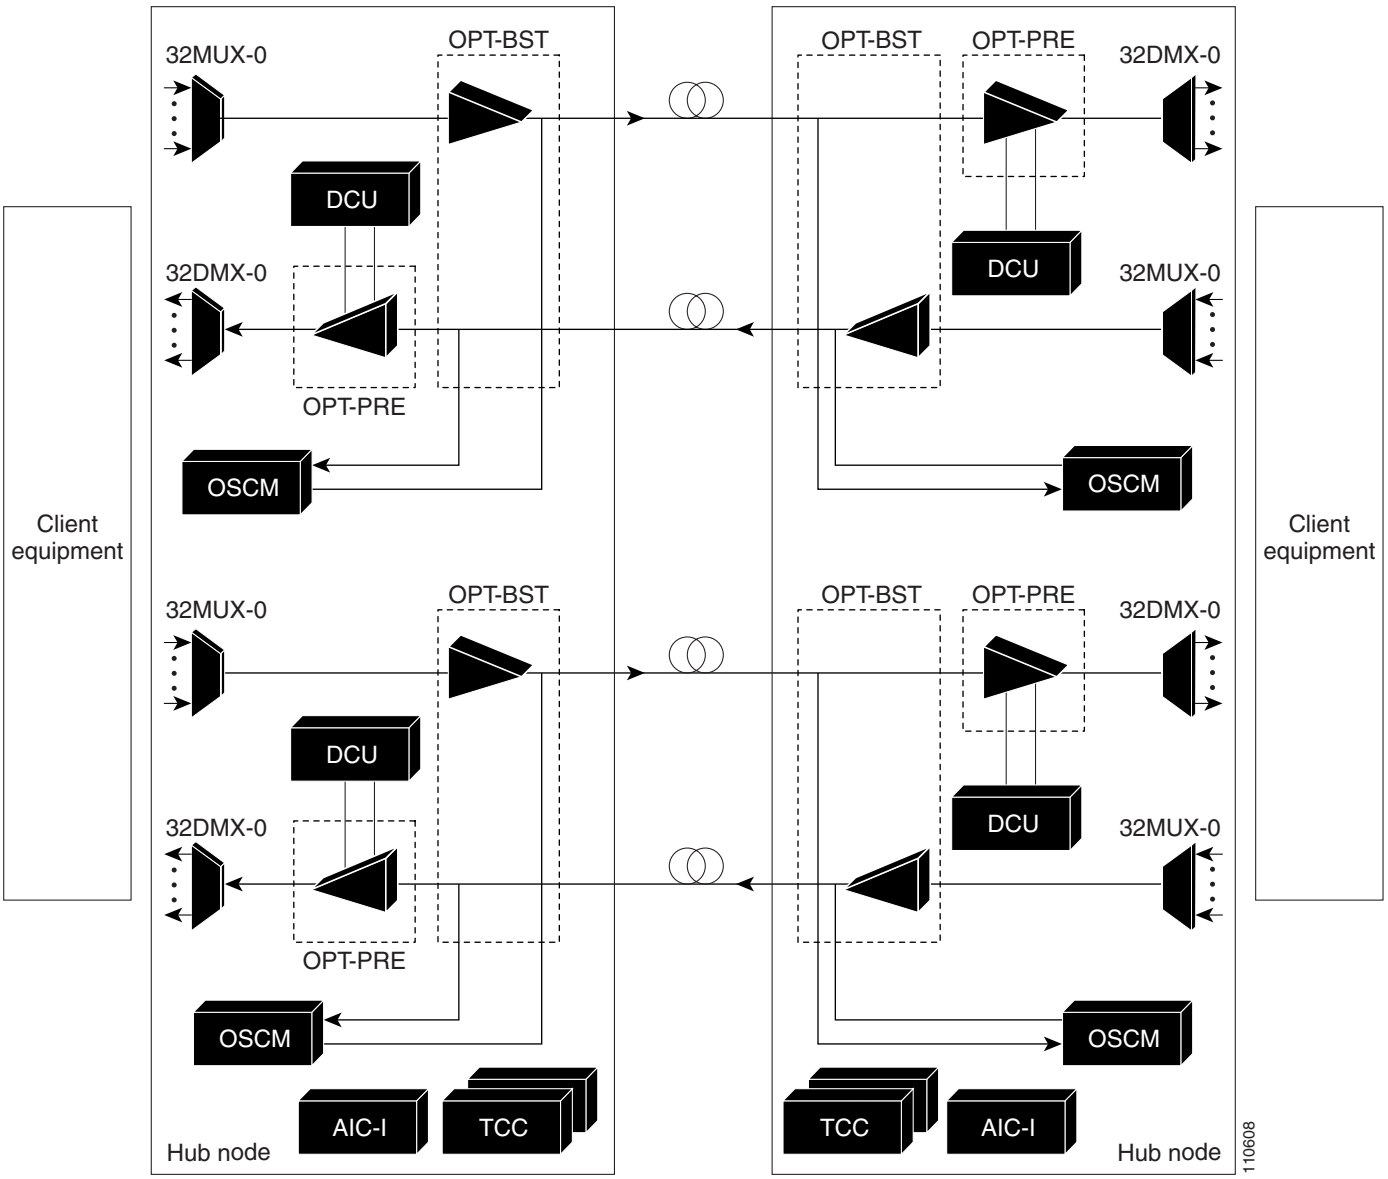

Figure 4-10 shows a functional block diagram of a 1+1 protected, single-span link with active OADM nodes. A 1+1 protected, single-span link with active OADM nodes can be used in a hybrid configurations.

