CHAPTER 10

Circuits and Tunnels

This chapter explains Cisco ONS 15454 SDH high-order and low-order circuits; low-order, data communication channel (DCC), and IP-encapsulated tunnels; and virtual concatenated (VCAT) circuits. To provision circuits and tunnels, refer to the *Cisco ONS 15454 SDH Procedure Guide*.

Chapter topics include:

10.1 Overview

You can create circuits across and within ONS 15454 SDH nodes and assign different attributes to circuits. For example, you can:

- Create one-way, two-way (bidirectional), or broadcast circuits. VC low-order path tunnels (VC_LO_PATH_TUNNEL) are automatically set to bidirectional and do not use multiple drops.
- Assign user-defined names to circuits.
- Assign different circuit sizes.
- Enable port grouping on low-order path tunnels. Three ports form a port group. For example, in one E3-12 or one DS3I-N-12 card, four port groups are available: Ports 1 to 3 = PG1, Ports 4 to 6 = PG2, Ports 7 to 9 = PG3, and Ports 10 to 12 = PG4.

Note CTC shows VC3-level port groups, but the XC10G creates only VC4-level port groups. VC4 tunnels must be used to transport VC3 signal rates.

Note Monitor circuits cannot be created on a VC3 circuit in a port group.

- Automatically or manually route VC high-order and low-order path circuits.
- Automatically route VC low-order path tunnels.
- Automatically create multiple circuits with autoranging. VC low-order path tunnels do not use autoranging.
- Provide full protection to the circuit path.
- Provide only protected sources and destinations for circuits.
- Define a secondary circuit source or destination that allows you to interoperate an ONS 15454 SDH subnetwork connection protection (SNCP) ring with third-party equipment SNCPs.

You can provision circuits at any of the following points:

- Before cards are installed. The ONS 15454 SDH allows you to provision slots and circuits before installing the traffic cards. (To provision an empty slot, right-click it and select a card from the shortcut menu.) However, circuits cannot carry traffic until you install the cards and place their ports in service. For card installation procedures and ring-related procedures, refer to the *Cisco ONS 15454 SDH Procedure Guide*.
- After cards are installed, but before their ports are in service (enabled). You must put the ports in service before circuits can carry traffic.
- After cards are installed and their ports are in service. Circuits carry traffic as soon as the signal is received.

10.2 Circuit Properties

The ONS 15454 SDH Circuits window, which appears in network, node, and card view, is where you can view information about circuits. The Circuits window (Figure 10-1) provides the following information:

- Name—The name of the circuit. The circuit name can be manually assigned or automatically generated.
- Type—Circuit types are HOP (high-order circuit), LOP (low-order circuit), VCT (VC low-order tunnel), VCA (VC low-order aggregation point), OCHNC (dense wavelength division multiplexing [DWDM] optical channel network connection), HOP_v (high-order virtual concatenated [VCAT] circuit), or LOP_v (low-order VCAT circuit).
- Size—The circuit size. Low-order circuits are VC12 and VC3. High-order circuit sizes are VC4, VC4-3c, VC4-4c, VC4-8c, VC4-16c, VC4-64c. OCHNC sizes are Equipped not specific, Multi-rate, 2.5 Gbps No FEC (forward error correction), 2.5 Gbps FEC, 10 Gbps No FEC, and 10 Gbps FEC. High-order VCAT circuits are VC4-2v and VC4-4c-2v. Low-order VCAT circuits are VC3-2v.
- OCHNC Wlen—For OCHNCs, the wavelength provisioned for the DWDM optical channel network connection.
- Direction—The circuit direction, either two-way (bidirectional) or one-way.

- OCHNC Dir—For OCHNCs, the direction of the DWDM optical channel network connection, either east to west or west to east.
- Protection—The type of circuit protection. See the "10.2.3 Circuit Protection Types" section on page 10-6.
- Status—The circuit status. See the "10.2.1 Circuit Status" section on page 10-4.
- Source—The circuit source in the format: node/slot/port "port name" virtual container/tributary unit group/tributary unit group/virtual container. (The port name appears in quotes.) Node and slot always display; port "port name"/virtual container/tributary unit group/tributary unit group/virtual container might display, depending on the source card, circuit type, and whether a name is assigned to the port. If the circuit size is a concatenated size (VC4-2c, VC4-4c, VC4-8c, etc.) VCs used in the circuit are indicated by an ellipsis, for example, VC4-7..9 (VCs 7, 8, and 9) or VC4-10..12 (VC 10, 11, and 12).
- Destination—The circuit destination in same format (node/slot/port "port name"/virtual container/tributary unit group/tributary unit group/virtual container) as the circuit source.
- # of VLANS—The number of VLANS used by an Ethernet circuit with end points on E-Series Ethernet cards in single-card or multicard mode.
- # of Spans—The number of inter-node links that constitute the circuit. Right-clicking the column displays a shortcut menu from which you can choose to show or hide circuit span detail.
- State—The circuit state. See the "10.2.2 Circuit States" section on page 10-4.

Figure 10-1 ONS 15454 SDH Circuit Window in Network View

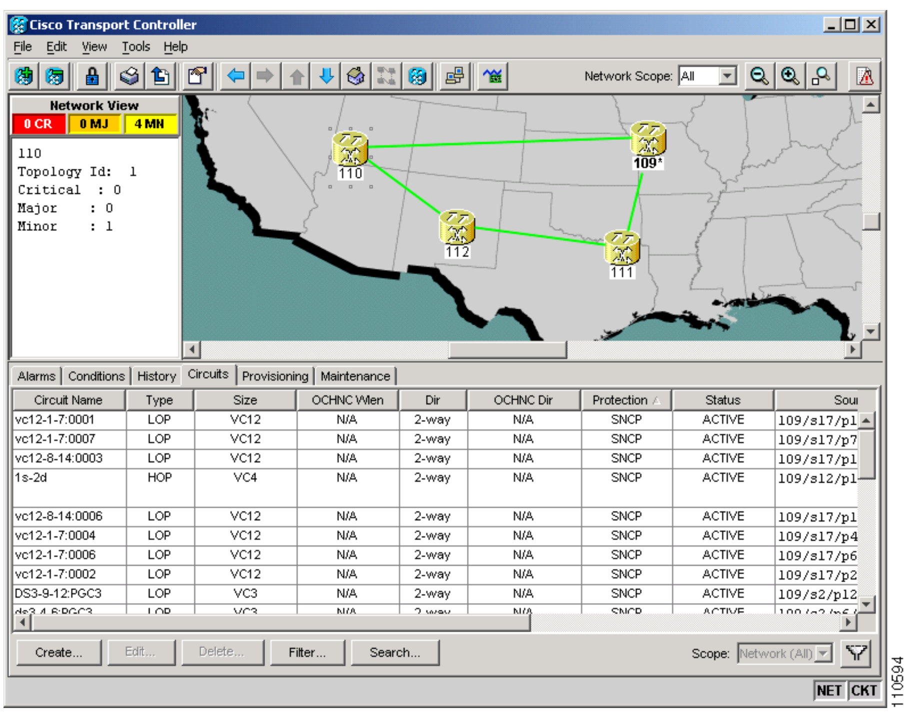

10.2.1 Circuit Status

The circuit statuses that appear in the Circuit window Status column are generated by CTC based on conditions along the circuit path. Table 10-1 shows the statuses that can appear in the Status column.

