



CHAPTER 11

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 synchronous transport signal (STS), virtual tributary (VT), and virtual concatenated (VCAT) circuits and VT, data communications channel (DCC), and IP-encapsulated tunnels. To provision circuits and tunnels, refer to the *Cisco ONS 15454 Procedure Guide*.

Chapter topics include:

- [11.1 Overview, page 11-2](#)
- [11.2 Circuit Properties, page 11-2](#)
- [11.3 Cross-Connect Card Bandwidth, page 11-12](#)
- [11.4 Portless Transmux, page 11-15](#)
- [11.5 DCC Tunnels, page 11-16](#)
- [11.7 Multiple Destinations for Unidirectional Circuits, page 11-18](#)
- [11.8 Monitor Circuits, page 11-18](#)
- [11.9 Path Protection Circuits, page 11-19](#)
- [11.10 BLSR Protection Channel Access Circuits, page 11-21](#)
- [11.11 BLSR STS and VT Squelch Tables, page 11-22](#)
- [11.12 IEEE 802.17 Resilient Packet Ring Circuit Display, page 11-23](#)
- [11.13 Section and Path Trace, page 11-24](#)
- [11.14 Path Signal Label, C2 Byte, page 11-25](#)
- [11.15 Automatic Circuit Routing, page 11-27](#)
- [11.16 Manual Circuit Routing, page 11-29](#)
- [11.17 Constraint-Based Circuit Routing, page 11-33](#)
- [11.18 Virtual Concatenated Circuits, page 11-34](#)
- [11.19 Bridge and Roll, page 11-39](#)
- [11.20 Merged Circuits, page 11-45](#)

- [11.21 Reconfigured Circuits, page 11-46](#)
- [11.22 VLAN Management, page 11-46](#)
- [11.23 Server Trails, page 11-46](#)

11.1 Overview

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

- Create one-way, two-way (bidirectional), or broadcast circuits.
- Assign user-defined names to circuits.
- Assign different circuit sizes.
- Automatically or manually route circuits.
- Automatically create multiple circuits with autoranging. VT 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 path protection configuration with third-party equipment path protection configurations.
- Set path protection circuits as revertive or nonrevertive.

You can provision circuits at either of the following points:

- Before cards are installed. The ONS 15454 allows you to provision slots and circuits before installing the traffic cards.
- After you preprovision the Small Form-factor Pluggables (SFPs) (also called provisionable port modules [PPMs]).
- After cards and SFPs are installed and ports are in service. Circuits do not actually carry traffic until the cards and SFPs are installed and the ports are In-Service and Normal (IS-NR); Out-of-Service and Autonomous, Automatic In-Service (OO-AU,AINS); or Out-of-Service and Management, Maintenance (OOS-MA,MT). Circuits carry traffic as soon as the signal is received.

11.2 Circuit Properties

The ONS 15454 Cisco Transport Controller (CTC) Circuits window, which appears in network, node, and card view, is where you can view information about circuits. The Circuits window ([Figure 11-1](#)) provides the following information:

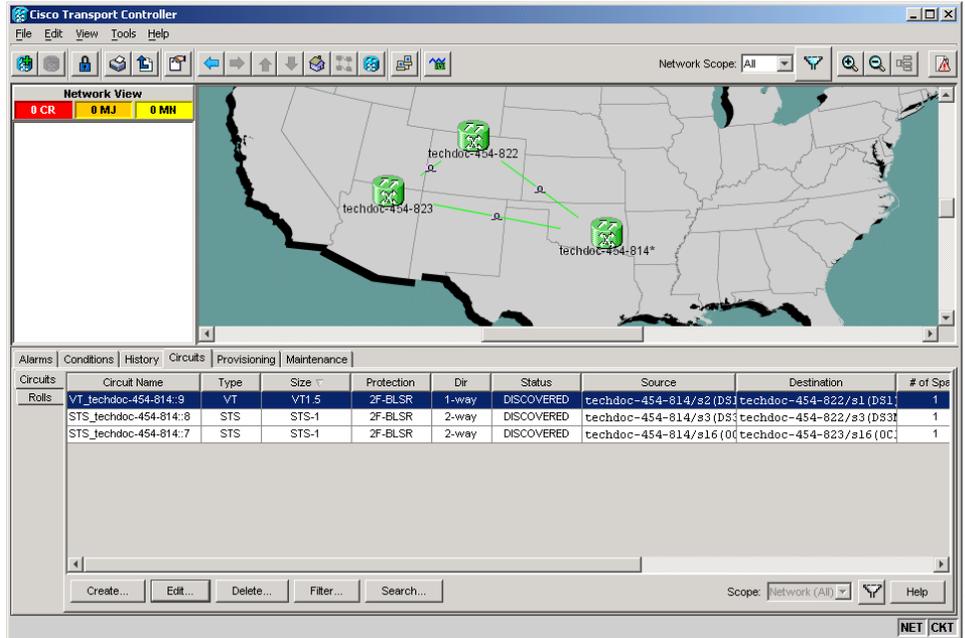
- Name—The name of the circuit. The circuit name can be manually assigned or automatically generated.
- Type—The circuit types are STS (STS circuit), VT (VT circuit), VTT (VT tunnel), VAP (VT aggregation point), OCHNC (dense wavelength division multiplexing [DWDM] optical channel network connection; refer to the *Cisco ONS 15454 DWDM Procedure Guide*), STS-V (STS VCAT circuit), or VT-V (VT VCAT circuit).
- Size—The circuit size. VT circuits are 1.5. STS circuit sizes are 1, 3c, 6c, 9c, 12c, 24c, 36c, 48c, and 192c. OCHNC sizes are Equipped non specific, Multi-rate, 2.5 Gbps No FEC (forward error correction), 2.5 Gbps FEC, 10 Gbps No FEC, and 10 Gbps FEC (OCHNC is DWDM only; refer to