Figure 4-10 1+1 Protected Single-Span Link with Active OADM Nodes

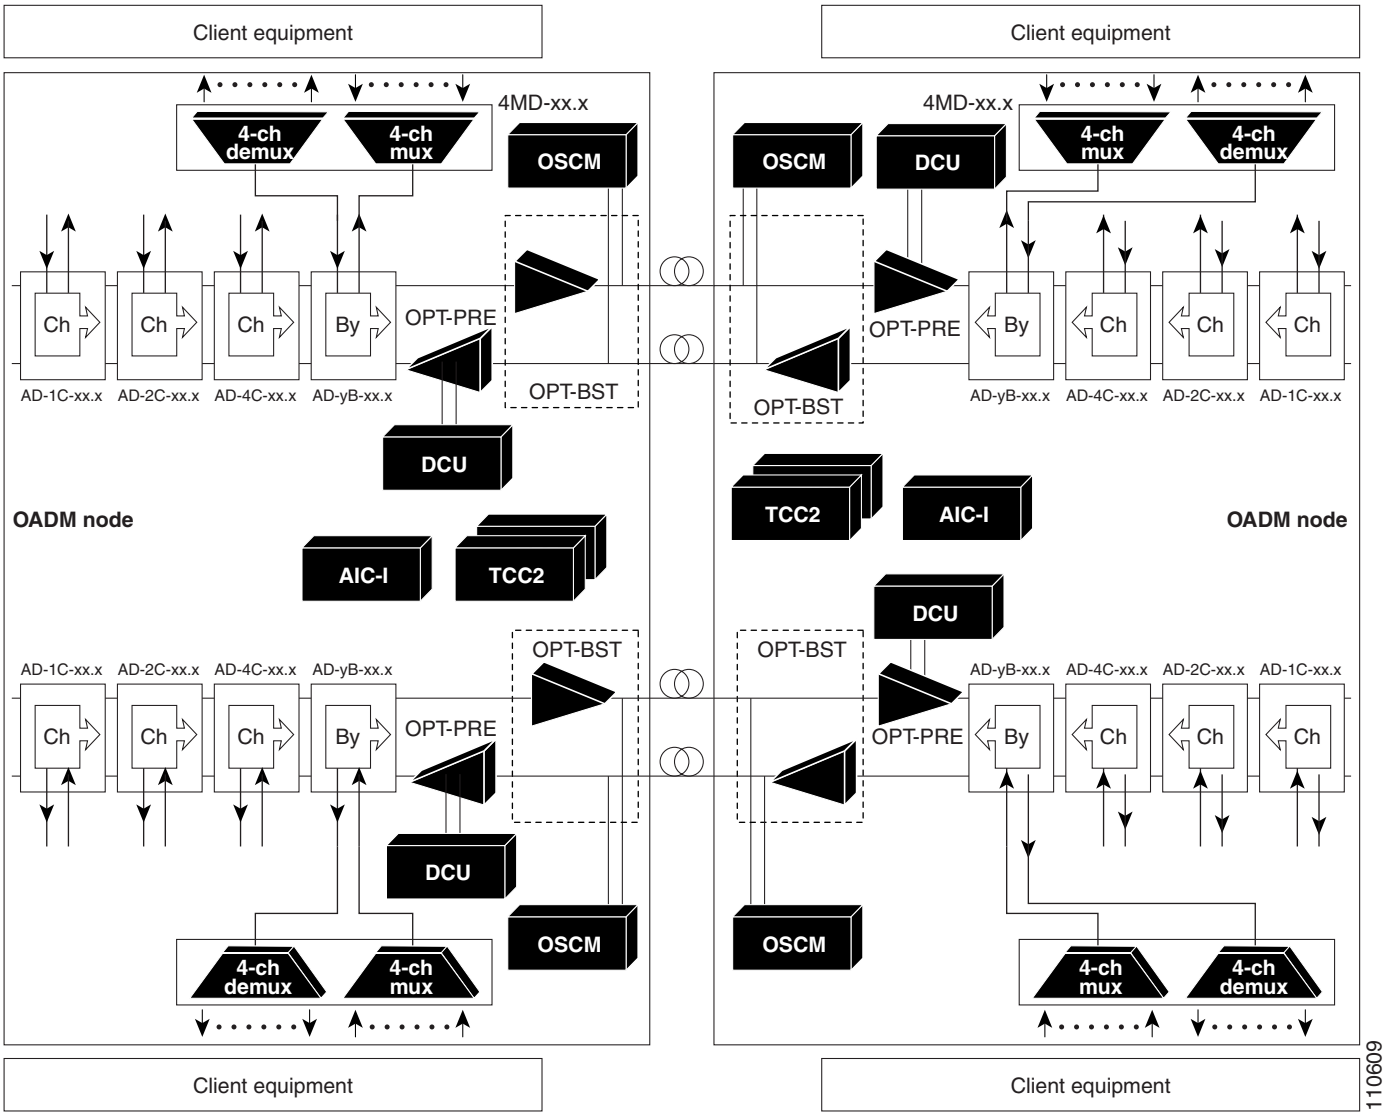

Figure 4-11 shows a functional block diagram of a 1+1 protected, single-span link with passive OADM nodes. 1+1 protected, single-span links with passive OADM nodes can be used in a hybrid configurations.

Figure 4-11 1+1 Protected Single-Span Link with Passive OADM Nodes

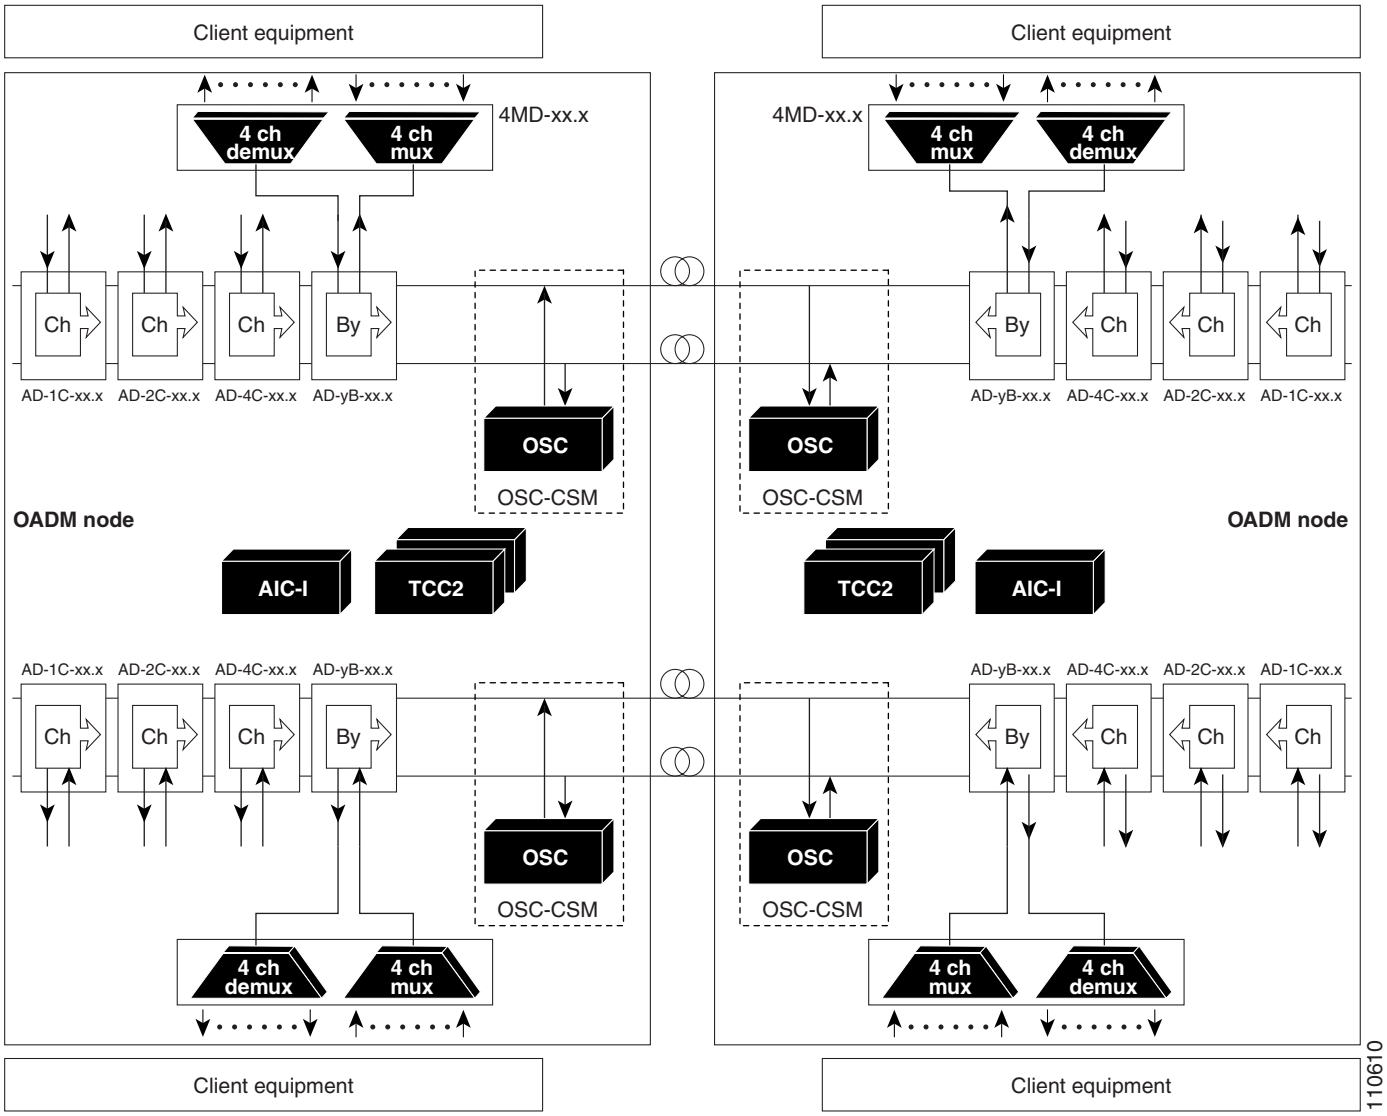

4.9 Hybrid Network

The hybrid network configuration is determined by the type of node that is used in an ONS 15454 network. Along with TDM nodes, the ONS 15454 supports the following hybrid node types: 1+1 protected flexible terminal, scalable terminal, hybrid terminal, hybrid OADM, hybrid line amplifier, and amplified TDM.

