



CHAPTER 12

Circuits and Tunnels



Note

The terms "Unidirectional Path Switched Ring" and "UPSR" may appear in Cisco literature. These terms do not refer to using Cisco ONS 15xxx products in a unidirectional path switched ring configuration. Rather, these terms, as well as "Path Protected Mesh Network" and "PPMN," refer generally to Cisco's path protection feature, which may be used in any topological network configuration. Cisco does not recommend using its path protection feature in any particular topological network configuration.

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:

- [12.1 Overview, page 12-2](#)
- [12.2 Circuit Properties, page 12-3](#)
- [12.3 Cross-Connect Card Bandwidth, page 12-12](#)
- [12.4 DCC Tunnels, page 12-12](#)
- [12.5 Multiple Destinations for Unidirectional Circuits, page 12-14](#)
- [12.6 Monitor Circuits, page 12-14](#)
- [12.7 SNCP Circuits, page 12-15](#)
- [12.8 MS-SPRing Protection Channel Access Circuits, page 12-16](#)
- [12.9 MS-SPRing VC4 Squelch Table, page 12-17](#)
- [12.10 IEEE 802.17 Resilient Packet Ring Circuit Display, page 12-17](#)
- [12.11 Section and Path Trace, page 12-18](#)
- [12.12 Path Signal Label, C2 Byte, page 12-19](#)
- [12.13 Automatic Circuit Routing, page 12-19](#)
- [12.14 Manual Circuit Routing, page 12-21](#)
- [12.15 Constraint-Based Circuit Routing, page 12-26](#)
- [12.16 Virtual Concatenated Circuits, page 12-26](#)
- [12.17 Bridge and Roll, page 12-32](#)
- [12.18 Merged Circuits, page 12-37](#)
- [12.19 Reconfigured Circuits, page 12-38](#)

- [12.20 Server Trails, page 12-38](#)

12.1 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 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. 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 you preprovision the small form-factor pluggables (SFPs) (also called pluggable port modules [PPMs]).
- When cards and SFPs are installed and ports are enabled. Circuits do not actually carry traffic until the cards and SFPs are installed and the ports are in the Unlocked-enabled state; the Locked-enabled,maintenance state; or the Unlocked-disabled,automaticInService state. Circuits carry traffic as soon as the signal is received.

12.2 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 12-1 on page 12-4](#)) 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), and LOP_v (low-order VCAT circuit).



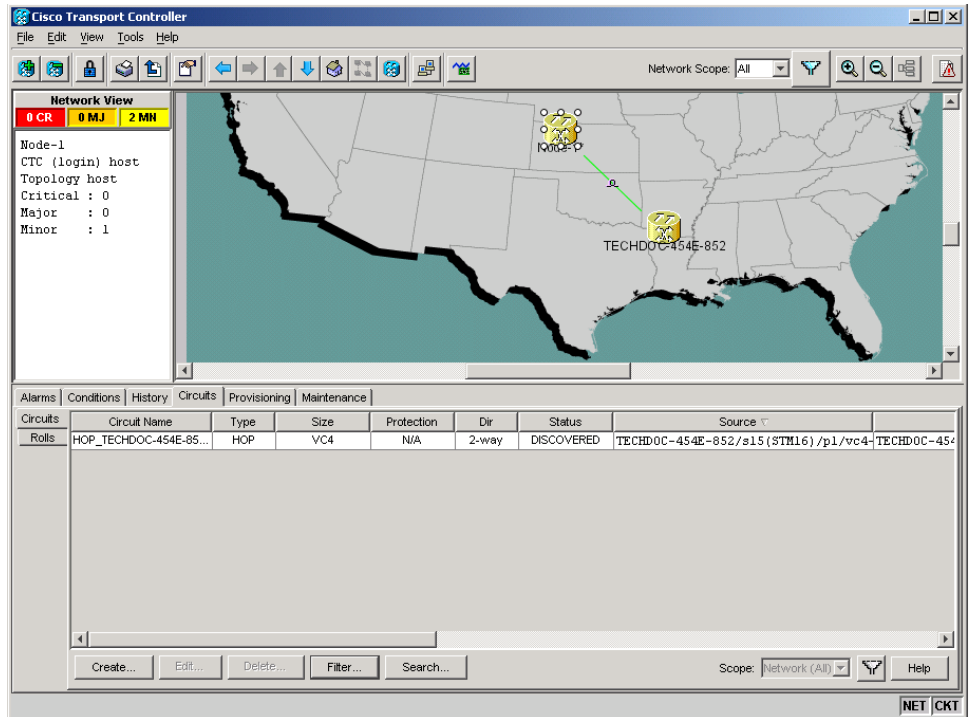
Note For OCHNC information, refer to the *Cisco ONS 15454 DWDM Procedure Guide*.

- Size—The circuit size. Low-order circuits are VC12, VC11 (XC-VXC-10G card only), and VC3. High-order circuit sizes are VC4, VC4-2c, VC4-3c, VC4-4c, VC4-6c, VC4-8c, VC4-12c, VC4-16c, and 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 and VC4-4c. OCHNCs are DWDM only, refer to the *Cisco ONS 15454 DWDM Procedure Guide* for more information. Low-order VCAT circuits are VC3 and VC12. For information on the number of supported members for each card, see [Table 12-14 on page 12-29](#).
- OCHNC Wlen—For OCHNCs, the wavelength provisioned for the DWDM optical channel network connection. (DWDM only; refer to the *Cisco ONS 15454 DWDM Procedure Guide* for more information).
- 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. (DWDM only; refer to the *Cisco ONS 15454 DWDM Procedure Guide* for more information).
- Protection—The type of circuit protection. See the [“12.2.4 Circuit Protection Types” section on page 12-9](#).
- Status—The circuit status. See the [“12.2.2 Circuit Status” section on page 12-6](#).
- 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. For the STM64-XFP and MRC-12 cards, the port appears as *port pluggable module (PPM)-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 internode links that constitute the circuit. Right-clicking the column shows a shortcut menu from which you can choose Span Details to show or hide circuit span detail. For each node in the span, the span detail shows the *node/slot/port/virtual container/tributary unit group/tributary unit group/virtual container*.

- State—The circuit state. See the “12.2.3 Circuit States” section on page 12-7.

The Filter button allows you to filter the circuits in network, node, or card view based on circuit name, size, type, direction, and other attributes. In addition, you can export the Circuit window data in HTML, comma-separated values (CSV), or tab-separated values (TSV) format using the Export command from the File menu.

Figure 12-1 ONS 15454 SDH Circuit Window in Network View



12.2.1 Concatenated VC4 Time Slot Assignments

Table 12-1 shows the available time slot assignments for concatenated VC4s when using CTC to provision circuits.

Table 12-1 VC4 Mapping Using CTC

Starting VC4	VC4	VC4-2c	VC4-3c	VC4-4c	VC4-6c	VC4-8c	VC4-12c	VC4-16c	VC4-64c
1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
3	Yes	Yes	No	No	Yes	Yes	Yes	No	No
4	Yes	No	Yes	No	Yes	Yes	Yes	No	No
5	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
6	Yes	Yes	Yes	No	Yes	Yes	No	No	No
7	Yes	Yes	Yes	No	Yes	Yes	No	No	No

Table 12-1 VC4 Mapping Using CTC (continued)

Starting VC4	VC4	VC4-2c	VC4-3c	VC4-4c	VC4-6c	VC4-8c	VC4-12c	VC4-16c	VC4-64c
8	Yes	No	No	No	Yes	Yes	No	No	No
9	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
10	Yes	Yes	Yes	No	Yes	No	No	No	No
11	Yes	Yes	No	No	Yes	No	No	No	No
12	Yes	No	No	No	No	No	No	No	No
13	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No
14	Yes	Yes	Yes	No	No	No	No	No	No
15	Yes	Yes	No	No	No	No	No	No	No
16	Yes	No	Yes	No	No	No	No	No	No
17	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
18	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
19	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
20	Yes	No	No	No	Yes	Yes	Yes	No	No
21	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
22	Yes	Yes	Yes	No	Yes	Yes	No	No	No
23	Yes	Yes	No	No	Yes	Yes	No	No	No
24	Yes	No	No	No	Yes	Yes	No	No	No
25	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
26	Yes	Yes	Yes	No	Yes	No	No	No	No
27	Yes	Yes	No	No	Yes	No	No	No	No
28	Yes	No	Yes	No	No	No	No	No	No
29	Yes	Yes	Yes	Yes	No	No	No	No	No
30	Yes	Yes	Yes	No	No	No	No	No	No
31	Yes	Yes	Yes	No	Yes	No	No	No	No
32	Yes	No	No	No	No	No	No	No	No
33	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
34	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
35	Yes	Yes	No	No	Yes	Yes	Yes	No	No
36	Yes	No	No	No	Yes	Yes	Yes	No	No
37	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
38	Yes	Yes	Yes	No	Yes	Yes	No	No	No
39	Yes	Yes	No	No	Yes	Yes	No	No	No
40	Yes	No	Yes	No	Yes	Yes	No	No	No
41	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
42	Yes	Yes	Yes	No	Yes	No	No	No	No