Table 10-1 Cisco ONS 15454 SDH Circuit Status

Status	Definition/Activity
CREATING	CTC is creating a circuit.
ACTIVE	CTC created a circuit. All components are in place and a complete path exists from the circuit source to the circuit destination.
DELETING	CTC is deleting a circuit.
INCOMPLETE	A CTC-created circuit is missing a cross-connect or network span, a complete path from source to destination(s) does not exist. In CTC, circuits are represented using cross-connects and network spans. If a network span is missing from a circuit, the circuit status is INCOMPLETE. However, an INCOMPLETE status does not necessarily mean a circuit traffic failure has occurred, because traffic may flow on a protect path. Network spans are in one of two states: up or down. On CTC circuit and network maps, up spans appear as green lines, and down spans appear as gray lines. If a failure occurs on a network span during a CTC session, the span remains on the network map but its color changes to gray to indicate that the span is down. If you restart your CTC session while the failure is active, the new CTC session cannot discover the span and its span line does not appear on the network map. Subsequently, circuits routed on a network span that goes down display as ACTIVE during the current CTC session, but they display as INCOMPLETE to users who log in after the span failure. The INCOMPLETE status does not apply to OCHNC circuit types.

10.2.2 Circuit States

State is a user-assigned designation that indicates whether the circuit should be in service or out of service. Table 10-2 lists the states that you can assign to circuits. To carry traffic, circuits must have a status of Active and a state of In Service (IS), Out of Service, Auto In Service (OOS-AINS), or Out of Service Maintenance (OOS-MT). The circuit source port and destination port must also be IS, OOS-AINS, or OOS-MT.

Note OOS-AINS and OOS-MT allow a signal to be carried, although alarm reporting is suppressed.

You can assign a state to circuits at two points:

- During circuit creation, you can assign a state to the circuit on the Create Circuit wizard.
- After circuit creation, you can change a circuit state on the Edit Circuit window or from the Tools > Circuits > Set Circuit State menu.

Table 10-2 Cisco ONS 15454 SDH Circuit States

State	Definition
IS	In Service; able to carry traffic.
OOS	Out of Service; unable to carry traffic. This state does not apply to OCHNC circuit types.
OOS-AINS	Out of Service, Auto In Service; alarm reporting is suppressed, but traffic is carried and loopbacks are allowed. Raised fault conditions, whether their alarms are reported or not, can be retrieved on the CTC Conditions tab or by using the TL1 RTRV-COND command. Low-order circuits in OOS-AINS generally switch to IS when source and destination ports are IS, OOS-AINS, or OOS-MT regardless of whether a physical signal is present. High-order circuits in OOS-AINS switch to IS when a signal is received. This state does not apply to OCHNC circuit types.
OOS-MT	Out of Service, Maintenance; alarm reporting is suppressed, but traffic is carried and loopbacks are allowed. Raised fault conditions, whether or not their alarms are reported, can be retrieved on the CTC Conditions tab or by using the TL1 RTRV-COND command. This state does not apply to OCHNC circuit types.

PARTIAL is appended to a circuit state whenever all circuit cross-connects are not in the same state. Table 10-3 shows the partial circuit states that can appear. Partial circuit states do not apply to OCHNC circuit types.

Table 10-3 Partial ONS 15454 SDH Circuit States

State	Definition
OOS-PARTIAL	At least one connection is OOS and at least one other is in a different state.
OOS-AINS-PARTIAL	At least one connection is OOS-AINS and at least one other is in IS state.
OOS-MT-PARTIAL	At least one connection is OOS-MT and at least one other is in a different state (other than OOS).

PARTIAL states can occur during automatic or manual transitions between states. PARTIAL can appear during a manual transition caused by an abnormal event such as a CTC crash, communication error, or if one of the cross-connects could not be changed. Refer to the *Cisco ONS 15454 SDH Troubleshooting Guide* for troubleshooting procedures.

Note Circuits that are assigned to a state other than OOS and that use E1, E3-12, or DS3I-N-12 cards for the source and destination ports can change to IS even though a signal is not present on the ports. Some cross-connects transition to IS, while others are OOS-AINS. The PARTIAL state might appear during a manual transition.

Circuits do not use the soak timer for transitional states, but ports do. The soak period is the amount of time that the port remains in the OOS-AINS state after a signal is continuously received. When provisioned as OOS-AINS, the ONS 15454 SDH monitors a circuit's cross-connects for an error-free signal. It changes the state of the circuit from OOS-AINS to IS or to OOS-AINS-PARTIAL as each cross-connect assigned to the circuit path is completed. Two common examples of state changes you see when provisioning circuits using CTC are as follows:

- When provisioning low-order circuits and tunnels as OOS-AINS, the circuit state transitions to IS shortly after the circuits are created when the circuit source and destination ports are IS, OOS-AINS, or OOS-MT. The source and destination ports on the low-order circuits remain in the OOS-AINS state until an alarm-free signal is received for the duration of the soak timer. When the soak timer expires, the low-order source port and destination port states change to IS.
- When provisioning high-order circuits as OOS-AINS, the circuit source and destination ports are OOS-AINS. As soon as an alarm-free signal is received, the circuit state changes to IS and the source and destination ports remain OOS-AINS for the duration of the soak timer. After the port soak timer expires, high-order source and destination ports change to IS.

To find the remaining port OOS-AINS soak time, choose the Maintenance > AINS Soak tabs in card view and click the Retrieve button. If the port is in the OOS-AINS state and has a good signal, the Time Until IS column shows the soak count down status. If the port is OOS-AINS and has a bad signal, the Time Until IS column indicates that the signal is bad. You must click the AINS Soak tab Retrieve button to obtain the latest time value.

10.2.3 Circuit Protection Types

The Protection column on the Circuit window shows the card (line) and SDH topology (path) protection used for the entire circuit path. Table 10-4 shows the protection type indicators that appear in this column.

Table 10-4 Circuit Protection Types

Protection Type	Description
N/A	Circuit protection is not applicable.
2F MS-SPR	The circuit is protected by a two-fiber multiplex section-shared protection rings (MS-SPRings).
4F MS-SPR	The circuit is protected by a four-fiber MS-SPRing.
MS-SPR	The circuit is protected by both a two-fiber and a four-fiber MS-SPRing.
SNCP	The circuit is protected by an SNCP.
SNCP-DRI	The circuit is protected by an SNCP dual-ring interconnection.
1+1	The circuit is protected by a 1+1 protection group.
Y-Cable	The circuit is protected by a transponder or muxponder card Y-cable protection group.
Splitter	The circuit is protected by the protect transponder (TXPP_MR_2.5G) splitter protection.
Protected	The circuit is protected by diverse SDH topologies, for example, an MS-SPRing and an SNCP, or an SNCP and a 1+1 protection group.
2F-PCA	The circuit is routed on a protection channel access (PCA) path on a two-fiber MS-SPRing; PCA circuits are unprotected.