Figure 4-12 shows ONS 15454 nodes in a hybrid TDM and DWDM configuration.

Figure 4-12 Hybrid Network Example

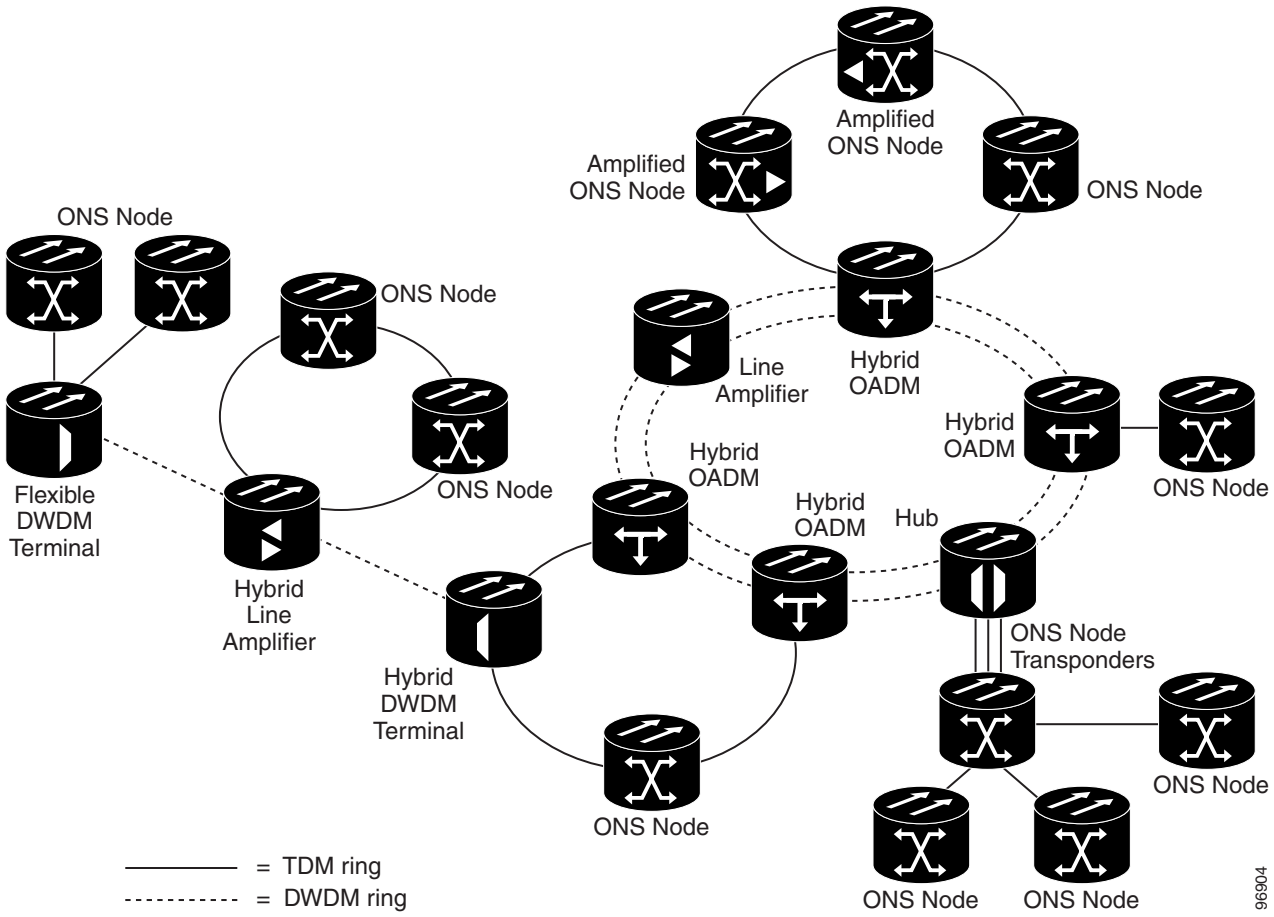

DWDM and TDM layers can be mixed in the same node; however they operate and are provisioned independently. The following TDM configurations can be added to a hybrid network:

- Point-to-point
- Linear add/drop multiplexing (ADM)
- Bidirectional line switched ring (BLSR)
- Path Protection

Figure 4-13 shows ONS 15454 nodes in a hybrid point-to-point configuration.

Figure 4-13 Hybrid Point-to-Point Network Example

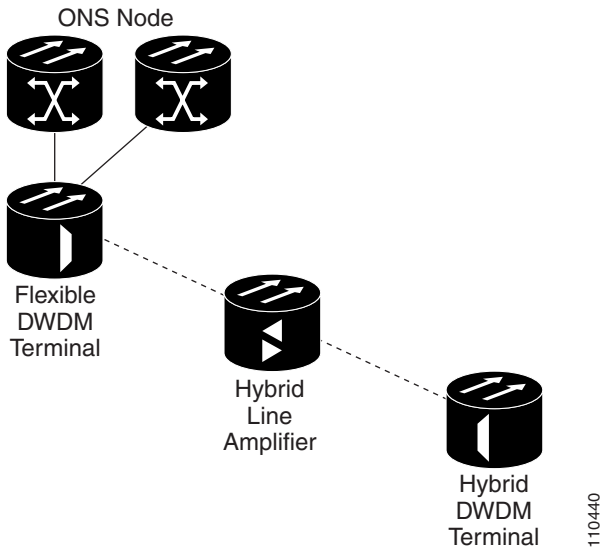

Figure 4-14 shows ONS 15454 nodes in a hybrid linear ADM configuration.