Table 12-1 VC4 Mapping Using CTC (continued)

Starting VC4	VC4	VC4-2c	VC4-3c	VC4-4c	VC4-6c	VC4-8c	VC4-12c	VC4-16c	VC4-64c
43	Yes	Yes	Yes	No	Yes	No	No	No	No
44	Yes	No	No	No	No	No	No	No	No
45	Yes	Yes	Yes	Yes	No	No	No	No	No
46	Yes	Yes	Yes	No	No	No	No	No	No
47	Yes	Yes	No	No	No	No	No	No	No
48	Yes	No	No	No	No	No	No	No	No
49	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
50	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
51	Yes	Yes	No	No	Yes	Yes	Yes	No	No
52	Yes	No	Yes	No	Yes	Yes	Yes	No	No
53	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
54	Yes	Yes	Yes	No	Yes	Yes	No	No	No
55	Yes	Yes	Yes	No	Yes	Yes	No	No	No
56	Yes	No	No	No	Yes	Yes	No	No	No
57	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
58	Yes	Yes	Yes	No	Yes	No	No	No	No
59	Yes	Yes	No	No	Yes	No	No	No	No
60	Yes	No	No	No	No	No	No	No	No
61	Yes	Yes	Yes	Yes	No	No	No	No	No
62	Yes	Yes	Yes	No	No	No	No	No	No
63	Yes	Yes	No	No	No	No	No	No	No
64	Yes	No	No	No	No	No	No	No	No

12.2.2 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 12-2](#) shows the statuses that can appear in the Status column.

Table 12-2 ONS 15454 SDH Circuit Status

Status	Definition/Activity
CREATING	CTC is creating a circuit.
DISCOVERED	CTC created a circuit. All components are in place and a complete path exists from circuit source to destination.
DELETING	CTC is deleting a circuit.

Table 12-2 ONS 15454 SDH Circuit Status (continued)

Status	Definition/Activity
PARTIAL	<p>A CTC-created circuit is missing a cross-connect or network span, a complete path from source to destination(s) does not exist, or an alarm interface panel (AIP) change occurred on one of the circuit nodes and the circuit is in need of repair. (AIPs store the node MAC address.)</p> <p>In CTC, circuits are represented using cross-connects and network spans. If a network span is missing from a circuit, the circuit status is PARTIAL. However, a PARTIAL status does not necessarily mean a circuit traffic failure has occurred, because traffic might flow on a protect path.</p> <p>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.</p> <p>Subsequently, circuits routed on a network span that goes down appear as DISCOVERED during the current CTC session, but appear as PARTIAL to users who log in after the span failure.</p>
DISCOVERED_TL1	A TL1-created circuit or a TL1-like, CTC-created circuit is complete. A complete path from source to destination(s) exists.
PARTIAL_TL1	A TL1-created circuit or a TL1-like, CTC-created circuit is missing a cross-connect or circuit span (network link), and a complete path from source to destination does not exist.
CONVERSION_PENDING	An existing circuit in a topology upgrade is set to this status. The circuit returns to the DISCOVERED status when the in-service topology upgrade is complete. For more information about in-service topology upgrades, see Chapter 11, “SDH Topologies and Upgrades.”
PENDING_MERGE	Any new circuits created to represent an alternate path in a topology upgrade are set to this status to indicate that it is a temporary circuit. These circuits can be deleted if an in-service topology upgrade fails. For more information about in-service topology upgrades, see Chapter 11, “SDH Topologies and Upgrades.”
DROP_PENDING	A circuit is set to this status when a new circuit drop is being added.

12.2.3 Circuit States

The circuit service state is an aggregate of the cross-connect states within the circuit.

- If all cross-connects in a circuit are in the Unlocked-enabled service state, the circuit service state is Unlocked.
- If all cross-connects in a circuit are in a Locked state (such as Locked-enabled,maintenance; Unlocked-disabled,automaticInService; or Locked-enabled,disabled service state) or the Unlocked-disabled,automaticInService state, the circuit service state is Locked.

- Partial is appended to the Locked circuit service state when circuit cross-connects state are mixed and not all in the Unlocked-enabled service state. The Locked-partial state can occur during automatic or manual transitions between states. The Locked-partial service state can appear during a manual transition caused by an abnormal event such as a CTC crash or 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. The Locked-partial circuit state does not apply to OCHNC circuit types.

You can assign a state to circuit cross-connects at two points:

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


Note

After you have created an initial circuit in a CTC session, the subsequent circuit states default to the circuit state of the initial circuit, regardless of which nodes in the network the circuits traverse or the node.ckt.state default setting.

During circuit creation, you can apply a service state to the drop ports in a circuit. You cannot transition a port from the Unlocked-enabled service state to the Locked-enabled,disabled state. You must first transition the port to the Locked-enabled,maintenance state before changing it to the Locked-enabled,disabled state. For more information about port service state transitions, see [Appendix B, “Administrative and Service States.”](#)

Circuits do not use the soak timer, but ports do. The soak period is the amount of time that the port remains in the Unlocked-disabled,automaticInService service state after a signal is continuously received. When the cross-connects in a circuit are in the Unlocked-disabled,automaticInService service state, the ONS 15454 SDH monitors the cross-connects for an error-free signal. It changes the state of the circuit from Locked to Unlocked or to Locked-partial as each cross-connect assigned to the circuit path is completed. This allows you to provision a circuit using TL1, verify its path continuity, and prepare the port to go into service when it receives an error-free signal for the time specified in the port soak timer.

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


Note

Although ML-Series cards do not use the Telcordia GR-1093-CORE state model, you can also set a soak timer for ML-Series cards ports. The soak period is the amount of time that the ML-Series port remains in the Down state after an error-free signal is continuously received before changing to the Up state. To find the remaining port soak time, choose the Maintenance > Ether/POS Port Soak tabs in ML-Series card view and click the Retrieve button.

For more information about cross-connect states, see [Appendix B, “Administrative and Service States.”](#)

12.2.4 Circuit Protection Types

The Protection column in the Circuit window shows the card (line) and SDH topology (path) protection used for the entire circuit path. [Table 12-3](#) shows the protection type indicators that appear in this column.

Table 12-3 *Circuit Protection Types*

Protection Type	Description
1+1	The circuit is protected by a 1+1 protection group.
2F MS-SPRing	The circuit is protected by a two-fiber MS-SPRing.
4F MS-SPRing	The circuit is protected by a four-fiber MS-SPRing.
2F-PCA	The circuit is routed on a protection channel access (PCA) path on a two-fiber MS-SPRing; PCA circuits are unprotected.
4F-PCA	The circuit is routed on a PCA path on a four-fiber MS-SPRing; PCA circuits are unprotected.
DRI	The circuit is protected by a dual-ring interconnection.
MS-SPRing	The circuit is protected by both a two-fiber and a four-fiber MS-SPRing.
N/A	A circuit with connections on the same node is not protected.
PCA	The circuit is routed on a PCA path on both two-fiber and four-fiber MS-SPRings; PCA circuits are unprotected.
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.
SNCP	The circuit is protected by an SNCP.
SPLITTER	The circuit is protected by the protect transponder (TXPP_MR_2.5G) splitter protection. For splitter information, refer to the <i>Cisco ONS 15454 DWDM Installation and Operations Guide</i> .
Unknown	A circuit has a source and destination on different nodes and communication is down between the nodes. This protection type appears if not all circuit components are known.
Unprot (black)	A circuit with a source and destination on different nodes is not protected.
Unprot (red)	A 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.
Y-Cable	The circuit is protected by a transponder or muxponder card Y-cable protection group. For more information, refer to the <i>Cisco ONS 15454 DWDM Installation and Operations Guide</i> .

12.2.5 Circuit Information in the Edit Circuit Window

You can edit a selected circuit using the Edit button on the Circuits window. The tabs that appear depend on the circuit chosen:

- General—Displays general circuit information and allows you to edit the circuit name.
- Drops—Allows you to add a drop to a unidirectional circuit. For more information, see the [“12.5 Multiple Destinations for Unidirectional Circuits”](#) section on page 12-14.

- **Monitors**—Displays possible monitor sources and allows you to create a monitor circuit. For more information, see the “[12.6 Monitor Circuits](#)” section on page 12-14.
- **SNCP Selectors**—Allows you to change SNCP selectors. For more information, see the “[12.7 SNCP Circuits](#)” section on page 12-15.
- **SNCP Switch Counts**—Allows you to change SNCP switch protection paths. For more information, see the “[12.7 SNCP Circuits](#)” section on page 12-15.
- **State**—Allows you to edit cross-connect service states.
- **Merge**—Allows you to merge aligned circuits. For more information, see the “[12.18 Merged Circuits](#)” section on page 12-37.

Using the Export command from the File menu, you can export data from the SNCP Selectors, SNCP Switch Counts, State, and Merge tabs in HTML, comma-separated values (CSV), or tab-separated values (TSV) format.