Table 10-4 Circuit Protection Types (continued)

Protection Type	Description
4F-PCA	The circuit is routed on a protection channel access path on a four-fiber MS-SPRing; PCA circuits are unprotected.
PCA	The circuit is routed on a protection channel access path on both two-fiber and four-fiber MS-SPRings; PCA circuits are unprotected.
Unprot (black)	The circuit is not protected.
Unprot (red)	The circuit created as a fully protected circuit is no longer protected due to a system change, such as removal of a MS-SPRing or 1+1 protection group.
Unknown	The circuit protection types appear in the Protection column only when all circuit components are known, that is, when the circuit status is ACTIVE or UPGRADABLE. If the circuit is in some other status, the protection type is "unknown."

10.2.4 Circuit Information in the Edit Circuit Window

The detailed circuit map on the Edit Circuit window allows you to view information about ONS 15454 SDH circuits. Routing information that appears includes:

- Circuit direction (unidirectional/bidirectional)
- The nodes, VC4s, TUG3, TUG2s, and VC12s through which the circuit passes, including slots and port numbers
- The circuit source and destination points
- OSPF Area IDs
- Link protection (SNCP, unprotected, MS-SPRing, 1+1) and bandwidth (STM-N)

For MS-SPRings, the detailed map shows the number of MS-SPRing fibers and the MS-SPRing ring ID. For SNCPs, the map shows the active and standby paths from circuit source to destination, and it also shows the working and protect paths. For VCAT circuits, the detailed map is not available for an entire VCAT circuit. However, you can view the detailed map to view the circuit route for each individual member.

You can also view alarms and states can also be viewed on the circuit map, including:

- Alarm states of nodes on the circuit route
- Number of alarms on each node organized by severity
- Port service states on the circuit route
- Alarm state/color of the most severe alarm on the port
- Loopbacks
- Path trace states
- Path selectors states

For example, in an SNCP, the working path is indicated by a green, bidirectional arrow, and the protect path is indicated by a purple, bidirectional arrow. Source and destination ports are shown as circles with an S and a D. Port states are indicated by colors, shown in Table 10-5.

Table 10-5 Port State Color Indicators

Port Color	State
Green	IS
Gray	OOS
Purple	OOS-AINS
Cyan (Blue)	OOS-MT

A notation within or by the squares on each node indicates switches and loopbacks, including:

- F = Force switch
- M = Manual switch
- L = Lockout switch
- T = Terminal loopback
- Arrow = Facility loopback

Move the mouse cursor over nodes, ports, and spans to see tooltips with information including the number of alarms on a node (organized by severity), port state of service (that is, IS, OOS, etc.), and the protection topology.

Right-click a node, port, or span on the detailed circuit map to initiate certain circuit actions:

- Right-click a unidirectional circuit destination node to add a drop to the circuit.
- Right-click a port containing a path trace capable card to initiate the path trace.
- Right-click an SNCP span to change the state of the path selectors in the SNCP circuit.

10.3 Cross-Connect Card Bandwidth

The XC10G is required to operate the ONS 15454 SDH. XC10Gs support high-order cross-connections (VC4 and above at STM-1, STM-4, STM-16, and STM-64 signal rates). The XC10G does not support any low-order circuits such as VC-11, VC-12, and VC3. The XC10G card cross connects standard VC4, VC4-4c, VC4-16c, and VC4-64c signal rates and nonstandard VC4-2c, VC4-3c, and VC4-8c signal rates, providing a maximum of 384 x 384 VC4 cross-connections. Any VC4 on any port can be connected to any other port, meaning that the VC cross-connection capacity is nonblocking. The XC10G card manages up to 192 bidirectional VC4 cross-connects.

VC4 tunnels must be used with the E3-12 and DS3i-N-12 cards to transport VC3 signal rates. Three ports form a port group. For example, in one E3-12 or one DS3i-N-12 card, there are four port groups: Ports 1 to 3 = PG1, Ports 4 to 6 = PG2, Ports 7 to 9 = PG3, and Ports 10 to 12 = PG4.

Note In SDH Software R3.4 and earlier, the XC10G does not support VC3 circuits for the E3-12 and DS3i-N-12 cards. You must create a VC tunnel. See the *Cisco ONS 15454 SDH Procedure Guide* for more information.

The XC-VXL-10G and XC-VXL-2.5G card support both low-order and high-order circuits (E-1, E-3, DS-3, STM-1, STM-4, STM-16, and STM-64 signal rates). They manage up to 192 bidirectional STM-1 cross-connects, 192 bidirectional E-3 or DS-3 cross-connects, or 1008 bidirectional E-1 cross-connects.

The XC10G, XC-VXL-10G, and XC-VXL-2.5G cards work with the TCC2 card to maintain connections and set up cross-connects within the node. You can create circuits using CTC.

Note Chapter 2, "Common Control Cards," contains detailed specifications of the XC10G, XC-VXL-10G, and XC-VXL-2.5G cards.

10.4 DCC Tunnels

SDH provides four DCCs for network element operations, administration, maintenance, and provisioning: one on the SDH regenerator section layer (SDCC) and three on the SDH multiplex section layer, also called the line DCC (LDCC). A section DCC (SDCC) and line DCC (LDCC) each provide 192 Kbps of bandwidth per channel. The aggregate bandwidth of the three LDCCs is 576 Kbps. When multiple DCC channels exist between two neighboring nodes, the ONS 15454 SDH balances traffic over the existing DCC channels.

You can tunnel third-party SDH equipment across ONS 15454 SDH networks using one of two tunneling methods, a traditional DCC tunnel or an IP-encapsulated tunnel.

10.4.1 Traditional DCC Tunnels

In traditional DCC tunnels, the ONS 15454 SDH uses regenerator section DCC for inter-ONS-15454-SDH data communications. It does not use the multiplex section DCCs; therefore, the multiplex section DCCs are available to tunnel DCCs from third-party equipment across ONS 15454 SDH networks. If D4 through D12 are used as data DCCs, they cannot be used for DCC tunneling.

A traditional DCC tunnel endpoint is defined by slot, port, and DCC, where DCC can be either the regenerator section DCC, Tunnel 1, Tunnel 2, or Tunnel 3. You can link a regenerator section DCC to a multiplex section DCC (Tunnel 1, Tunnel 2, or Tunnel 3) and a multiplex section DCC to a regenerator section DCC. You can also link multiplex section DCCs to multiplex section DCCs and link regenerator section DCCs to regenerator section DCCs. To create a DCC tunnel, you connect the tunnel end points from one ONS 15454 SDH STM-N port to another.

Cisco recommends a maximum of 84 DCC tunnel connections for an ONS 15454 SDH. Table 10-6 shows the DCC tunnels that you can create.

Table 10-6 DCC Tunnels

DCC	SDH Layer	SDH Bytes	STM-1 (All Ports)	STM-4, STM-16, STM-64
SDCC	Regenerator Section	D1 to D3	Yes	Yes
Tunnel 1	Multiplex Section	D4 to D6	No	Yes
Tunnel 2	Multiplex Section	D7 to D9	No	Yes
Tunnel 3	Multiplex Section	D10 to D12	No	Yes

Figure 10-2 shows a DCC tunnel example. Third-party equipment is connected to STM-1 cards at Node 1/Slot 3/Port 1 and Node 3/Slot 3/Port 1. Each ONS 15454 SDH node is connected by STM-16 trunk (span) cards. In the example, three tunnel connections are created, one at Node 1 (STM-1 to STM-16), one at Node 2 (STM-16 to STM-16), and one at Node 3 (STM-16 to STM-1).