Figure 4-14 Hybrid Linear ADM Network Example

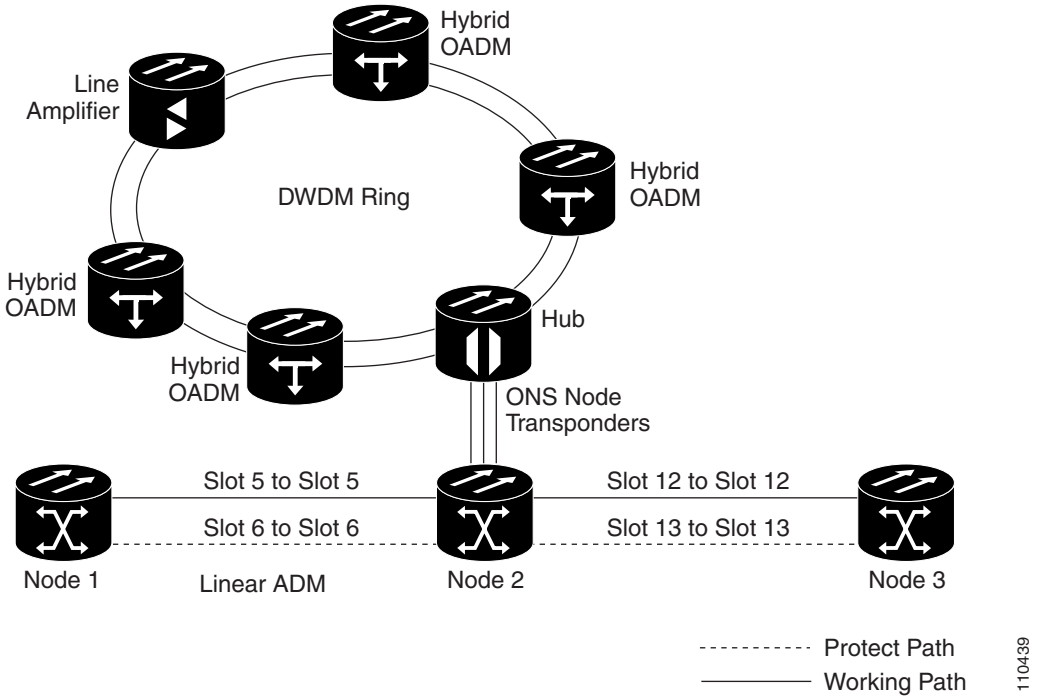

Figure 4-15 shows ONS 15454 nodes in a hybrid BLSR configuration.

Figure 4-15 Hybrid BLSR Network Example

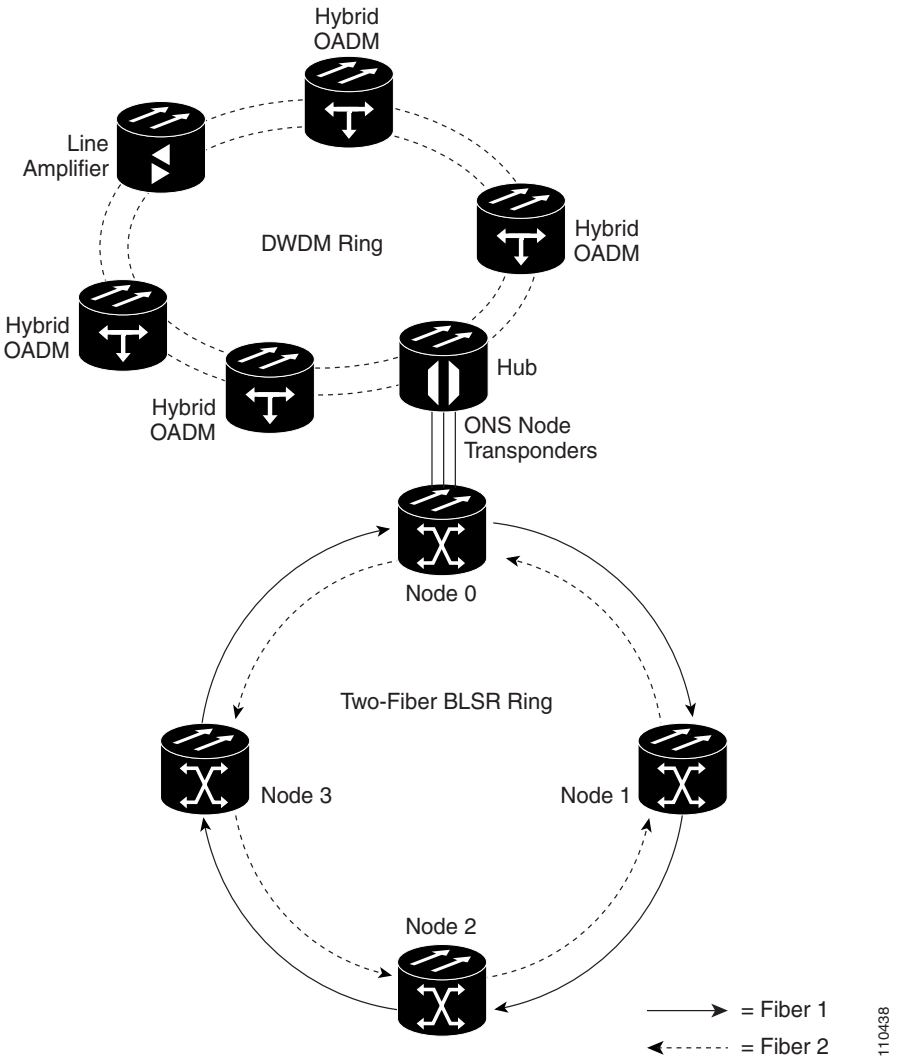

Figure 4-16 shows ONS 15454 nodes in a hybrid path protection configuration.

Figure 4-16 Hybrid Path Protection Network Example

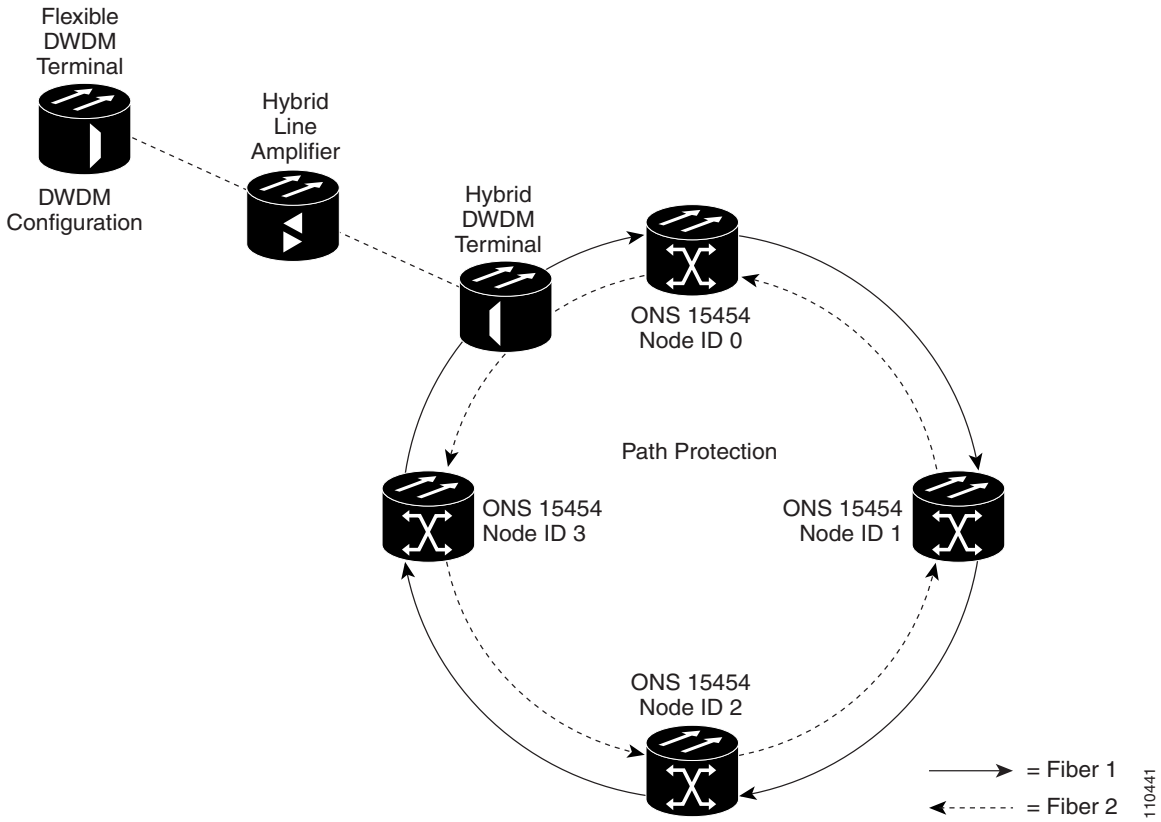

4.10 Transponder and Muxponder Protection Topologies

The ONS 15454 supports Y-cable and splitter protection for transponder (TXP) and muxponder (MXP) cards. Figure 4-17 shows Y-cable and splitter protection.