The Show Detailed Map checkbox in the Edit Circuit window updates the graphical view of the circuit to show more detailed routing information, such as:

- Circuit direction (unidirectional/bidirectional)
- The nodes, VC4s, VC3/TUG3, TUG2s, VC12s, and VC11s through which the circuit passes, including slots and port numbers
- The circuit source and destination points
- Open Shortest Path First (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 SNCP rings, the map shows the active and standby paths from circuit source to destination, and it also shows the working and protect paths. Selectors appear as pentagons on the detailed circuit map. The map indicates nodes set up as dual-ring interconnect nodes. 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 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 service states are indicated by colors, shown in [Table 12-4](#).

Table 12-4 Port State Color Indicators

Port Color	Service State
Green	Unlocked-enabled
Gray	Locked-enabled,disabled

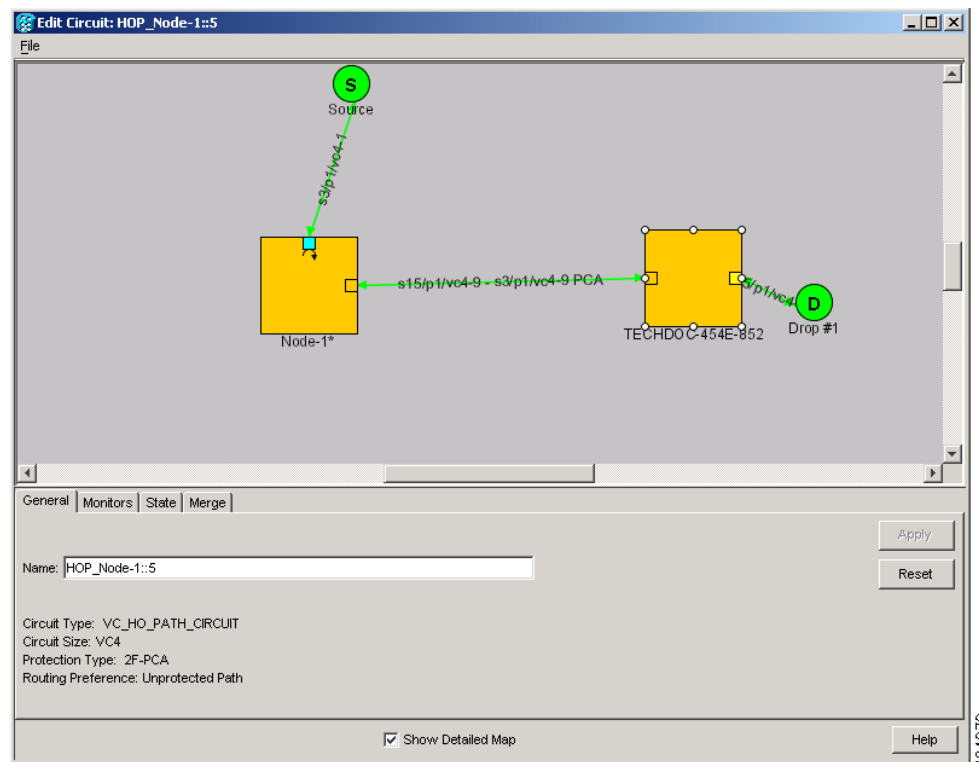
Table 12-4 Port State Color Indicators (continued)

Port Color	Service State
Violet	Unlocked-disabled,automaticInService
Blue (Cyan)	Locked-enabled,maintenance

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

- F = Force switch
- M = Manual switch
- L = Lockout switch
- Arrow = Facility (outward) or terminal (inward) loopback

Figure 12-2 shows an example of a 2F-PCA circuit with a card in terminal loopback in the Edit Circuits window.

Figure 12-2 Terminal Loopback in the Edit Circuits Window

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 service state, 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.

12.3 Cross-Connect Card Bandwidth

The XC-VXL-10G, XC-VXL-2.5G, and XC-VXC-10G cards support both low-order and high-order circuits, although only the XC-VXC-10G card supports VC-11 (low-order) circuits. The XC-VXL-10G and XC-VXL-2.5G cards 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 XC-VXC-10G card manages up to 576 bidirectional STM-1 cross-connects, 576 bidirectional E-3 or DS-3 cross-connects, or 1344 bidirectional E-1 cross-connects.

The XC-VXL-10G, XC-VXL-2.5G, and XC-VXC-10G cards work with the TCC2/TCC2P 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 XC-VXL-10G, XC-VXL-2.5G, and XC-VXC-10G cards.

12.4 DCC Tunnels

SDH provides four DCCs for network element operation, administration, maintenance, and provisioning: one on the SDH regenerator section layer (RS-DCC) and three on the SDH multiplex section layer, also called multiplex-section DCC (MS-DCC). A regenerator-section DCC (RS-DCC) and multiplex-section DCC (MS-DCC) each provide 192 Kbps of bandwidth per channel. The aggregate bandwidth of the three RS-DCCs 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.

12.4.1 Traditional DCC Tunnels

In traditional DCC tunnels, the ONS 15454 SDH uses RS-DCC for inter-ONS-15454-SDH data communications. It does not use the multiplex section DCCs; therefore, the MS-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 RS-DCC, Tunnel 1, Tunnel 2, or Tunnel 3. You can link an RS-DCC to a MS-DCC (Tunnel 1, Tunnel 2, or Tunnel 3) and a MS-DCC to an RS-DCC. You can also link MS-DCCs to MS-DCCs and link RS-DCCs to RS-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 12-5 shows the DCC tunnels that you can create.

Table 12-5 DCC Tunnels

DCC	SDH Layer	SDH Bytes	STM-1		STM-4, STM-16, STM-64
			4 Ports	8 Ports	
RS-DCC	Regenerator Section	D1 to D3	Yes	Yes	Yes
Tunnel 1	Multiplex Section	D4 to D6	No	Yes	Yes

Table 12-5 DCC Tunnels

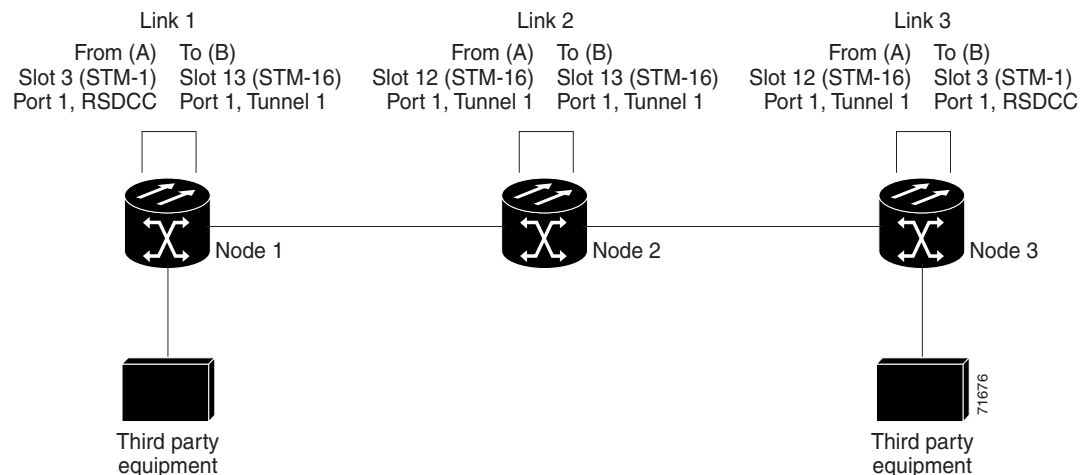
DCC	SDH Layer	SDH Bytes	STM-1		STM-4, STM-16, STM-64
			4 Ports	8 Ports	
Tunnel 2	Multiplex Section	D7 to D9	No	Yes	Yes
Tunnel 3	Multiplex Section	D10 to D12	No	Yes	Yes

Figure 12-3 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 12-3 Traditional 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 RS-DCC terminations.
- An RS-DCC that is terminated cannot be used as a DCC tunnel endpoint.
- An RS-DCC that is used as a DCC tunnel endpoint cannot be terminated.
- All DCC tunnel connections are bidirectional.

**Note**

An MS-DCC cannot be used for tunneling if a data DCC is assigned.

12.4.2 IP-Encapsulated Tunnels

An IP-encapsulated tunnel puts an RS-DCC 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 RS-DCC 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 DISCOVERED status can be converted.



Caution

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


12.5 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 have its 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.

12.6 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 12-4](#) 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 12-4](#) is used to monitor VC4 traffic received by Port 1 of the STM-1 card.

Figure 12-4 VC4 Monitor Circuit Received at an STM-1 Port



Note

Monitor circuits cannot be used with Ethernet circuits.

12.7 SNCP Circuits

Use the Edit Circuits window to change SNCP selectors and switch protection paths. In the SNCP Selectors subtab on the Edit Circuits window, you can:

- View the SNCP circuit's working and protection paths.