Note A DCC does not function on a mixed network of ONS 15454 SDH nodes and ONS 15454 nodes. DCC tunneling is required for ONS 15454 SDH nodes transporting data through ONS 15454 nodes.

Figure 10-2 DCC Tunnel

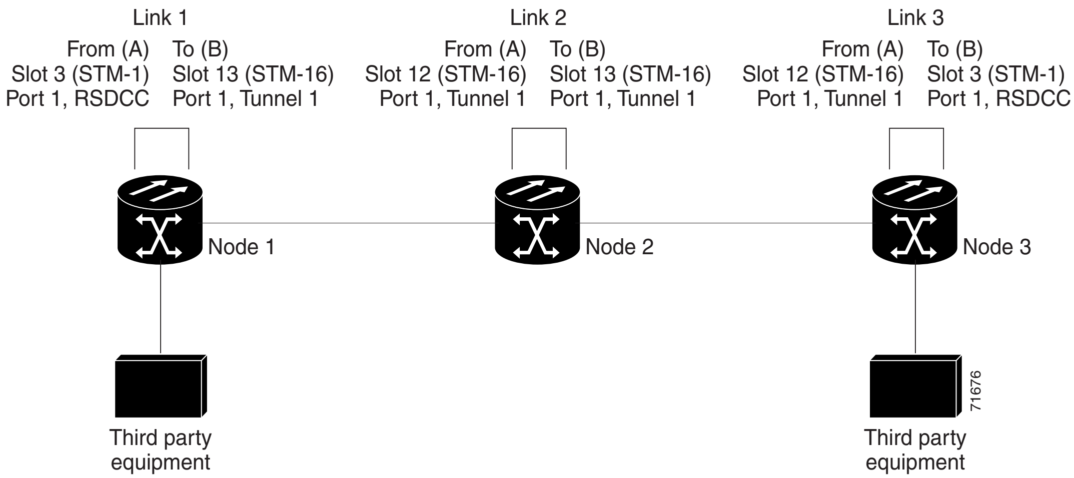

When you create DCC tunnels, keep the following guidelines in mind:

- Each ONS 15454 SDH can have up to 84 DCC tunnel connections.
- Each ONS 15454 SDH can have up to 84 regenerator section DCC terminations.
- A regenerator section DCC that is terminated cannot be used as a DCC tunnel endpoint.
- A regenerator section DCC that is used as a DCC tunnel endpoint cannot be terminated.
- All DCC tunnel connections are bidirectional.

Note A multiplex section DCC cannot be used for tunneling if a data DCC is assigned.

10.4.2 IP-Encapsulated Tunnels

An IP-encapsulated tunnel puts an SDCC in an IP packet at a source node and dynamically routes the packet to a destination node. A traditional DCC tunnel is configured as one dedicated path across a network and does not provide a failure recovery mechanism if the path is down. An IP-encapsulated tunnel is a virtual path, which adds protection when traffic travels between different networks.

IP-encapsulated tunneling has the potential of flooding the DCC network with traffic resulting in a degradation of performance for CTC. The data originating from an IP tunnel can be throttled to a user-specified rate, which is a percentage of the total SDCC bandwidth.

Each ONS 15454 SDH supports up to ten IP-encapsulated tunnels. You can convert a traditional DCC tunnel to an IP-encapsulated tunnel or an IP-encapsulated tunnel to a traditional DCC tunnel. Only tunnels in the Active state can be converted.

Caution Converting from one tunnel type to the other is service-affecting.

10.5 Multiple Destinations for Unidirectional Circuits

Unidirectional circuits can have multiple destinations for use in broadcast circuit schemes. In broadcast scenarios, one source transmits traffic to multiple destinations, but traffic is not returned back to the source.

When you create a unidirectional circuit, the card that does not have its backplane receive (Rx) input terminated with a valid input signal generates a loss of signal (LOS) alarm. To mask the alarm, create an alarm profile suppressing the LOS alarm and apply it to the port that does not have its Rx input terminated.

10.6 Monitor Circuits

Monitor circuits are secondary circuits that monitor traffic on primary bidirectional circuits. Monitor circuits can be created on E1 or STM-N cards. Figure 10-3 shows an example of a monitor circuit. At Node 1, a VC4 is dropped from Port 1 of an STM-1 card. To monitor the VC4 traffic, test equipment is plugged into Port 2 of the STM-1 card and a monitor circuit to Port 2 is provisioned in CTC. Circuit monitors are one-way. The monitor circuit in Figure 10-3 is used to monitor VC4 traffic received by Port 1 of the STM-1 card.

Figure 10-3 VC4 Monitor Circuit Received at an STM-1 Port

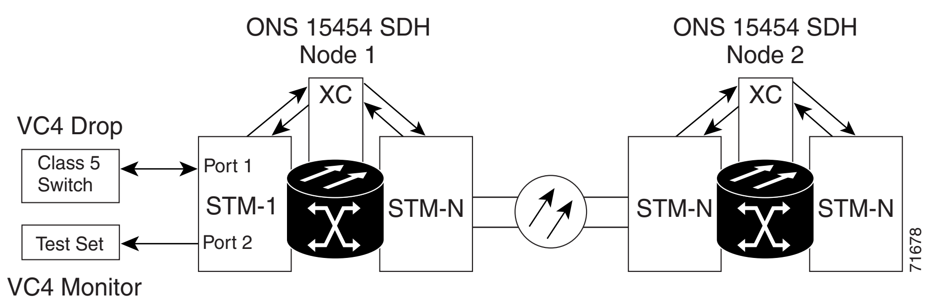

Note Monitor circuits cannot be used with Ethernet circuits.

10.7 SNCP Circuits

Use the Edit Circuits window to change SNCP selectors and switch protection paths (Figure 10-4). In this window, you can:

- View the SNCP circuit's working and protection paths.
- Edit the reversion time.
- Edit the Signal Fail/Signal Degrade thresholds.
- Change PDI-P settings.
- Perform maintenance switches on the circuit selector.
- View switch counts for the selectors.