the *Cisco ONS 15454 DWDM Procedure Guide*). VCAT circuits are VT1.5-*nv*, STS-1-*nv*, STS-3c-*nv*, and STS-12c-*nv*, where *n* is the number of members. For time slot availability on concatenated STSs, see the “[11.2.1 Concatenated STS Time Slot Assignments](#)” section on [page 11-4](#).

- OCHNC Wlen—For OCHNCs, the wavelength provisioned for the optical channel network connection. For more information, refer to the *Cisco ONS 15454 DWDM Procedure Guide*.
- Direction—The circuit direction, either two-way or one-way.
- OCHNC Dir—For OCHNCs, the direction of the optical channel network connection, either east to west or west to east. For more information, refer to the *Cisco ONS 15454 DWDM Procedure Guide*.
- Protection—The type of circuit protection. See the “[11.2.4 Circuit Protection Types](#)” section on [page 11-9](#) for a list of protection types.
- Status—The circuit status. See the “[11.2.2 Circuit Status](#)” section on [page 11-6](#).
- Source—The circuit source in the format: *node/slot/port “port name”/STS/VT*. (The port name appears in quotes.) Node and slot always appear; *port “port name”/STS/VT* might appear, depending on the source card, circuit type, and whether a name is assigned to the port. For the OC192-XFP and MRC-12 cards, the port appears as *port pluggable module (PPM)-port*. If the circuit size is a concatenated size (3c, 6c, 12c, etc.), STSs used in the circuit are indicated by an ellipsis, for example, S7..9, (STSs 7, 8, and 9) or S10..12 (STS 10, 11, and 12).
- Destination—The circuit destination in the same format as the circuit source.
- # of VLANs—The number of VLANs used by an Ethernet circuit.
- # 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 (card type)/port/STS/VT*.
- State—The circuit state. See the “[11.2.3 Circuit States](#)” section on [page 11-7](#).

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 11-1 ONS 15454 Circuit Window in Network View



11.2.1 Concatenated STS Time Slot Assignments

Table 11-1 shows the available time slot assignments for concatenated STSs when using CTC to provision circuits.

Table 11-1 STS Mapping Using CTC

Starting STS	STS-3c	STS-6c	STS-9c	STS-12c	STS-18c	STS-24c	STS-36c	STS-48c	STS-192c
1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
4	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
7	Yes	Yes	No	No	Yes	Yes	Yes	No	No
10	Yes	No	Yes	No	Yes	Yes	Yes	No	No
13	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
16	Yes	Yes	Yes	No	Yes	Yes	No	No	No
19	Yes	Yes	Yes	No	Yes	Yes	No	No	No
22	Yes	No	No	No	Yes	Yes	No	No	No
25	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
28	Yes	Yes	Yes	No	Yes	No	No	No	No
31	Yes	Yes	No	No	Yes	No	No	No	No
34	Yes	No	No	No	No	No	No	No	No
37	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No

Table 11-1 STS Mapping Using CTC (continued)