**Note**

The ONS 15310-MA and the CTX2500 card do not support SNCP switching for VC3 circuits containing BIP errors. The SF/SD alarm is not raised for VC3 circuits.

- Edit the reversion time.
- Set the hold-off timer.
- Edit the Signal Fail (SF)/Signal Degrade (SD) bit error rate (BER) thresholds.

**Note**

On the SNCP Selectors tab, the SF Ber Level and SD Ber Level columns display “N/A” for those nodes that do not support VC low-order signal BER monitoring. In Software Release 6.0, only the Cisco ONS 15310-CL supports VC low-order signal BER monitoring.

On the SNCP Switch Counts subtab, you can:

- Perform maintenance switches on the circuit selector.
- View switch counts for the selectors.

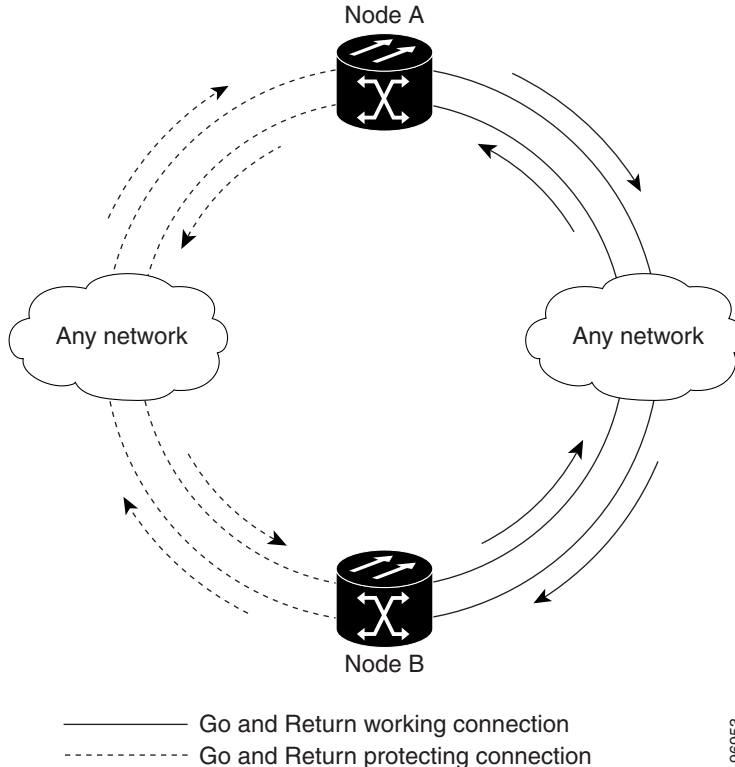
12.7.1 Open-Ended SNCP Circuits

If ONS 15454 SDH nodes 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 SDH nodes. That circuit routes the two circuit signals across the network to ONS 15454 SDH nodes 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.

12.7.2 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 12-5). The working path will always be the shortest path. If a fault occurs, neither the working fibers nor the protection fibers are 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 12-5 SNCP Go-and-Return Routing



12.8 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:

- 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.

12.9 MS-SPRing VC4 Squelch Table

MS-SPRing VC4 squelch tables show VC4s that will be squelched for every isolated node.

The MS-SPRing Squelch Table window displays the following information:

- **VC4 Number**—Shows the MS-SPRing VC4 numbers. For two-fiber MS-SPRings, the number of VC4s is half the MS-SPRing OC-N, for example, an STM-16 MS-SPRing squelch table will show 8 VC4s. For four-fiber MS-SPRings, the number of VC4s in the table is the same as the MS-SPRing STM-N.
- **West Source**—If traffic is received by the node on its west span, the MS-SPRing node ID of the source appears. (To view the MS-SPRing node IDs for all nodes in the ring, click the Ring Map button.)
- **West Dest**—If traffic is sent on the node's west span, the MS-SPRing node ID of the destination appears.
- **East Source**—If traffic is received by the node on its east span, the MS-SPRing node ID of the source appears.
- **East Dest**—If traffic is sent on the node's east span, the MS-SPRing node ID of the destination appears.

**Note**

MS-SPRing squelching is performed on VC4s that carry VC4 circuits only.

12.10 IEEE 802.17 Resilient Packet Ring Circuit Display

Resilient Packet Ring (RPR), as described in IEEE 802.17, is a metropolitan area network (MAN) technology supporting data transfer among stations interconnected in a dual-ring configuration. The IEEE 802.17b spatially-aware sublayer amendment is not yet ratified but is expected to add support for bridging to IEEE 802.17. Since the amendment is not yet ratified, no equipment is currently IEEE 802.17b compliant. The RPR-IEEE for ONS 15454 ML-Series cards is based on the expected IEEE 802.17b-based standard.

For more information about IEEE 802.17 RPR, refer to the *Cisco ONS 15454 and Cisco ONS 15454 SDH Ethernet Card Software Feature and Configuration Guide, Release 8.0*.

CTC provides a graphical representation (map) of IEEE 802.17 RPR circuits between ML-Series cards with a list of the following information:

- Circuit name
- Type
- Size
- OCHNC Wlen
- Direction
- Protection
- Status
- Source
- Destination
- # of VLANs

- # of Spans
- State
- Loopback



Note CTC does not support the display of Cisco proprietary RPR circuit topologies.



Note CTC does not support provisioning or maintenance of IEEE RPR rings. You must use Cisco IOS.

12.11 Section and Path Trace

SDH J1 and J2 path trace are repeated, fixed-length strings composed of 64 consecutive bytes. You can use the strings to monitor interruptions or changes to circuit traffic. The STM64-XFP and MRC-12 cards support J0 section trace. [Table 12-6](#) shows the ONS 15454 SDH cards that support J1 path trace. Cards that are not listed in the table do not support the J1 byte.

Table 12-6 ONS 15454 SDH Cards Capable of J1 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

[Table 12-7](#) shows cards that support J2 path trace.

Table 12-7 ONS 15454 SDH Cards Capable of J2 Path Trace

J2 Function	Cards
Transmit and Receive	E1-42
Receive Only	STM1E-12

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

12.12 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 12-8](#) provides the C2 byte hex values.

Table 12-8 STM Path Signal Label Assignments for Signals

Hex Code	Content of the STM Synchronous Payload Envelope (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 container-3 (C-3)
0x12	Asynchronous mapping of 139,264 kbps into container-4 (C-4)
0x13	Mapping for asynchronous transfer mode (ATM)
0x14	Mapping for Distributed Queue Dual Bus (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 indication 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 cross-connect card 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.

12.13 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.

12.13.1 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.

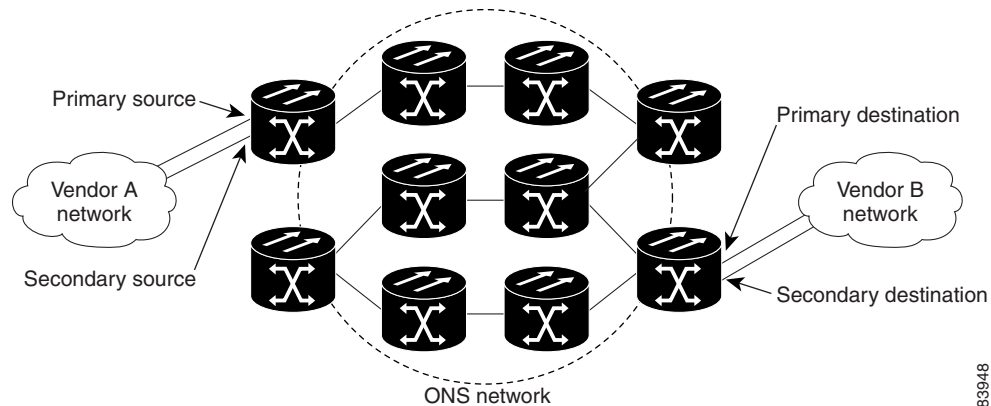
If CTC cannot find a link that meets these requirements, an error appears.

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.

12.13.2 Secondary Sources and Destinations

CTC supports secondary sources and destinations (drops). Secondary sources and destinations typically interconnect two “foreign” networks ([Figure 12-6](#)). Traffic is protected while it goes through a network of ONS 15454 SDH nodes.

Figure 12-6 Secondary Sources and Destinations



83948

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 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.



Note DRI and open-ended SNCP nodes allow primary and secondary sources and destinations on different nodes.

- Secondary sources and destinations are permitted only for regular high-order or low-order connections (not for low-order tunnels and multiscard 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.