Figure 10-4 Editing SNCP Selectors

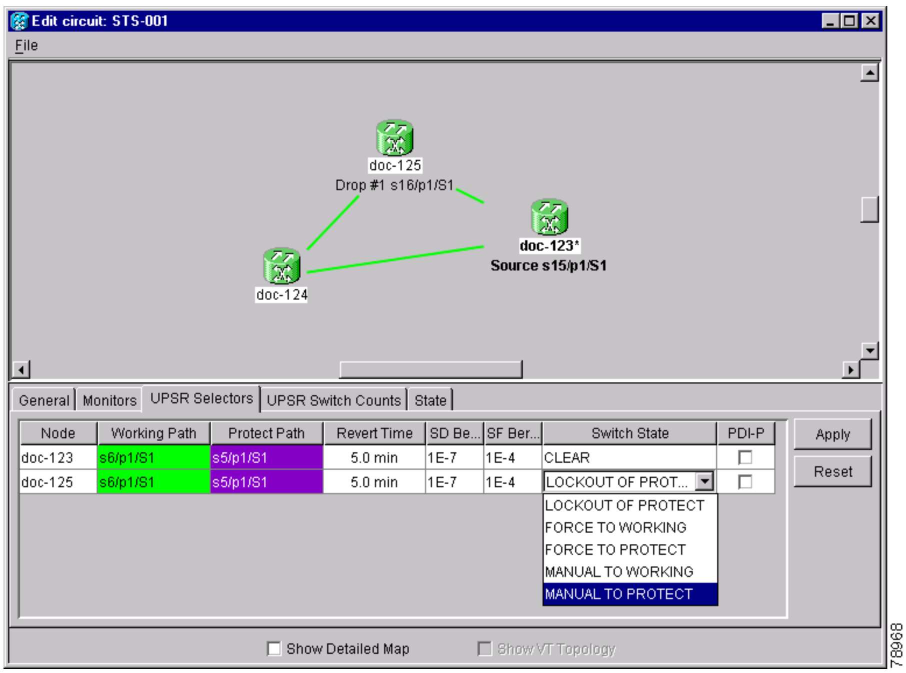

10.7.1 Open-Ended SNCP Circuits

If ONS 15454 SDHs are connected to a third-party network, you can create an open-ended SNCP circuit to route a circuit through it. To do this, you create three circuits. One circuit is created on the source ONS 15454 SDH network. This circuit has one source and two destinations, one at each ONS 15454 SDH that is connected to the third-party network. The second circuit is created on the third-party network so that the circuit travels across the network on two paths to the ONS 15454 SDHs. That circuit routes the two circuit signals across the network to ONS 15454 SDHs that are connected to the network on other side. At the destination node network, the third circuit is created with two sources, one at each node connected to the third-party network. A selector at the destination node chooses between the two signals that arrive at the node, similar to a regular SNCP circuit.

10.7.2 Go-and-Return SNCP Routing

The go-and-return SNCP routing option allows you to route the SNCP working path on one fiber pair and the protect path on a separate fiber pair (Figure 10-5). The working path will always be the shortest path. If a fault occurs, both the working and protection fibers are not affected. This feature only applies to bidirectional SNCP circuits. The go-and-return option appears on the Circuit Attributes panel of the Circuit Creation wizard.

Figure 10-5 SNCP Go-and-Return Routing

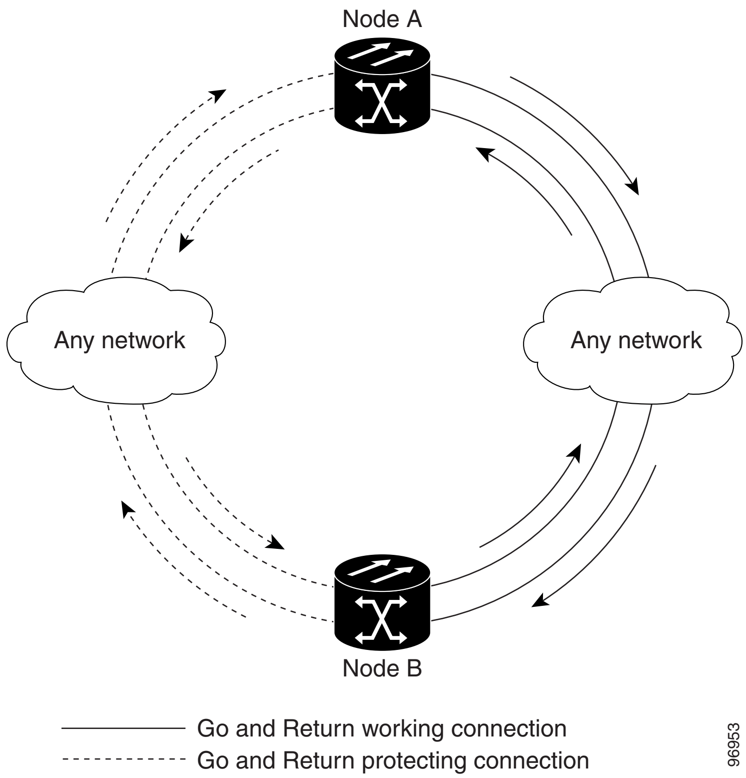

10.8 MS-SPRing Protection Channel Access Circuits

You can provision circuits to carry traffic on MS-SPRing protection channels when conditions are fault free. Traffic routed on MS-SPRing PCA circuits, called extra traffic, has lower priority than the traffic on the working channels and has no means for protection. During ring or span switches, PCA circuits are preempted and squelched. For example, in a two-fiber STM-16 MS-SPRing, STMs 9 to 16 can carry extra traffic when no ring switches are active, but PCA circuits on these STMs are preempted when a ring switch occurs. When the conditions that caused the ring switch are remedied and the ring switch is removed, PCA circuits are restored if the MS-SPRing is provisioned as revertive.

Provisioning traffic on MS-SPRing protection channels is performed during circuit provisioning. The Protection Channel Access check box appears whenever Fully Protected Path is unchecked on the circuit creation wizard. Refer to the *Cisco ONS 15454 SDH Procedure Guide* for more information. When provisioning PCA circuits, two considerations are important to keep in mind:

- If MS-SPRings are provisioned as nonrevertive, PCA circuits are not restored automatically after a ring or span switch. You must switch the MS-SPRing manually.
- PCA circuits are routed on working channels when you upgrade a MS-SPRing from a two-fiber to a four-fiber or from one STM-N speed to a higher STM-N speed. For example, if you upgrade a two-fiber STM-16 MS-SPRing to an STM-64, STMs 9 to 16 on the STM-16 MS-SPRing become working channels on the STM-64 MS-SPRing.

10.9 J1 Path Trace

The J1 Path Trace is a repeated, fixed-length string comprised of 64 consecutive J1 bytes. You can use the string to monitor interruptions or changes to circuit traffic. Table 10-7 shows the ONS 15454 SDH cards that support path trace. Cards that are not listed in the table do not support the J1 byte.

Table 10-7 ONS 15454 SDH Cards Capable of Path Trace

J1 Function	Cards
Transmit and receive	E3-12
	DS3I-N-12
	G-Series
	ML-Series
Receive only	OC3 IR 4/STM1 SH 1310
	OC12/STM4-4
	OC48 IR/STM16 SH AS 1310
	OC48 LR/STM16 LH AS 1550
	OC192 LR/STM64 LH 1550

The J1 path trace transmits a repeated, fixed-length string. If the string received at a circuit drop port does not match the string the port expects to receive, an alarm is raised. Two path trace modes are available:

- Automatic—The receiving port assumes that the first J1 string it receives is the baseline J1 string.
- Manual—The receiving port uses a string that you manually enter as the baseline J1 string.

10.10 Path Signal Label, C2 Byte

One of the overhead bytes in the SDH frame is the C2 byte. The SDH standard defines the C2 byte as the path signal label. The purpose of this byte is to communicate the payload type being encapsulated by the high-order path overhead (HO-POH). The C2 byte functions similarly to EtherType and Logical Link Control (LLC)/Subnetwork Access Protocol (SNAP) header fields on an Ethernet network; it allows a single interface to transport multiple payload types simultaneously. Table 10-8 provides the C2 byte hex values.