Figure 4-17 Y-Cable and Splitter Protection

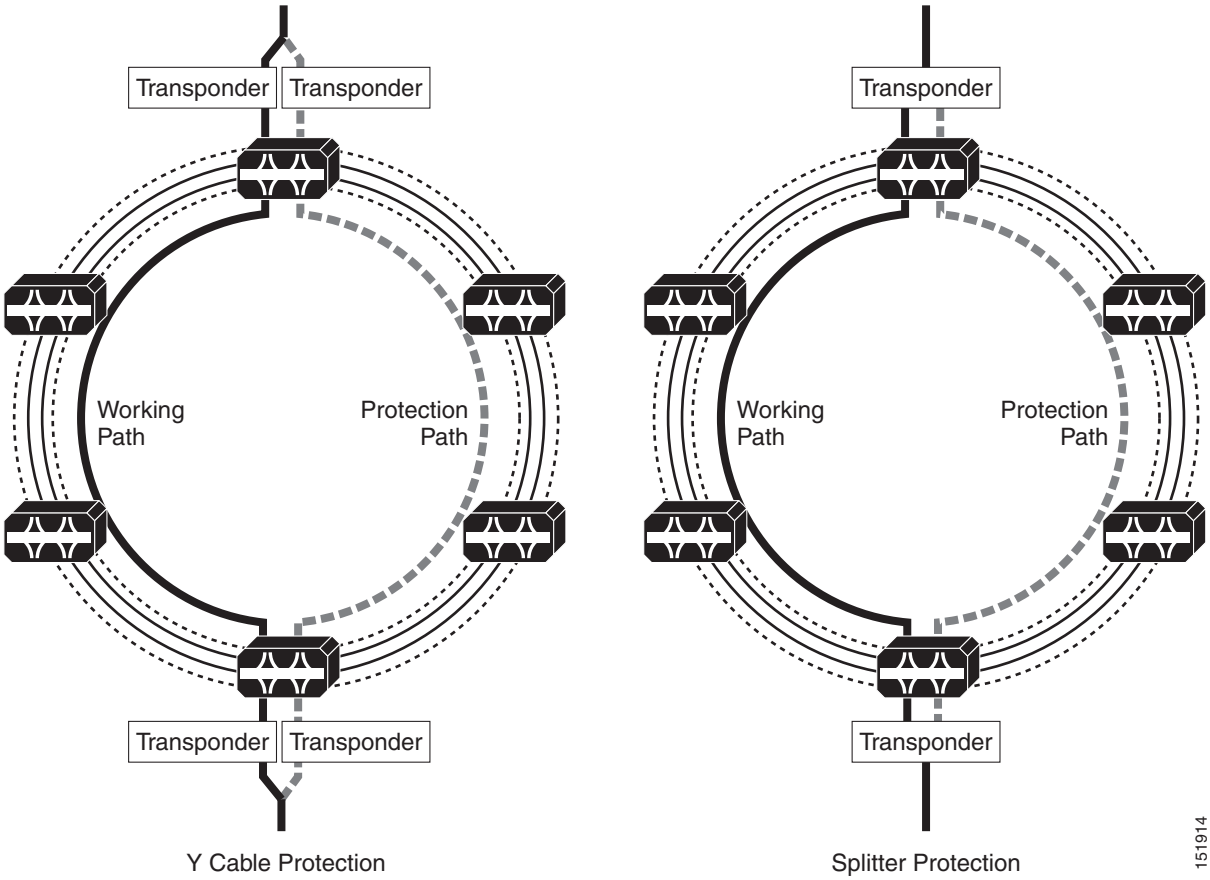

4.10.1 Y-Cable Protection

The Y-cable protection scheme employs two Y cables, which are hardware combiner/splitters. A signal injected into the stem of the Y is duplicated into both arms of the Y with 50 percent attenuation in each arm. Signals injected into the arms of the Y are summed into the stem of the Y.

A Y-cable protection group requires two DWDM cards with the arms of the Y-cables connected to the client ports on the DWDM cards, and the stems of the Y-cables connected to the client source, such as an OC-N card. When a TXP Y-cable protection group is required, the two TXP cards must be installed in the same shelf assembly in adjacent slots.

Figure 4-18 shows a functional block diagram of Y-cable protection.

Figure 4-18 Y-Cable Protection

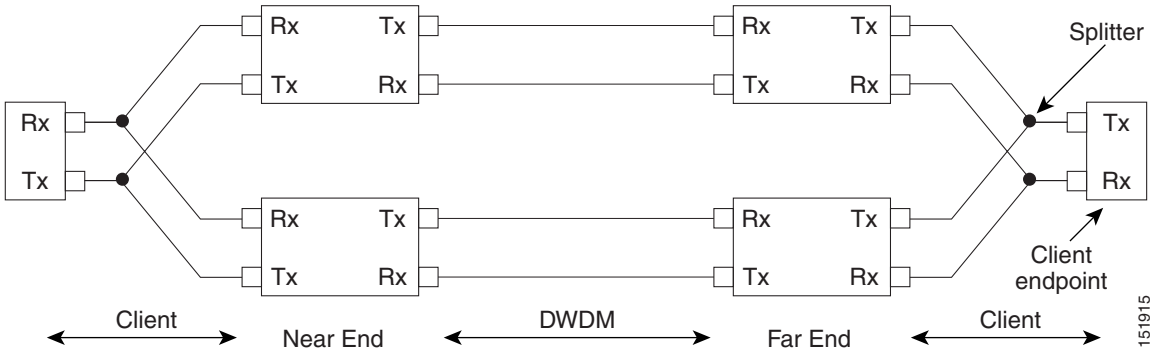

A Y-cable protection group has two paths:

- Transmit (TX) path, defined as the client RX and the trunk TX
- Receive (RX) path, defined as the trunk RX and the client TX

The basic behavior of the Y-cable group is that an incoming client signal is bridged to the two TX paths, and one RX path is selected for the outgoing client signal. Thus, a Y-cable protection group only protects against defects in the RX path. Figure 4-19 shows the RX path for the near-end Y-cable protection group.

Figure 4-19 Rx Path for Near-End Y-Cable Protection

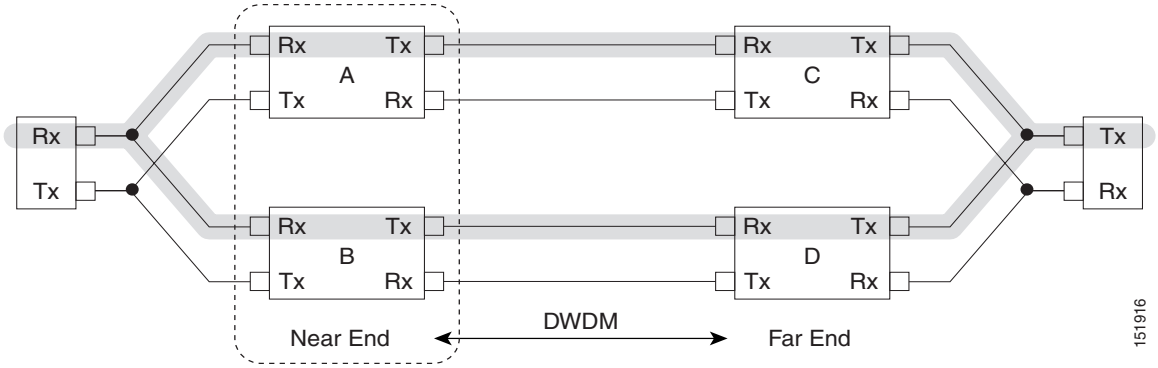

To protect against all defects, a pair of Y-cable protection groups is required. Each protects against defects in its own RX path. Figure 4-20 shows how the near-end and the far-end jointly protect against defects.