Note Automatic routing and its associated subfields are not available if both the Automatic Circuit Routing NE default and the Network Circuit Automatic Routing Overridable NE default are set to FALSE. For a full description of these defaults see [Appendix C, “Network Element Defaults.”](#)

12.14 Manual Circuit Routing

Routing circuits manually allows you to:

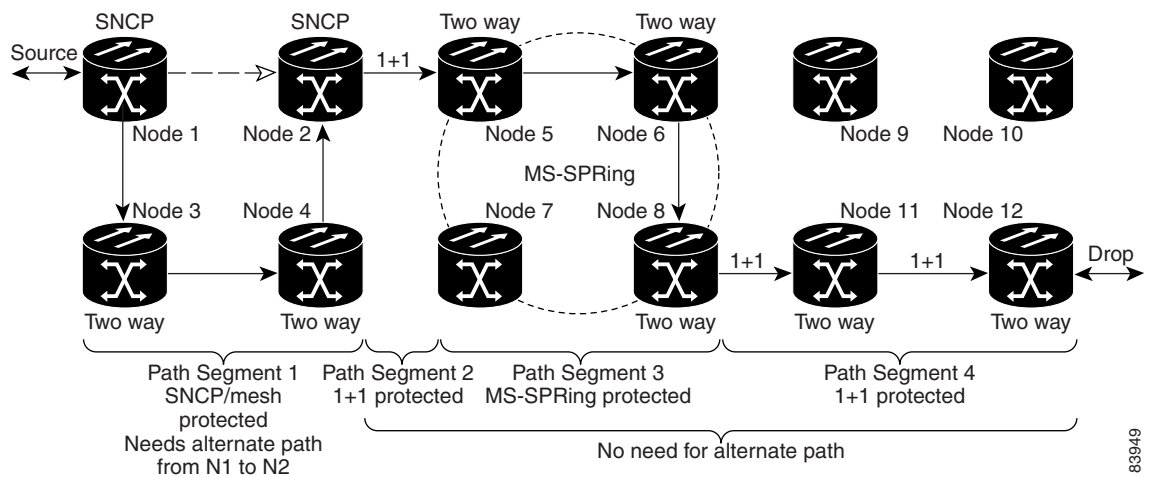
- Choose a specific path, not necessarily the shortest path.
- Choose a specific VC4/VC3/TUG3/TUG2/VC12/VC11 on each link along the route.

- Create a shared packet ring for multcard EtherSwitch circuits.
- Choose a protected path for multcard EtherSwitch circuits, allowing virtual SNCP segments.

CTC imposes the following rules on manual routes:

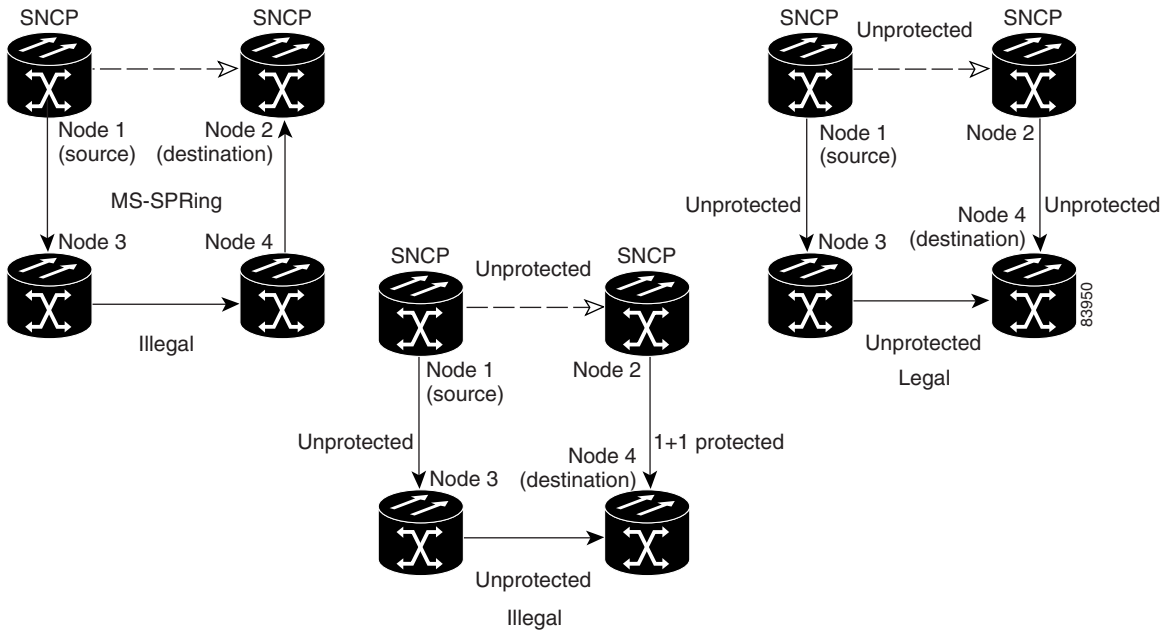
- All circuits, except multcard EtherSwitch circuits in a shared packet ring, should have links with a direction that flows from source to destination. This is true for multcard EtherSwitch circuits that are not in a shared packet ring.
- If you enabled Fully Path Protected, choose a diverse protect (alternate) path for every unprotected segment (Figure 12-7).

Figure 12-7 *Alternate Paths for Virtual SNCP Segments*




- For multcard 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 12-8). The same rule applies at the SNCP bridge.

Figure 12-8 *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 the circuit from source to destination back to source (Figure 12-9). Otherwise, a route (set of links) chosen with loops is invalid.

Figure 12-9 *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 12-10).

Figure 12-10 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 12-9 through Table 12-12 on page 12-25 summarize the available node connections. Any other combination is invalid and generates an error.

Table 12-9 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	—	—
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 12-10 Unidirectional Circuit

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Drops
One way	1	1	—	—
SNCP headend	1	2	—	—

Table 12-10 Unidirectional Circuit (continued)

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Drops
SNCP headend	—	2	1	—
SNCP drop and continue	2	—	—	1+

Table 12-11 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 12-12 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
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.

12.15 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.

12.16 Virtual Concatenated Circuits

Virtual concatenated (VCAT) circuits, also called VCAT groups (VCGs), transport traffic using noncontiguous time division multiplexing (TDM) time slots, avoiding the bandwidth fragmentation problem that exists with contiguous concatenated circuits. The cards that support VCAT circuits are the CE-Series, FC_MR-4 (both enhanced and line rate mode), and ML-Series cards.

In a VCAT circuit, circuit bandwidth is divided into smaller circuits called VCAT members. The individual members act as independent TDM circuits. All VCAT members should be the same size and must originate/terminate at the same end points. For two-fiber MS-SPRing configurations, some members can be routed on protected time slots and others on PCA time slots.

12.16.1 VCAT Circuit States

The state of a VCAT circuit is an aggregate of its member circuits. You can view whether a VCAT member is In Group or Out of Group in the VCAT State column on the Edit Circuits window.

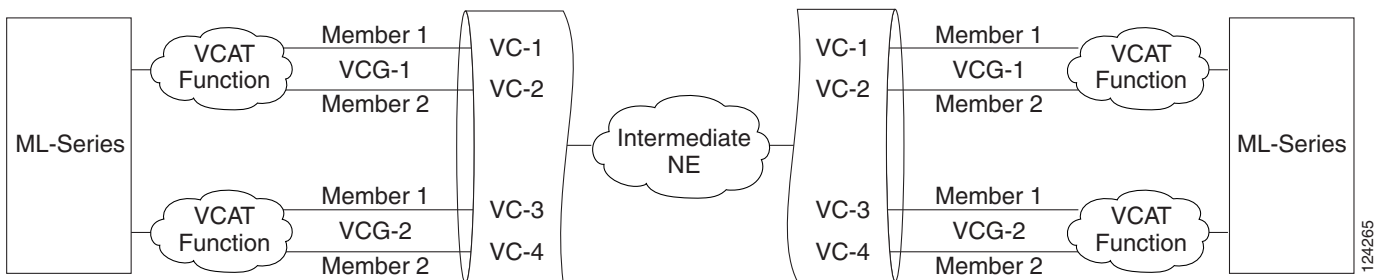
- If all member circuits are Unlocked, the VCAT circuit is Unlocked.
- If all In Group member circuits are Locked, the VCAT circuit state is Locked.
- If no member circuits exist or are all Out of Group, the state of a VCAT circuit is Locked.
- A VCAT circuit is Locked-partial when In Group member states are mixed and all members are not in the Unlocked state.

12.16.2 VCAT Member Routing

The automatic and manual routing selection applies to the entire VCAT circuit, that is, all members are manually or automatically routed. Bidirectional VCAT circuits are symmetric, which means that the same number of members travel in each direction. With automatic routing, you can specify the constraints for individual members; with manual routing, you can select different spans for different members.

Two types of automatic and manual routing are available for VCAT members: common fiber routing and split routing. CE-Series, FC_MR-4 (both line rate and enhanced mode), and ML-Series cards support common fiber routing. In common fiber routing, all VCAT members travel on the same fibers, which eliminates delay between members. Three protection options are available for common fiber routing: Fully Protected, PCA, and Unprotected. Each member can use a different protection scheme; however, CTC checks the combination to make sure a valid route exists and if it does not, the user must modify the protection type. [Figure 12-11](#) shows an example of common fiber routing.