Table 10-8 STM Path Signal Label Assignments for Signals

Hex Code	Content of the STM SPE
0x00	Unequipped
0x01	Equipped—nonspecific payload
0x02	Tributary unit group (TUG) structure
0x03	Locked tributary unit (TU-n)
0x04	Asynchronous mapping of 34,368 kbps or 44,736 kbps into the container-3 (C-3)
0x12	Asynchronous mapping of 139,264 kbps into the container-4 (C-4)
0x13	Mapping for Asynchronous Transfer Mode (ATM)
0x14	Mapping for DQDB
0x15	Asynchronous mapping for Fiber Distributed Data Interface (FDDI)
0xFE	0.181 Test signal (TSS1 to TSS3) mapping SDH network (see ITU-T G.707)
0xFF	Virtual container-alarm indications signal (VC-AIS)

If a circuit is provisioned using a terminating card, the terminating card provides the C2 byte. A low-order path circuit is terminated at the XC10G or XC-VXL and the cross-connect card generates the C2 byte (0x02) downstream to the VC terminating cards. The cross-connect generates the C2 value (0x02) to the terminating card. If an STM-N circuit is created with no terminating cards, the test equipment must supply the path overhead in terminating mode. If the test equipment is in “pass through mode,” the C2 values usually change rapidly between 0x00 and 0xFF. Adding a terminating card to an STM-N circuit usually fixes a circuit having C2 byte problems.

10.11 Automatic Circuit Routing

If you select automatic routing during circuit creation, CTC routes the circuit by dividing the entire circuit route into segments based on protection domains. For unprotected segments of circuits provisioned as fully protected, CTC finds an alternate route to protect the segment, creating a virtual SNCP. Each segment of a circuit path is a separate protection domain. Each protection domain is protected in a specific protection scheme including card protection (1+1, 1:1, etc.) or SDH topology (SNCP, MS-SPRing, etc.).

The following list provides principles and characteristics of automatic circuit routing:

- Circuit routing tries to use the shortest path within the user-specified or network-specified constraints. Low-order tunnels are preferable for low-order circuits because low-order tunnels are considered shortcuts when CTC calculates a circuit path in path-protected mesh networks.

- If you do not choose Fully Path Protected during circuit creation, circuits can still contain protected segments. Because circuit routing always selects the shortest path, one or more links and/or segments can have some protection. CTC does not look at link protection while computing a path for unprotected circuits.
- Circuit routing does not use links that are down. If you want all links to be considered for routing, do not create circuits when a link is down.
- Circuit routing computes the shortest path when you add a new drop to an existing circuit. It tries to find the shortest path from the new drop to any nodes on the existing circuit.
- If the network has a mixture of low-order-capable nodes and low-order-incapable nodes, CTC might automatically create a low-order tunnel. Otherwise, CTC asks you whether or not a low-order tunnel is needed.

10.11.1 Bandwidth Allocation and Routing

Within a given network, CTC routes circuits on the shortest possible path between source and destination based on the circuit attributes, such as protection and type. CTC considers using a link for the circuit only if the link meets the following requirements:

- The link has sufficient bandwidth to support the circuit.
- The link does not change the protection characteristics of the path.
- The link has the required time slots to enforce the same time slot restrictions for MS-SPRing.

If CTC cannot find a link that meets these requirements, an error appears.

The same logic applies to low-order circuits on low-order tunnels. Circuit routing typically favors low-order tunnels because low-order tunnels are shortcuts between a given source and destination. If the low-order tunnel in the route is full (no more bandwidth), CTC asks whether you want to create an additional low-order tunnel.

10.11.2 Secondary Sources and Destinations

CTC supports secondary sources and destinations. Secondary sources and destinations typically interconnect two “foreign” networks (Figure 10-6). Traffic is protected while it goes through a network of ONS 15454 SDHs.

Figure 10-6 Secondary Sources and Destinations

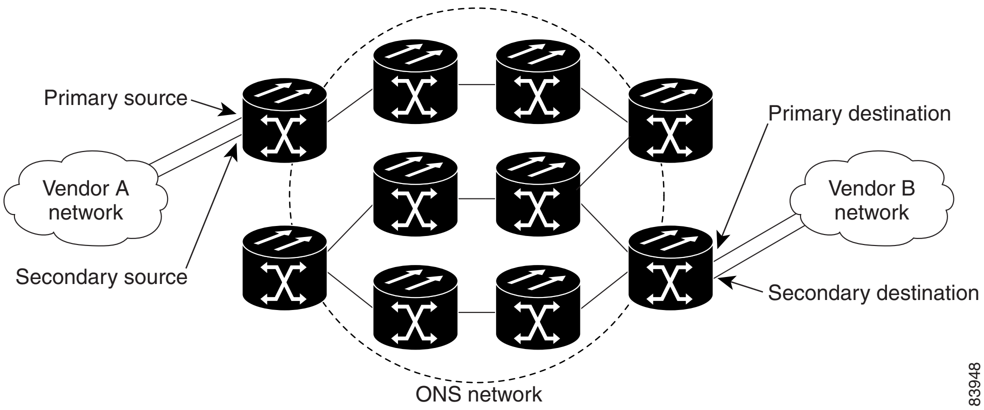

Several rules apply to secondary sources and destinations:

- CTC does not allow a secondary destination for unidirectional circuits, because you can always specify additional destinations (drops) after you create the circuit.
- Primary and secondary sources should be on the same node.
- Primary and secondary destinations should be on the same node.
- Secondary sources and destinations are permitted only for regular high-order or low-order connections (not for low-order tunnels and multicard EtherSwitch circuits).
- For point-to-point (straight) Ethernet circuits, only VC endpoints can be specified as multiple sources or drops.

For bidirectional circuits, CTC creates an SNCP connection at the source node that allows traffic to be selected from one of the two sources on the ONS 15454 SDH network. If you check the Fully Path Protected option during circuit creation, traffic is protected within the ONS 15454 SDH network. At the destination, another SNCP connection is created to bridge traffic from the ONS 15454 SDH network to the two destinations. A similar but opposite path exists for the reverse traffic flowing from the destinations to the sources.

For unidirectional circuits, an SNCP drop-and-continue connection is created at the source node.

10.12 Manual Circuit Routing

Routing circuits manually allows you to:

- Choose a specific path, not necessarily the shortest path.
- Choose a specific VC4/TUG3/TUG2/VC12 on each link along the route.
- Create a shared packet ring for multicard EtherSwitch circuits.
- Choose a protected path for multicard EtherSwitch circuits, allowing virtual SNCP segments.

CTC imposes the following rules on manual routes:

- All circuits, except multicard EtherSwitch circuits in a shared packet ring, should have links with a direction that flows from source to destination. This is true for multicard EtherSwitch circuits that are not in a shared packet ring.
- If you enabled Fully Path Protected, choose a diverse protect (alternate) path for every unprotected segment (Figure 10-7).