Figure 4-20 Near-End and Far-End Y-Cable Protection

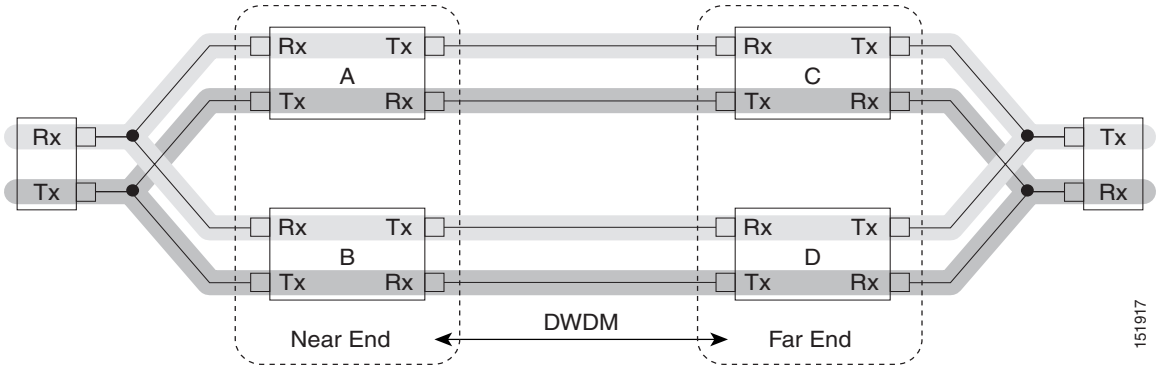

A Y-cable protection group is defined by two client ports on two different cards. One client port is designated as working and the other is designated as protect. Some of the rules in a Y-cable protection are as follows:

- The cards must have the same equipment type.
- The cards must have the same payload data type.
- The cards must have the same termination mode.
- The client ports must have the same payload type.
- The client ports must have the same facility number.

For example, a Y-cable protection group can include MXP Client 2 on the working and protect cards, but cannot include Client 2 on the working card if Client 3 is on the protect card. TXP cards have a single client (facility 1), so this requirement is satisfied by default.

Zero, one, two, or all of the client ports on an MXP card can be in Y-cable protection groups. Some clients can be in a protection group while others are unprotected.

The client ports on a card that are in Y-cable protection groups are either all working or all protect. You cannot mix working and protect client ports on the same card. For convenience, the trunk ports adopt the designation (working or protect) of the client ports.

The Y-cable protection groups on an MXP card switch independently. A Y-cable protection group performs protection switching by disabling the transmitter on the standby client port and enabling the active client port.

The protection group does not enable the active transmitter, because the active transmitter may have been disabled for other reasons. The port is disabled if it has an OOS-DBLD service state, is squelched, or is shutdown by automatic laser shutdown (ALS). The protection group releases a signal that the active transmitter is disabled. This activity changes the RX path but not the TX path. A Y-cable protection group can only protect its RX path.

A Y-cable protection group enables its client receivers unless the client facilities have an OOS-DSBLD service state. This means that client receivers (and trunk transmitters) are operational regardless of the active/standby status of the card. Traffic might not be lost if both client lasers in a Y-cable protection group are enabled. If the output powers of the two lasers are not identical, then the receiver at the stem of the Y-cable can opt for the stronger client laser and ignore the weaker signal.

4.10.2 Splitter Protection

A splitter protection group consists of a single 2.5-Gbps transponder splitter (TXPP_MR_2.5G card). The protection group is defined by the two trunk ports on the splitter card. One trunk port is designated working and the other is designated as protect. A splitter card has a single trunk laser and a hardware splitter that duplicates the trunk signal out of the card's two trunk ports. The switch in the card receives one of the two trunk input signals and the received signal is connected to the client ports.

A splitter protection group has two TX paths and two RX paths on the same card; the paths share client ports. The TX path is defined as the client RX and the trunk TX, and the RX path is defined as the trunk RX and the client TX. In a splitter group, an incoming client signal is bridged to the TX paths, and one RX path is selected for the outgoing client signal.

A splitter protection group performs switching by enabling the receiver of the active trunk port and then routing the active trunk traffic to the client ports. The protection group does not disable the transmitter on the standby trunk port.

4.10.3 Switch Criteria

Cisco Transport Controller (CTC), the ONS 15454 software interface, performs protection switches based on priority, trunk and client line conditions, switch commands, unidirectional/bidirectional switching, and other criteria.

4.10.3.1 Switch Priority

Switch priorities are defined in Table 4-1.

Table 4-1 Switch Priorities

Request/State	Abbreviation	Priority
Lockout of Protection	LO	8 (highest)
Forced Switch	FS	7
Signal Fail	SF	6
Signal Degrade	SD	5
Manual Switch	MS	4
Wait to Restore	WTR	3
Do Not Revert	DNR	2
No Request	NR	1 (lowest)

All switch criteria are assigned a numerical priority, which is reversed from ITU-T G.873.1 to avoid confusion when comparing priorities. In this document, a higher priority is numerically greater than a lower priority.

If the protection channel and the working channel have conditions with the same priority, and the priority is greater than Do Not Revert (DNR), then the condition on the protection channel takes precedence.

In practice, only two priorities can exist independently and simultaneously on the working and protection channels: signal fail (SF) and signal degrade (SD). This requirement means that if both channels have, for example, an SF condition without any higher conditions present, then the protection group chooses the working channel. Traffic switches away from the highest-priority condition.

4.10.3.2 Line Conditions on the Trunk

The following line conditions on the trunk generate the priorities given:

- OTUn-LOS on a trunk, if ITU-T G.709 is enabled, has an SF priority.
- OTUn-LOF on a trunk, if ITU-T G.709 is enabled, has an SF priority.
- OTUn-LOM on a trunk, if ITU-T G.709 is enabled, has an SF priority.
- OTUn-AIS on a trunk, if ITU-T G.709 is enabled, has an SF priority.
- ODUn-AIS on a trunk, if ITU-T G.709 is enabled, has an SF priority.
- OTU BER SF on a trunk, if ITU-T G.709 is enabled, has an SF priority.
- OTU BER SD on a trunk, if ITU-T G.709 is enabled, has an SD priority.
- TIM on OTU SM TTI on a trunk, if ITU-T G.709 is enabled, has an SF priority.
- TIM on ODU PM TTI on a trunk, if ITU-T G.709 is enabled, has an SF priority.
- S-LOF on a trunk, if ITU-T G.709 is disabled and the trunk is OCn, has an SF priority.
- S-LOS on a trunk, if ITU-T G.709 is disabled and the trunk is OCn, has an SF priority.
- SF on a trunk, if ITU-T G.709 is disabled and the trunk is OCn, has an SF priority.
- SD on a trunk, if ITU-T G.709 is disabled and the trunk is OCn, has an SD priority.
- AIS-L on a trunk, if ITU-T G.709 is disabled and the trunk is OCn, has an SF priority.
- RS-LOF on a trunk, if ITU-T G.709 is disabled and the trunk is STMn, has an SF priority.
- RS-LOS on a trunk, if ITU-T G.709 is disabled and the trunk is STMn, has an SF priority.
- SF on a trunk, if ITU-T G.709 is disabled and the trunk is STMn, has an SF priority.
- SD on a trunk, if ITU-T G.709 is disabled and the trunk is STMn, has an SD priority
- MS-AIS on a trunk, if ITU-T G.709 is disabled and the trunk is STMn, has an SF priority.
- TIMS-S on J0 on a trunk, if ITU-T G.709 is disabled and TIM is enabled, has an SF priority.
- RS-TIM-S on J0 on a trunk, if ITU-T G.709 is disabled and TIM is enabled, has an SF priority.
- CARLOSS on a trunk, if ITU-T G.709 is disabled and the payload is GigE, has an SF priority.
- SIGLOSS on a trunk, if ITU-T G.709 is disabled and the payload is Fibre Channel (any speed), has an SF priority.
- GE-OOSYNC on a trunk, if ITU-T G.709 is disabled and the payload is 10GigE, has an SF priority.
- OOS on a trunk, if ITU-T G.709 is disabled and the payload is GigE, has an SF priority.
- OOS on a trunk, if ITU-T G.709 is disabled and the payload is Fibre Channel, has an SF priority.
- SYNCLOS on a trunk, if ITU-T G.709 is disabled and the payload is Fibre Channel, has an SF priority.