Figure 12-11 VCAT Common Fiber Routing



CE-Series cards also support split fiber routing, which allows the individual members to be routed on different fibers or each member to have different routing constraints. This mode offers the greatest bandwidth efficiency and also the possibility of differential delay, which is handled by the buffers on the terminating cards. Four protection options are available for split fiber routing: Fully Protected, PCA, Unprotected, and DRI. [Figure 12-12](#) shows an example of split fiber routing.

Figure 12-12 VCAT Split Fiber Routing

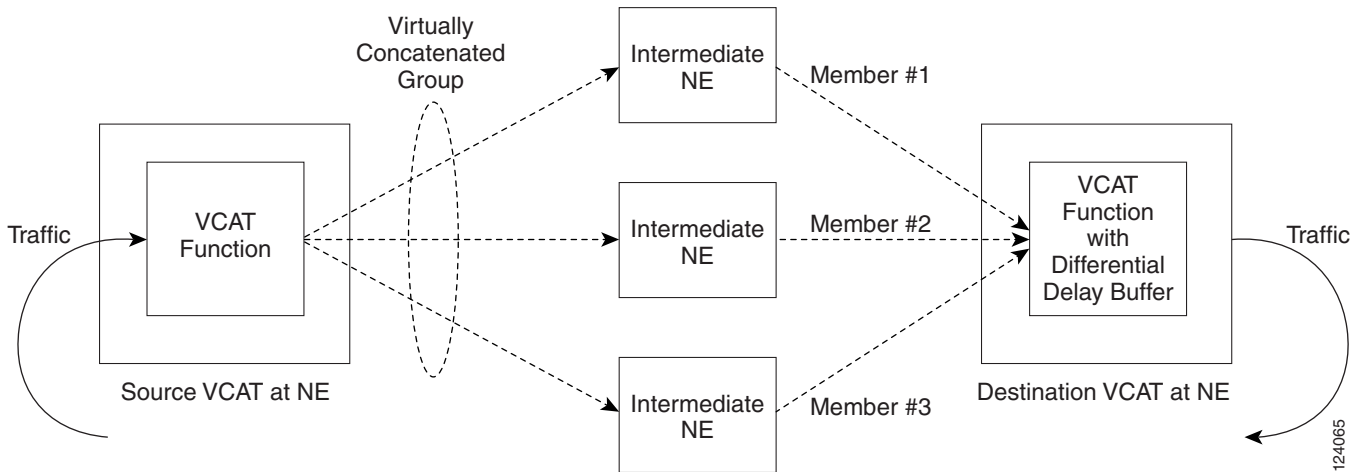


Table 12-13 Switch Times

Type of circuit	For CE100T-8 in ms
CCAT	60
HO VCAT	90
HO LCAS ¹	90
LO VCAT	202
LO LCAS	202

1. The calculated number for HO LCAS includes all the inherent delays of the protocol. Also the CE-100-T numbers are for a group size of only three members.

**Note**

The switch time values shown in Table 12-13 does not include differential delay. The maximum differential delay for CE100T-8 is 48ms. This differential delay is added to the switch time to get the maximum time.

12.16.3 Link Capacity Adjustment

The CE-100T-8 supports the link capacity adjustment scheme (LCAS), which is a signaling protocol that allows dynamic bandwidth adjustment of VCAT circuits. When a member fails, a brief traffic hit occurs. LCAS temporarily 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 automatically added back into the VCAT circuit without affecting traffic. You can select LCAS during VCAT circuit creation.

Instead of LCAS, the FC_MR-4 (enhanced mode), CE-1000-4, CE-MR-10, and ML-Series cards support Software-Link Capacity Adjustment Scheme (SW-LCAS), which uses legacy SONET failure indicators like the AIS-P and RDI-P to detect member failure. If used, 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 SONET defects such as AIS-P, LOP-P, etc. SW-LCAS is optional. You can select SW-LCAS during VCAT circuit creation. The FC_MR-4 card in line rate mode does not support SW-LCAS.

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 a four-fiber MS-SPRing, member protection cannot be mixed.

In addition, you can create non-LCAS VCAT circuits, which do not use SW-LCAS. While SW-LCAS member cross-connects can be in different service states, all In Group non-LCAS members must have cross-connects in the same service state. A non-LCAS circuit can mix Out of Group and In Group members if the In Group members are in the same service state. Non-LCAS members do not support the Locked-enabled, outOfGroup service state; to put a non-LCAS member in the Out of Group VCAT state, use Locked-enabled, disabled.

**Note**

Protection switching for LCAS, SW-LCAS, and non-LCAS VCAT circuits might exceed 60ms.

12.16.4 VCAT Circuit Size

Table 12-14 lists supported circuit rates and number of members for each card.

Table 12-14 ONS 15454 SDH Card VCAT Circuit Rates and Members

Card	Circuit Rate	Number of Members
CE-100T-8	VC12	1-64
	VC3	1-3 ¹
CE-1000-4	VC4	1-7 ¹
CE-MR-10	VC3	1-21
	VC4	1-7
	VC12	1-63
FC_MR-4 (Line rate mode)	VC4	8 (1-Gbps port)
		16 (2-Gbps port)
FC_MR-4 (Enhanced mode)	VC4	1-8 (1-Gbps port)
		1-16 (2-Gbps port)
ML-Series	VC3, VC4, VC4-4c	2
ML-MR-10	VC3	1-192 ²
	VC4	1-63
	VC12	1-63

1. A VCAT circuit with a CE-Series card as a source or destination and an ML-Series card as a source or destination can have only two members.
2. The maximum number of VCAT circuit members available for SNCP circuits are 126.

Use the Members tab in the Edit Circuit window to add or delete members from a VCAT circuit. The capability to add or delete members depends on the card and whether the VCAT circuit is LCAS, SW-LCAS, or non-LCAS.

- CE-100T-8 cards—Before deleting a member of an LCAS VCAT circuit, Cisco recommends that you put the member in the Locked-enabled,outOfGroup service state. If you create non-LCAS VCAT circuits on the CE-Series card, adding members to the circuit is possible, but service-affecting. You cannot delete members from non-LCAS VCAT circuits without affecting the entire VCAT circuit.
- CE-1000-4 and CE-MR-10 cards—You can add or delete SW-LCAS VCAT members, although it might affect service. Before deleting a member, Cisco recommends that you put the member in the Locked-enabled,outOfGroup service state. If you create non-LCAS VCAT circuits, adding and deleting members to the circuit is possible, but service-affecting.
- FC_MR-4 (enhanced mode) card—You can add or delete SW-LCAS VCAT members, although it might affect service. Before deleting a member, Cisco recommends that you put the member in the Locked-enabled,outOfGroup service state. You cannot add or delete members from VCAT circuits without SW-LCAS.
- FC_MR-4 (line rate mode) card—All VCAT circuits using FC_MR-4 (line rate mode) cards have a fixed number of members; you cannot add or delete members.
- ML-Series card—All VCAT circuits using ML-Series cards have a fixed number of members; you cannot add or delete members.

Table 12-15 summarizes the VCAT capabilities for each card.

Table 12-15 ONS 15454 SDH VCAT Card Capabilities

Card	Mode	Add a Member	Delete a Member	Support Locked-enabled, outOfGroup
CE-100T-8	LCAS	Yes ¹	Yes ¹	Yes
	SW-LCAS	No	No	No
	Non-LCAS	Yes ²	Yes ²	No
CE-1000-4	LCAS	No	No	No
	SW-LCAS	Yes	Yes	Yes
	Non-LCAS	Yes ²	Yes ²	No
CE-MR-10	LCAS	Yes	Yes	Yes
	SW-LCAS	Yes	Yes	Yes
	Non-LCAS	Yes	Yes	No
FC_MR-4 (enhanced mode)	SW-LCAS	Yes	Yes	Yes
	Non-LCAS	No	No	No
FC_MR-4 (line mode)	Non-LCAS	No	No	No
ML-Series	SW-LCAS	No	No	No
	Non-LCAS	No	No	No
ML-MR-10	LCAS	Yes	Yes	Yes
	SW-LCAS	Yes	Yes	Yes
	Non-LCAS	Yes	Yes	Yes

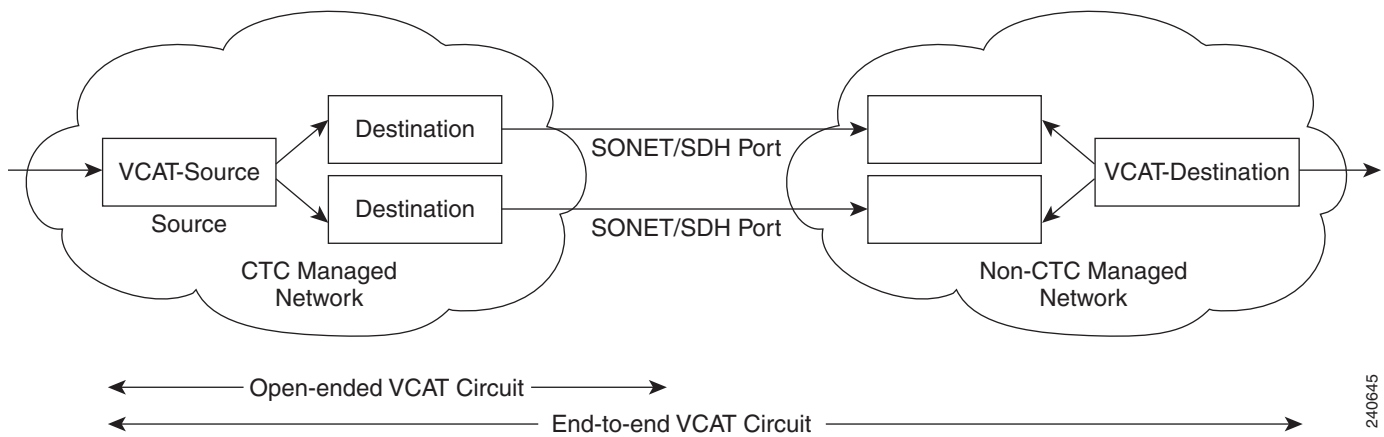
1. When adding or deleting a member from an LCAS VCAT circuit, Cisco recommends that you first put the member in the OOS-MA,OOG service state to avoid service disruptions.