Figure 10-7 Alternate Paths for Virtual SNCP Segments

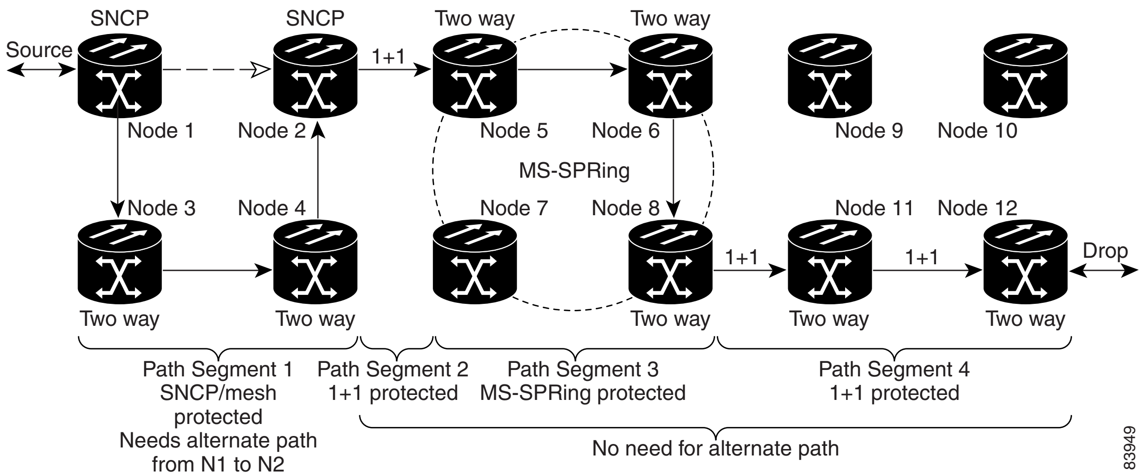

- For multicard EtherSwitch circuits, the Fully Path Protected option is ignored.
- For a node that has an SNCP selector based on the links chosen, the input links to the SNCP selectors cannot be 1+1 or MS-SPRing protected (Figure 10-8). The same rule applies at the SNCP bridge.

Figure 10-8 Mixing 1+1 or MS-SPRing Protected Links with an SNCP

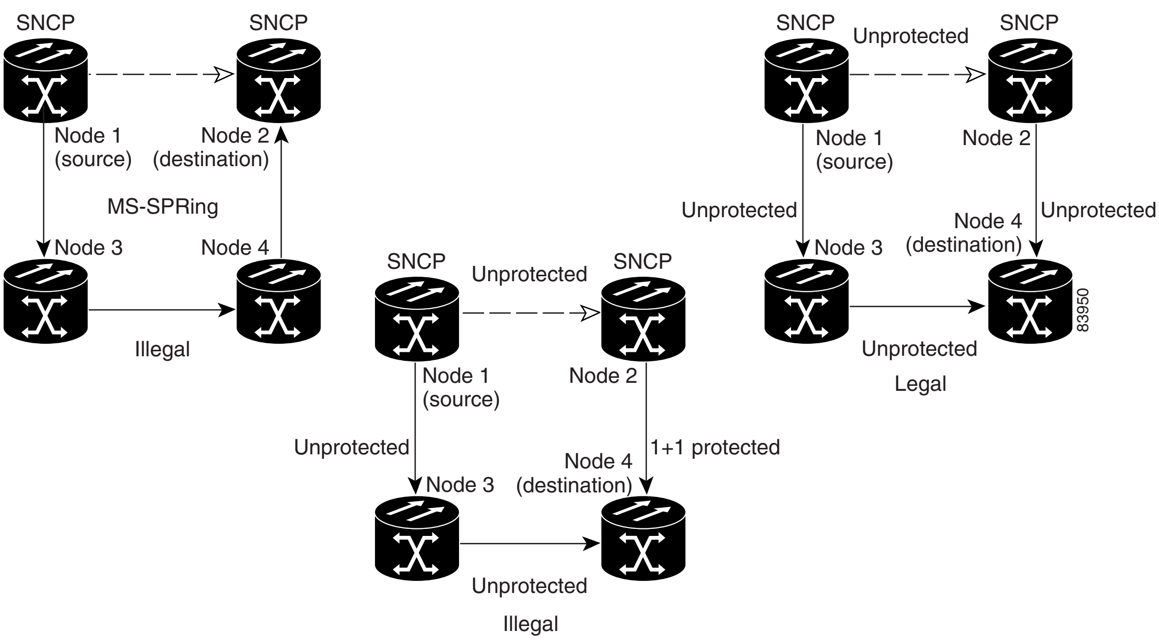

- Choose the links of multicard EtherSwitch circuits in a shared packet ring to route from source to destination back to source (Figure 10-9). Otherwise, a route (set of links) chosen with loops is invalid.

Figure 10-9 Ethernet Shared Packet Ring Routing

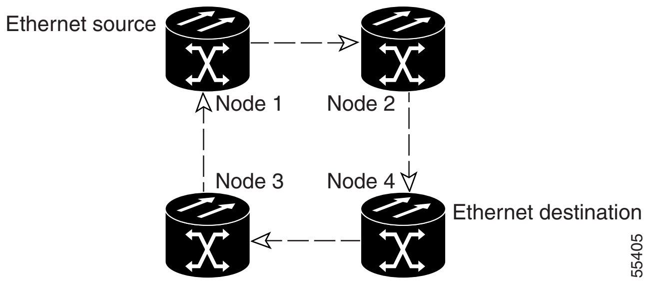

- Multicard EtherSwitch circuits can have virtual SNCP segments if the source or destination is not in the SNCP domain. This restriction also applies after circuit creation; therefore, if you create a circuit with SNCP segments, Ethernet drops cannot exist anywhere on the SNCP segment (Figure 10-10).

Figure 10-10 Ethernet and SNCP

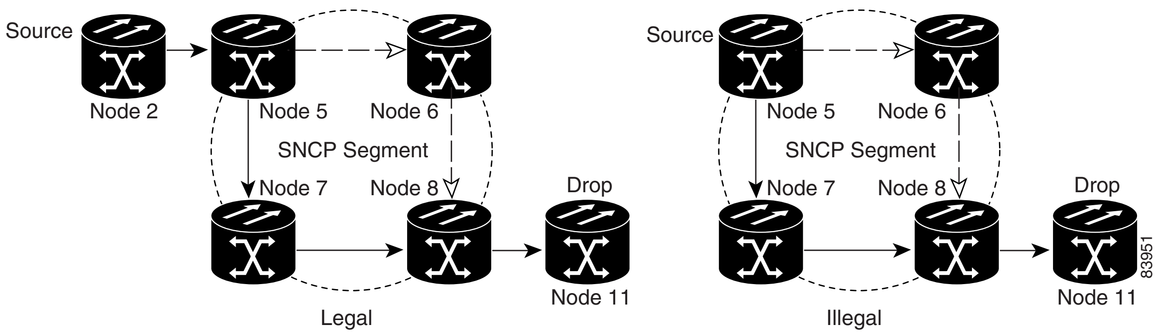

- Low-order tunnels cannot be the endpoint of an SNCP segment. A SNCP segment endpoint is where the SNCP selector resides.

If you provision full path protection, CTC verifies that the route selection is protected at all segments. A route can have multiple protection domains with each domain protected by a different scheme.

Table 10-9 through Table 10-12 on page 10-20 summarize the available node connections. Any other combination is invalid and generates an error.