4.10.3.3 Line Conditions on the Client

Most of the defects on client ports are corrected by switching at the far-end protection group.

In a Y-cable protection group, a line alarm indication signal (AIS-L) on the client signal has an SF priority if generic framing procedure (GFP) is not used and the client framing is SONET.

4.10.3.4 Switch Commands

Switch commands have the following priorities:

- A Lockout of Protect (Lockon Working) switch command has an LO priority on the protect port.
- A Force From Working (Force to Protect) switch command has an FS priority on the working port.
- A Force From Protect (Force To Working) switch command has an FS priority on the protect port.
- A Manual From Working (Manual To Protect) switch command has an MS priority on the working port.
- A Manual From Protect (Manual To Working) switch command has an MS priority on the protect port.
- A Clear command cancels (unlocks) any switch command.

4.10.3.5 Unidirectional and Bidirectional Switching

Y-cable and splitter protection support unidirectional switching. In unidirectional switching, the near-end protection group switches without regard for the status of the far-end protection group.

Therefore, the near-end working facilities can be active at the same time the far-end protection facilities are in standby. This does not mean that a defect at the far end will not cause a near-end switch. A defect at the far end might result in a condition at the near end, which then causes a switch, but the switch is caused by the near-end condition.

Bidirectional protection cannot be provisioned unless the near end and far end have the same hardware and data modes and the trunks are connected as working-to-working and protect-to-protect.

Any other configuration results in an undefined behavior. Y-cable and splitter protection groups are not required to detect misconfigured bidirectional protection.

A Y-cable or splitter protection group switches unidirectionally unless at least one trunk signal is intact and there is an operational card terminating the trunk signal at the far end.

4.10.3.6 Other Switch Criteria

This section details switch criteria other than line conditions. A card is said to become operational after it has received and processed the first provisioning message after a warm or cold boot. A card ceases to be operational when it is reset, either with a soft reset request or a hardware reset.

A soft reset of a card does not cause a protection switch or a traffic disruption greater than the disruption induced by the soft reset of an unprotected card. If one or both cards are not operational, then any disruption of traffic on the active card will cause traffic loss until both cards have become operational again.

Switch conditions with a priority lower than FS are ignored by a Y-cable protection group unless both cards are operational. Switch conditions cannot be used to restore traffic while one member of the protection group is reset. For instance, if the working/active card is soft reset and a Forced Switch to the protect card is issued, the protect client laser will turn on but the working client laser will not turn off. Traffic will be lost until the working card becomes operational and can process the Forced Switch request.

The LO and FS switch conditions are accepted by the shelf regardless of the operational status of the cards in a Y-cable protection group or a splitter protection group. A nonoperational card cannot process the switch condition; the provisioning has been accepted by the shelf controller and will be issued to the cards when it becomes operational.

The LO and FS switch conditions are implemented immediately by the operational cards in a Y-cable protection group even if one card is nonoperational. This might cause traffic loss.

Y-cable protection switching is inhibited during a shelf controller reset. Protection switching does not resume until the cards receive their first provisioning message from the active shelf controller.

Each card in a Y-cable protection group begins its provisioning hold-off timer after processing the first provisioning message. For proper behavior, both cards should be provisioned within the provisioning hold-off timer interval. A card missing condition has an SF+ priority. This gives the card missing condition a higher priority than any span alarm.

- A card MEA condition has an SF priority.
- An SFP failure condition has an SF priority.
- An SFP mismatch (failure to support the data rate or mode) condition has an SF priority.
- An SFP missing condition has an SF priority.
- A wavelength mismatch condition has an SF priority.
- A port that is OOS-DSBLD has an SF+ priority.
- The OOS-DSBLD condition has a higher priority than any span alarm.
- A port that is shutdown by ALS has an SF priority.

4.10.3.7 Switch Stability

Y-cable and splitter protection groups use a variety of timers to prevent oscillation, as detailed in the following requirements. No timer is provisionable.

The protection groups implement soak-to-clear timers. A soak-to-clear timer starts when a switch condition clears. While the timer is running, the protection group behaves as though the switch condition is still present. If the switch condition recurs before the timer has expired, the timer is canceled. When the switch condition clears, the timer is restarted.

The durations of soak-to-clear timers are not user-provisionable and are unrelated to the soak times for alarms and conditions. A soak-to-clear timer is not started when a switch condition clears if the switch condition has a lower priority than the currently active switch condition For example, an SD BER soak-to-clear timer will not start if SD BER clears while AIS is present, since AIS has a higher priority than SD BER. All line defects with an SF priority, except for SF BER, share a single one-second soak-to-clear timer. SF BER and SD BER line conditions have a 10-second soak-to-clear timer.

The protection group does not switch sooner than 1.5 seconds after the last switch (the switch hold-off timer). This timer prevents rapid oscillation of the protection group.

A Y-cable protection group does not switch for the first 5 seconds after it is created unless both cards in the protection group become active before 5 seconds elapses.This delay allows both cards in the protection group to be provisioned before any switching decisions are made.

A Y-cable protection group does not switch sooner than 0.5 seconds after a client or trunk facility moves from the OOS-SDBLD state. This hold-off timer allows the cards to ignore transients caused by a port going in-service. The ALS condition has a 3 second soak-to-clear timer.

4.10.3.8 Revertive and Nonrevertive Attributes

Both revertive and nonrevertive switching is supported; the default switch mode is nonrevertive.

Network element (NE) defaults contain a revertive attribute for Y-cable and splitter protection. When applicable in revertive mode, the revert delay timer, also called the Wait-To-Restore (WTR) timer, is software provisionable for Y-cable and splitter protection. The WTR timer is provisionable between 0.5 and 12 minutes, in 0.5 minute increments, and it has a default value of 5 minutes.

When applicable, the NE defaults contain a WTR attribute for Y-cable and splitter protection. When a Y-cable protection group is deleted, a dialog box will appear warning of possible traffic loss.

4.10.3.9 Communications Channels

In a Y-cable protection group, only the working client cam be provisioned with a section data communications channel (SDCC) or line data communications channel (LDCC), and only the working client port can be provisioned as a timing reference (as permitted by payload). The working and protect trunks can be provisioned separately with communication channels (SDCC, LDCC, or generic communications channel [GCC], as permitted by payload type). The communication channels are not protected.