- For CE-Series cards, you can add or delete members after creating a VCAT circuit with no protection. During the time it takes to add or delete members (from seconds to minutes), the entire VCAT circuit will be unable to carry traffic.

12.16.5 Open-Ended VCAT

For applications where the complete end-to-end VCAT circuit is not in a CTC managed network, CTC will only see either the source or the destination of the Virtual Concatenated Group (VCG) and some of the intermediate nodes. [Figure 12-13](#) shows an end-to-end VCAT circuit. The termination points of the end-to-end VCAT circuit, with VCAT functionality, are referred to as the VCAT-Source and VCAT-Destination. The termination points of the CTC managed circuit, which is the Open-Ended VCAT circuit, is referred to as simply the Source and Destination.

Figure 12-13 Open-Ended VCAT





Open-ended VCAT circuits can originate or terminate on any pair of OC-N ports and you can route open-ended VCAT circuits using any of the cards and ports supported by VCAT. The CTC circuit creation wizard provides an additional check box in the VCAT attributes pane to enable Open-VCAT circuit creation. Enabling the check box differentiates open-ended VCAT from regular VCAT Circuits.

The routing preferences for an open-ended VCAT circuit must be specified in the initial stages of circuit provisioning. For example, if the circuit is independent fiber routing, then multiple OC-N ports can be involved. Alternatively, the source of an open-VCAT circuit should always be a card capable of participating in a VCG. This allows CTC to determine which routing preferences are permissible.

12.16.5.1 Open-Ended VCAT Protection

[Table 12-16](#) summarizes the protection options for open-ended VCAT circuits. Note that members can have different routing preferences.

Table 12-16 Protection options for Open-Ended VCAT Circuits

Routing Preferences	Routing Mode	Protection Options
Common fiber	Manual/Auto	<ul style="list-style-type: none"> Fully protected (Line only) Unprotected PCA
Split fiber	Manual/Auto	<ul style="list-style-type: none"> Fully protected (Line only) Unprotected PCA DRI  <hr/> <p>Note Path protection is not supported.</p> <hr/>
Split fiber with secondary destinations	Manual/Auto	<ul style="list-style-type: none"> Fully protected  <hr/> <p>Note Line protection is not supported.</p> <hr/> <ul style="list-style-type: none"> DRI

12.17 Bridge and Roll

The CTC Bridge and Roll wizard reroutes live traffic without interrupting service. The bridge process takes traffic from a designated “roll from” facility and establishes a cross-connect to the designated “roll to” facility. When the bridged signal at the receiving end point is verified, the roll process creates a new cross-connect to receive the new signal. When the roll completes, the original cross-connects are released. You can use the bridge and roll feature for maintenance functions such as card or facility replacement, or for load balancing. You can perform a bridge and roll on the following ONS platforms: ONS 15600, ONS 15454, ONS 15454 SDH, and ONS 15310-CL, and ONS 15310-MA.

12.17.1 Rolls Window

The Rolls window lists information about a rolled circuit before the roll process is complete. You can access the Rolls window by clicking the Circuits > Rolls tabs in either network or node view.

Figure 12-14 shows the Rolls window.

Figure 12-14 Rolls Window

Rolls	Roll From Circuit	Roll To Circuit	Roll State	Roll Valid Signal	Roll Mode	Roll Path	Roll From Path	Roll To Path
	Roll Circuit	Roll Circuit	ROLL_PENDING	false	Auto	TECHDOC	TECHDOC	TECHDOC

The Rolls window information includes:

- Roll From Circuit—The circuit with connections that will no longer be used when the roll process is complete.
- Roll To Circuit—The circuit that will carry the traffic when the roll process is complete. The Roll To Circuit is the same as the Roll From Circuit if a single circuit is involved in a roll.
- Roll State—The roll status; see the “12.17.2 Roll Status” section on page 12-34 for information.
- Roll Valid Signal—If the Roll Valid Signal status is true, a valid signal was found on the new port. If the Roll Valid Signal status is false, a valid signal was not found. It is not possible to get a true Roll Valid Signal status for a one-way destination roll.
- Roll Mode—The mode indicates whether the roll is automatic or manual.

CTC implements a roll mode at the circuit level. TL1 implements a roll mode at the cross-connect level. If a single roll is performed, CTC and TL1 behave the same. If a dual roll is performed, the roll mode specified in CTC might be different than the roll mode retrieved in TL1. For example, if you select Automatic, CTC coordinates the two rolls to minimize possible traffic hits by using the Manual mode behind the scenes. When both rolls have a good signal, CTC signals the nodes to complete the roll.

- Automatic—When a valid signal is received on the new path, CTC completes the roll on the node automatically. One-way source rolls are always automatic. When the valid signal status is true, the Automatic mode switches the traffic to the Roll To Path and completes the roll automatically.
- Manual—You must complete a manual roll after a valid signal is received. One-way destination rolls are always manual. When the valid signal status is true, the Manual mode switches the traffic to the Roll To Path.
- Roll Path—The fixed point of the roll object.
- Roll From Path— The old path that is being rerouted.
- Roll To Path—The new path where the Roll From Path is rerouted.
- Complete—Completes a manual roll after a valid signal is received. You can complete a manual roll if it is in a ROLL_PENDING status and you have not yet completed the roll or have not cancelled its sibling roll. You cannot cancel the roll after you complete the roll.
- Force Valid Signal—Forces a roll onto the Roll To Circuit destination without a valid signal. If you choose Force Valid Signal, traffic on the circuit that is involved in the roll will be dropped when the roll is completed.

- **Finish**—Completes the circuit processing of both manual and automatic rolls and changes the circuit status from ROLL_PENDING to DISCOVERED. After a roll, the Finish button also removes any cross-connects that are no longer used from the Roll From Circuit field. The roll process ends when you finish the roll.
- **Cancel**—Cancels the roll process. When the roll mode is Manual, cancel roll is only allowed before you click the Complete button. When the roll mode is Auto, cancel roll is only allowed before a good signal is detected by the node or before you click the Force Valid Signal button.

12.17.2 Roll Status

Table 12-17 lists the roll statuses. You can only reroute circuits that have a DISCOVERED status. (See Table 12-2 on page 12-6 for a list of circuit statuses.) You cannot reroute circuits that are in the ROLL_PENDING status.

Table 12-17 Roll Statuses

State	Description
ROLL_PENDING	The roll is awaiting completion or cancellation.
ROLL_COMPLETED	The roll is complete. Click the Finish button.
ROLL_CANCELLED	The roll has been canceled.
TL1_ROLL	A TL1 roll was initiated. Note If a roll is created using TL1, a CTC user cannot complete or cancel the roll. Also, if a roll is created using CTC, a TL1 user cannot complete or cancel the roll. You must use the same interface to complete or change a roll.
INCOMPLETE	This state appears when the underlying circuit becomes incomplete. To correct this state, you must fix the underlying circuit problem before the roll state will change. For example, a circuit traveling on Nodes A, B, and C can become INCOMPLETE if Node B is rebooted. The cross-connect information is lost on Node B during a reboot. The Roll State on Nodes A and C will change to INCOMPLETE.


12.17.3 Single and Dual Rolls

Circuits have an additional layer of roll types: single and dual. A single roll on a circuit is a roll on one of its cross-connects. Use a single roll to:

- Change either the source or destination of a selected circuit (Figure 12-15 and Figure 12-16, respectively).
- Roll a segment of the circuit onto another chosen circuit (Figure 12-17 on page 12-35). This roll also results in a new destination or a new source.


In Figure 12-15, you can select any available VC4 on Node 1 for a new source.

Figure 12-15 Single Source Roll




In [Figure 12-16](#), you can select any available VC4 on Node 2 for a new destination.