Table 10-9 Bidirectional VC/TUG/Regular Multicard EtherSwitch/Point-to-Point (Straight) Ethernet Circuits

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Drops
SNCP	—	2	1	—
SNCP	2	—	—	1
SNCP	2	1	—	—
SNCP	1	2	—	—
SNCP	1	—	—	2
SNCP	—	1	2	—
Double SNCP	2	2	—	—

Table 10-9 Bidirectional VC/TUG/Regular Multicard EtherSwitch/Point-to-Point (Straight) Ethernet Circuits (continued)

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Drops
Double SNCP	2	—	—	2
Double SNCP	—	2	2	—
Two way	1	1	—	—
Ethernet	0 or 1	0 or 1	Ethernet node source	—
Ethernet	0 or 1	0 or 1	—	Ethernet node drop

Table 10-10 Unidirectional Circuit

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Drops
One way	1	1	—	—
SNCP head end	1	2	—	—
SNCP head end	—	2	1	—
SNCP drop and continue	2	—	—	1+

Table 10-11 Multicard Group Ethernet Shared Packet Ring Circuit

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Drops
Intermediate nodes only				
SNCP	2	1	—	—
SNCP	1	2	—	—
Double SNCP	2	2	—	—
Two way	1	1	—	-—
Source or destination nodes only				
Ethernet	1	1	—	—

Table 10-12 Bidirectional Low-Order Tunnels

Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Drops	Connection Type
Intermediate nodes only				
2	1	—	—	SNCP
1	2	—	—	SNCP
2	2	—	—	Double SNCP

Table 10-12 Bidirectional Low-Order Tunnels (continued)

Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Drops	Connection Type
1	1	—	—	Two way
Source nodes only				
—	1	—	—	Low-order tunnel endpoint
Destination nodes only				
1	—	—	—	Low-order tunnel endpoint

Although virtual SNCP segments are possible in low-order tunnels, low-order tunnels are still considered unprotected. If you need to protect low-order circuits, use two independent low-order tunnels that are diversely routed or use a low-order tunnel that is routed over 1+1, MS-SPRing, or a mixture of 1+1 and MS-SPRing links.

10.13 Constraint-Based Circuit Routing

When you create circuits, you can choose Fully Protected Path to protect the circuit from source to destination. The protection mechanism used depends on the path CTC calculates for the circuit. If the network is composed entirely of MS-SPRing or 1+1 links, or the path between source and destination can be entirely protected using 1+1 or MS-SPRing links, no path-protected mesh network (Extended SNCP) or virtual SNCP protection is used.

If Extended SNCP protection is needed to protect the path, set the level of node diversity for the Extended SNCP portions of the complete path in the Circuit Creation dialog box:

- Nodal Diversity Required—Ensures that the primary and alternate paths of each Extended SNCP domain in the complete path have a diverse set of nodes.
- Nodal Diversity Desired—CTC looks for a node diverse path; if a node-diverse path is not available, CTC finds a link-diverse path for each Extended SNCP domain in the complete path.
- Link Diversity Only—Creates only a link-diverse path for each Extended SNCP domain.

When you choose automatic circuit routing during circuit creation, you have the option to require or exclude nodes and links in the calculated route. You can use this option to:

- Simplify manual routing, especially if the network is large and selecting every span is tedious. You can select a general route from source to destination and allow CTC to fill in the route details.
- Balance network traffic; by default CTC chooses the shortest path, which can load traffic on certain links while other links have most of their bandwidth available. By selecting a required node or a link, you force the CTC to use (or not use) an element, resulting in more efficient use of network resources.

CTC considers required nodes and links to be an ordered set of elements. CTC treats the source nodes of every required link as required nodes. When CTC calculates the path, it makes sure the computed path traverses the required set of nodes and links and does not traverse excluded nodes and links.

The required nodes and links constraint is only used during the primary path computation and only for Extended SNCP domains/segments. The alternate path is computed normally; CTC uses excluded nodes/links when finding all primary and alternate paths on Extended SNCPs.

10.14 Virtual Concatenated Circuits

Virtual concatenated (VCAT) circuits, also called VCAT groups (VCGs), transport traffic using noncontiguous time division multiplexing (TDM) timeslots, avoiding the bandwidth fragmentation problem that exists with concatenated circuits. In a VCAT circuit, circuit bandwidth is divided into smaller circuits called VCAT members. The individual members act as independent TDM circuits.

Intermediate nodes treat the VCAT members as normal circuits that are independently routed and protected by the SDH network. At the terminating nodes, these member circuits are multiplexed into a contiguous stream of data. All VCAT members should be the same size and must originate/terminate at the same end points. Each member can be line protected, unprotected, or use PCA. If a member is unprotected, all members must be unprotected. Path protection is not supported.

Note Software Release 4.6 supports two members in VCAT circuits created using ML-Series cards and eight members in VCAT circuits created using the FC_MR-4 card.

The automatic and manual routing selection applies to the entire VCAT circuit, that is, all members are manually routed or automatically routed. In Software R4.6, bidirectional VCAT circuits are symmetric, which means that the same number of members travel in each direction. Software Release 4.6 supports common fiber routing, where all VCAT members travel on the same fibers, thus eliminating delay between members. Figure 10-11 shows an example of common fiber routing.

Figure 10-11 VCAT on Common Fiber

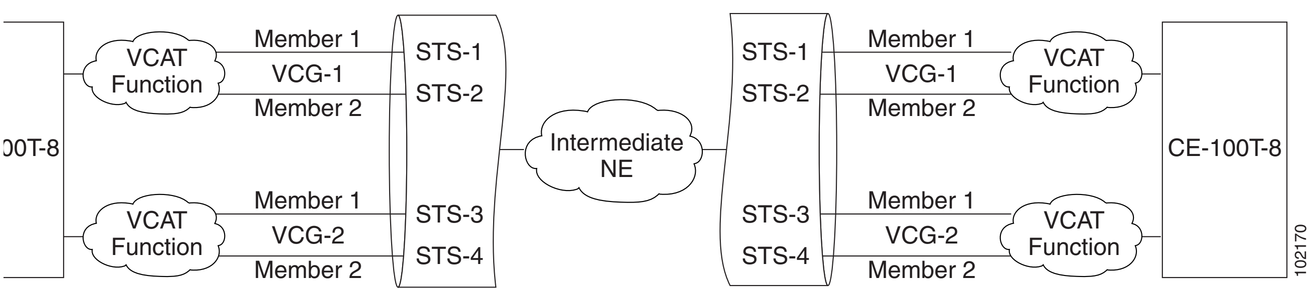

The Software–Link Capacity Adjustment Scheme (Sw-LCAS) uses legacy SDH failure indicators like the AIS-P and RDI-P to detect member failure. Sw-LCAS removes the failed member from the VCAT circuit for the duration of the failure, leaving the remaining members to carry the traffic. When the failure clears, the member circuit is added back into the VCAT circuit. Sw-LCAS cannot autonomously remove members that have defects in the H4/Z7 byte. Sw-LCAS is only available for legacy SDH defects such as AIS-P, LOP-P, etc. Sw-LCAS is optional. You can select Sw-LCAS during VCAT circuit creation.

Note Sw-LCAS allows circuit pairing for ML-Series cards over two-fiber MS-SPRing. With circuit pairing, a VCAT circuit is set up between two ML-Series cards; one is a protected circuit (line protection) and the other is PCA. For four-fiber MS-SPRing, member protection cannot be mixed.