4.10.3.10 Inherited Port Properties

Selected port properties of the protection port are inherited from the working port. In this section, the word port refers to a Y-cable client port or a splitter trunk port:

- The maximum Ethernet Frame Size of the protect port is inherited from the working port.
- The Port Type (SONET or SDH) of the protect port is inherited from the working port.
- The Termination Mode of the protect port is inherited from the working port.
- The SF BER threshold of the protect port is inherited from the working port.
- The SD BER threshold of the protect port is inherited from the working port.
- The SyncMsgIn and SendDoNotUse attributes of the protect port are inherited from the working port.
- Section trace provisioning of the protect port is inherited from the working port.
- The line thresholds of the protect port are inherited from the working port.
- The SDCC/LDCC/GCC provisioning of the protect port is inherited from the working port.
- The ALS provisioning of the protect port is inherited from the working port.

ALS is not permitted on the client ports of Y-cable protection groups. This requirement applies only to splitter protection groups.

4.10.3.11 Switch Status Reporting

Y-cable and splitter protection groups indicate to management software the active/standby status of facilities and cards involved in the protection group. A facility has an active/standby status within the protection group and it has a status that is reported to the management software. These two do not always coincide. Internally, the protection group always has one active facility and one standby facility. In some circumstances, the protection group reports both facilities as standby.

The reported status of any port on a nonoperational card is undefined. While a card is reset, its status might or might not be reported properly. Because the card does not report any status, the report to the user is a function of the management software, not the protection group.

A Y-cable protection group reports a separate status for the TX path and the RX path, for every facility. The active/standby status of the protection group is reported as the status of the RX path. If the status of the far-end protection group is known, then the status of the far-end protection group is reported as the status of the near-end TX path.

The ability of a protection group to know the status of the far-end protection group is a function of the equipment type and the trunk type. If the status of the far-end protection group is not known, the status of the near-end protection group shall be reported as the status of the TX path.

A Y-cable protection group has at most one active client port. A port in a Y-cable or splitter protection group is reported as standby if it has an OOS-DSBLD service state, regardless of its status within the protection group. A port in a Y-cable or splitter protection group is reported as active if it does not have an OOS-DSBLD state and if it carries overhead traffic (GCC, SDCC, LDCC, or E1 bytes), regardless of its status within the protection group.

A client port in a Y-cable protection group is reported as active if it does not have an OOS-DSBLD service state and if it is active within the protection group. A trunk port in a Y-cable protection group is reported as active if it does not have an OOS-DSBLD service state and if any client port on the same card is active.

Transponders have exactly one client port, and the relationship of client to trunk is clear. Muxponder cards have multiple client ports, which means that multiple protection groups are present. If any client port on a muxponder is active, and if the trunk is in-service, the trunk is also reported as active.

Client ports and trunk ports on unprotected cards (cards not part of any protection group) are reported as active if they do not have an OOS-DSBLD service state.

4.10.3.12 Switch Conditions

Protection groups generate conditions and transient conditions to provide a status to the node management software. Common conditions include:

- The protection group raises a MAN-REQ condition against the working facility while a Manual Switch to Protection switch command is in effect.
- The protection group raises a MAN-REQ condition against the protection facility while a Manual Switch to Working switch command is in effect.
- The protection group raises a FORCED-REQ condition against the working facility while a Forced Switch to Protection switch command is in effect.
- The protection group raises a FORCED-REQ condition against the protection facility while a Forced Switch to Working switch command is in effect.
- The protection group raises a LOCKOUT-REQ condition against the protection facility while a Lockout of Protection switch command is in effect.
- The protection group signals an APS-CLEAR condition when a switch command is preempted by a higher-priority switch condition.
- The protection group signals a FAILTOSW condition while a switch command is inhibiting a protection switch due to a lower-priority line condition.
- The protection group raises a WTR condition against the working facility while the Wait To Restore timer is running.
- The protection group, if it is in revertive mode, raises a WKSWPR condition against the working facility while the protection facility is active.
- The protection group, if it is in nonrevertive mode, signals a WKSWPR condition against the working facility when the protection facility becomes active.

- The protection group, if it is nonrevertive, signals a WKSWBK condition against the working facility when the working facility becomes active.

4.10.3.13 Protection Switching Performance Requirements

Protection switching is executed within 50 ms of a defect appearing at the near end. Loss of light on the client outputs of a Y-cable protection group does not exceed 20 ms during a switch.

During a protection switch, the standby client transmitter turns off, and the active client transmitter turns on. If the standby transmitter turns off before the active transmitter is fully on, a loss of light occurs at the stem of the Y-cable. This loss of light does not last longer than 20 ms. If a payload cannot tolerate a loss of light less than 20 ms, then that payload cannot be used with Y-cable protection.

4.10.4 Usability Requirements

The following section discusses regeneration groups, automatic laser shutdown, and client signal failures.

4.10.4.1 Regeneration Groups

A regeneration group boosts the power and improves the signal-to-noise (S/N) ratio in a DWDM signal. The purpose is to extend the reach of a DWDM signal between two termination points. In an ideal condition, regeneration is totally transparent to the endpoints. However, some regeneration techniques fall short of this ideal condition and might modify, delay, or even drop overhead signals (ITU-T G.709, GFP, or other section-level signaling protocols).

The behavior of Y-cable and splitter protection groups is unchanged by the presence of a single peer-to-peer regeneration group in one or both of the DWDM spans. This requirement cannot be met if the regeneration corrupts the overhead bytes that are necessary for protection switching.

4.10.4.2 Automatic Laser Shutdown

ALS disables the transmitter of a facility if the receiver of the same facility detects a loss of light. ALS exists as a human-safety standard. After ALS shuts down the transmitter, it is not restarted until the loss of light condition clears. To facilitate restarting lasers when both ends of a span are shut down by ALS, the facility can be provisioned to send short test pulses of light. ALS is not permitted on the client ports of a Y-cable protection group.

4.10.4.3 Client Signal Failures

Y-cable protection groups can protect against failures of the client RX signal at the end. The far-end client RX failure can be in the fiber (in an “arm” of the Y), in the equipment (for example, the SFP), or in the provisioning (client OOS-DSBLD). These failure types require special handling because they are out-of-band with respect to the normal client payload. The term used for these failures is client signal fail (CSF). This has the same meaning as GFP-CSF, but does not imply that GFP-CSF is used for the signaling. Client signal failures include:

- An S-LOF on a client port, if the client is OC-N, is signaled to the downstream client port as a CSF.
- An S-LOS on a client port, if the client is OC-N, is signaled to the downstream client port as a CSF.

- An RS-LOF on a client port, if the client is STM-N, is signaled to the downstream client port as a CSF.
- An RS-LOS on a client port, if the client is STM-N, is signaled to the downstream client port as a CSF.
- A GE-OOSYNC on a client port, if the client is GigE or 10GigE, is signaled to the downstream client port as a CSF.
- An SFP missing condition on a client port is signaled to the downstream silent port as a CSF.
- An SFP mismatch (failure to support client data rate) condition on a client port is signaled to the downstream client port as a CSF.
- An SFP failure condition on a client port is signaled to the downstream client port as a CSF.
- An OOS-DSBLD condition on a client port is signaled to the downstream client port as a CSF.

4.10.5 In-service Upgrade

A Y-cable group switches normally during a software activation or software revert if both cards in the Y-cable group are running the same software release.

This behavior is different from that of OC-N 1+1 protection groups, which will not switch until the software activation is complete. The Y-cable group is able to switch before either card has booted to the new release and after both cards have booted to the new release. This requirement does not preempt the other requirements that both cards be operational and that an active TCC2/TCC2P be installed.

4.11 Path Diversion Support for Client Protection

The ONS 15454 DWDM system provides the capability to provision unprotected wavelengths on a per-wavelength basis and supports the reuse of unprotected wavelengths on adjacent spans.

Figure 4-21 provides examples of unprotected wavelengths.

Figure 4-21 Example of Unprotected Wavelengths