Figure 12-16 Single Destination Roll




[Figure 12-17](#) shows one circuit rolling onto another circuit at the destination. The new circuit has cross-connects on Node 1, Node 3, and Node 4. CTC deletes the cross-connect on Node 2 after the roll.

Figure 12-17 Single Roll from One Circuit to Another Circuit (Destination Changes)



[Figure 12-18](#) shows one circuit rolling onto another circuit at the source.

Figure 12-18 Single Roll from One Circuit to Another Circuit (Source Changes)



**Note**

Create a Roll To Circuit before rolling a circuit with the source on Node 3 and the destination on Node 4.

A dual roll involves two cross-connects. It allows you to reroute intermediate segments of a circuit, but keep the original source and destination. If the new segments require new cross-connects, use the Bridge and Roll wizard or create a new circuit and then perform a roll.

**Caution**

Only single rolls can be performed using TL1. Dual rolls require the network-level view that only CTC or CTM provide.

Dual rolls have several constraints:

- You must complete or cancel both cross-connects rolled in a dual roll. You cannot complete one roll and cancel the other roll.
- When a Roll To circuit is involved in the dual roll, the first roll must roll onto the source of the Roll To circuit and the second roll must roll onto the destination of the Roll To circuit.

Figure 12-19 illustrates a dual roll on the same circuit.

Figure 12-19 Dual Roll to Reroute a Link





Figure 12-20 illustrates a dual roll involving two circuits.

Figure 12-20 Dual Roll to Reroute to a Different Node

**Note**

If a new segment is created on Nodes 3 and 4 using the Bridge and Roll wizard, the created circuit has the same name as the original circuit with the suffix `_ROLL**`. The circuit source is on Node 3 and the circuit destination is on Node 4.

12.17.4 Two Circuit Bridge and Roll

When using the bridge and roll feature to reroute traffic using two circuits, the following constraints apply:

- DCC must be enabled on the circuits involved in a roll before roll creation.
- A maximum of two rolls can exist between any two circuits.
- If two rolls are involved between two circuits, both rolls must be on the original circuit. The second circuit should not carry live traffic. The two rolls loop from the second circuit back to the original circuit. The roll mode of the two rolls must be identical (either automatic or manual).
- If a single roll exists on a circuit, you must roll the connection onto the source or the destination of the second circuit and not an intermediate node in the circuit.

12.17.5 Protected Circuits

CTC allows you to roll the working or protect path regardless of which path is active. You can upgrade an unprotected circuit to a fully protected circuit or downgrade a fully protected circuit to an unprotected circuit with the exception of an SNCP circuit. When using bridge and roll on SNCP circuits, you can roll the source or destination or both path selectors in a dual roll. However, you cannot roll a single path selector.

12.18 Merged Circuits

A circuit merge combines a single selected circuit with one or more circuits. You can merge VCTs, VCA circuits, VLAN-assigned circuits, VCAT members, orderwire and user data channel overhead circuits, CTC-created traffic circuits, and TL1-created traffic circuits. To merge circuits, you choose a circuit on the CTC Circuits tab and the circuits that you want to merge with the chosen (master) circuit on the Merge tab in the Edit Circuits window. The Merge tab shows only the circuits that are available for merging with the master circuit:

- Circuit cross-connects must create a single, contiguous path.
- Circuit types must be compatible. For example, you can combine a HOP with a VCA circuit to create a longer VCA circuit, but you cannot combine a LOP with a HOP.
- Circuit directions must be compatible. You can merge a one-way and a two-way circuit, but not two one-way circuits in opposing directions.
- Circuit sizes must be identical.
- VLAN assignments must be identical.
- Circuit end points must send or receive the same framing format.
- The merged circuits must become a DISCOVERED circuit.

If all connections from the master circuit and all connections from the merged circuits align to form one complete circuit, the merge is successful. If all connections from the master circuit and some, but not all, connections from the other circuits align to form a single complete circuit, CTC notifies you and gives you the chance to cancel the merge process. If you choose to continue, the aligned connections merge successfully into the master circuit, and unaligned connections remain in the original circuits. All connections in the completed master circuit use the original master circuit name.

All connections from the master circuit and at least one connection from the other selected circuits must be used in the resulting circuit for the merge to succeed. If a merge fails, the master circuit and all other circuits remain unchanged. When the circuit merge completes successfully, the resulting circuit retains the name of the master circuit.

12.19 Reconfigured Circuits

You can reconfigure multiple circuits, which is typically necessary when a large number of circuits are in the PARTIAL state. When reconfiguring multiple circuits, the selected circuits can be any combination of DISCOVERED, PARTIAL, DISCOVERED_TL1, or PARTIAL_TL1 circuits. You can reconfigure VCTs, VCA circuits, VLAN-assigned circuits, VCAT circuits, CTC-created circuits, and TL1-created circuits. The Reconfigure command maintains the names of the original cross-connects.

Use the CTC Tools > Circuits > Reconfigure Circuits command to reconfigure selected circuits. During reconfiguration, CTC reassembles all connections of the selected circuits into circuits based on path size, direction, and alignment. Some circuits might merge and others might split into multiple circuits. If the resulting circuit is a valid circuit, it appears as a DISCOVERED circuit. Otherwise, the circuit appears as a PARTIAL or PARTIAL_TL1 circuit.



Note

If CTC cannot reconfigure all members in a VCAT circuit, the reconfigure operation fails for the entire VCAT circuit and it remains in the PARTIAL or PARTIAL_TL1 status. If CTC does reconfigure all members in a VCAT circuit, the VCAT circuit may still remain in the PARTIAL or PARTIAL_TL1 status. This occurs if the ports defined in the VCAT termination do not match the source/drop ports of the member circuits or if one or two VCAT terminations are missing.



Note

PARTIAL tunnel and PARTIAL VLAN-capable circuits do not split into multiple circuits during reconfiguration.

12.20 Server Trails

A server trail is a non-DCC (logical or virtual) link across a third-party network that connects two CTC network domains. A server trail allows A-Z circuit provisioning when no DCC is available. You can create server trails between two distant optical or STM-1E ports. The end ports on a server trail can be different types (for example, an STM-4 port can be linked to an STM-1 port). Server trails are not allowed on DCC-enabled ports.

The server trail link is bidirectional and can be VC3, VC11, VC12, VC4, VC4-2c, VC4-3c, VC4-4c, VC4-6c, VC4-8c, VC4-12c, VC4-16c, VC4-32c, and VC4-64c; you cannot change an existing server trail to another size. It must be deleted and recreated. A circuit provisioned over a server trail must match the type and size of the server trail it uses. For example, an VC4-3c server trail can carry only VC4-3c circuits and not three VC4 circuits.



Note

There is no OSPF or any other management information exchange between NEs over a server trail.

12.20.1 Server Trail Protection Types

The server trail protection type determines the protection type for any circuits that traverse it. A server trail link can be one of the following protection types:

- Preemptible—PCA circuits will use server trails with the Preemptible attribute.
- Unprotected—In Unprotected Server Trail, CTC assumes that the circuits going out from that specific port will not be protected by provider network and will look for a secondary path from source to destination if you are creating a protected circuit.
- Fully Protected—In Fully Protected Server Trail, CTC assumes that the circuits going out from that specific port will be protected by provider network and will not look for a secondary path from source to destination.

**Note**

Only SNCP protection is available on server trails. MS-SPRing protection is not available on server trail.

12.20.2 VCAT Circuit Routing over Server Trails

An VC4-3c server trail can be used to route VC4-3c circuits and an VC4 server trail can be used to route VC4 circuits. Similarly, a VC3 server trail can be used to route VC3 circuits.

For example, to route a VC4-3c-2v circuit over a server trail, you must enable split fiber routing and create two VC4-3c server trails and route each member manually or automatically over each server trail. To route a VC4-12c-2v circuit over a server trail, you must enable split fiber routing and create two VC4 server trails and route each member manually or automatically over each server trail.

**Note**

Server trails can only be created between any two optical ports or STM-1E ports.

VCAT circuits can be created over server trails in the following ways:

- Manual routing
- Automatic routing
 - Diverse routing: This method enables VCAT circuit routing over diverse server trail links.

**Note**

When creating circuits or VCATs, you can choose a server trail link during manual circuit routing. CTC may also route circuits over server trail links during automatic routing. VCAT common-fiber automatic routing is not supported.

For a detailed procedure on how to route a VCAT circuit over a server trail, refer “Chapter 6, Create Circuits and VT Tunnels, Section NTP-A264, Create an Automatically Routed VCAT Circuit and Section NTP-A265, Create a Manually Routed VCAT Circuit” in the *Cisco ONS 15454 Procedure Guide*.

12.20.2.1 Shared Resource Link Group

The Shared Resource Link Group (SRLG) attribute can be assigned to a server trail link using a commonly shared resource such as port, fiber or span. For example, if two server trail links are routed over the same fiber, an SRLG attribute can be assigned to these links. SRLG is used by Cisco Transport Manager (CTM) to specify link diversity. If you create multiple server trails from one port, you can assign the same SRLG value to all the links to indicate that they originate from the same port.