Starting STS	STS-3c	STS-6c	STS-9c	STS-12c	STS-18c	STS-24c	STS-36c	STS-48c	STS-192c
40	Yes	Yes	Yes	No	No	No	No	No	No
43	Yes	Yes	No	No	No	No	No	No	No
46	Yes	No	Yes	No	No	No	No	No	No
49	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
52	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
55	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
58	Yes	No	No	No	Yes	Yes	Yes	No	No
61	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
64	Yes	Yes	Yes	No	Yes	Yes	No	No	No
67	Yes	Yes	No	No	Yes	Yes	No	No	No
70	Yes	No	No	No	Yes	Yes	No	No	No
73	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
76	Yes	Yes	Yes	No	Yes	No	No	No	No
79	Yes	Yes	No	No	Yes	No	No	No	No
82	Yes	No	Yes	No	No	No	No	No	No
85	Yes	Yes	Yes	Yes	No	No	No	No	No
88	Yes	Yes	Yes	No	No	No	No	No	No
91	Yes	Yes	Yes	No	Yes	No	No	No	No
94	Yes	No	No	No	No	No	No	No	No
97	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
100	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
103	Yes	Yes	No	No	Yes	Yes	Yes	No	No
106	Yes	No	No	No	Yes	Yes	Yes	No	No
109	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
112	Yes	Yes	Yes	No	Yes	Yes	No	No	No
115	Yes	Yes	No	No	Yes	Yes	No	No	No
118	Yes	No	Yes	No	Yes	Yes	No	No	No
121	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
124	Yes	Yes	Yes	No	Yes	No	No	No	No
127	Yes	Yes	Yes	No	Yes	No	No	No	No
130	Yes	No	No	No	No	No	No	No	No
133	Yes	Yes	Yes	Yes	No	No	No	No	No
136	Yes	Yes	Yes	No	No	No	No	No	No
139	Yes	Yes	No	No	No	No	No	No	No
142	Yes	No	No	No	No	No	No	No	No

Table 11-1 STS Mapping Using CTC (continued)

Starting STS	STS-3c	STS-6c	STS-9c	STS-12c	STS-18c	STS-24c	STS-36c	STS-48c	STS-192c
145	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
148	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
151	Yes	Yes	No	No	Yes	Yes	Yes	No	No
154	Yes	No	Yes	No	Yes	Yes	Yes	No	No
157	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
160	Yes	Yes	Yes	No	Yes	Yes	No	No	No
163	Yes	Yes	Yes	No	Yes	Yes	No	No	No
166	Yes	No	No	No	Yes	Yes	No	No	No
169	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
172	Yes	Yes	Yes	No	Yes	No	No	No	No
175	Yes	Yes	No	No	Yes	No	No	No	No
178	Yes	No	No	No	No	No	No	No	No
181	Yes	Yes	Yes	Yes	Yes	No	No	No	No
184	Yes	Yes	Yes	No	Yes	No	No	No	No
187	Yes	Yes	No	No	Yes	No	No	No	No
190	Yes	No	No	No	Yes	No	No	No	No

11.2.2 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 11-2](#) shows the statuses that can appear in the Status column.

Table 11-2 ONS 15454 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 11-2 ONS 15454 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 destinations 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 destinations 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 destinations does not exist.
CONVERSION_PENDING	An existing circuit in a topology upgrade is set to this state. The circuit returns to the DISCOVERED state once the topology upgrade is complete. For more information about topology upgrades, see Chapter 12, “SONET 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 a topology upgrade fails. For more information about topology upgrades, see Chapter 12, “SONET Topologies and Upgrades.”
DROP_PENDING	A circuit is set to this status when a new circuit drop is being added.
ROLL_PENDING	A circuit roll is awaiting completion or cancellation.

11.2.3 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 In-Service and Normal (IS-NR) service state, the circuit service state is In-Service (IS).

- If all cross-connects in a circuit are in an Out-of-Service (OOS) service state, such as Out-of-Service and Management, Maintenance (OOS-MA,MT); Out-of-Service and Management, Disabled (OOS-MA,DSBLD); or Out-of-Service and Autonomous, Automatic In-Service (OOS-AU,AINS) service state, the circuit service state is Out-of-Service (OOS).
- PARTIAL is appended to the OOS circuit service state when circuit cross-connects state are mixed and not all in IS-NR. The OOS-PARTIAL state can occur during automatic or manual transitions between states. For example, OOS-PARTIAL appears if you assign the IS,AINS administrative state to a circuit with DS-1 or DS3XM cards as the source or destination. Some cross-connects transition to the IS-NR service state, while others transition to OOS-AU,AINS. OOS-PARTIAL 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 Troubleshooting Guide* for troubleshooting procedures. The OOS-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 in 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 drop port from the IS-NR service state to the OOS-MA,DSBLD service state; you must first put the port in the OOS-MA,MT state before changing it to the OOS-MA,DSBLD 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 OOS-AU,AINS service state after a signal is continuously received. When the cross-connects in a circuit are in the OOS-AU,AINS service state, the ONS 15454 monitors the cross-connects for an error-free signal. It changes the state of the circuit from OOS to IS or to OOS-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. Two common examples of state changes you see when provisioning circuits using CTC are:

- When assigning the IS,AINS administrative state to cross-connects in VT circuits and VT tunnels, the source and destination ports on the VT circuits remain in the OOS-AU,AINS service state until an alarm-free signal is received for the duration of the soak timer. When the soak timer expires and an alarm-free signal is found, the VT source port and destination port service states change to IS-NR and the circuit service state becomes IS.
- When assigning the IS,AINS administrative state to cross-connects in STS circuits, the circuit source and destination ports transition to the OOS-AU,AINS service state. When an alarm-free signal is received, the source and destination ports remain OOS-AU,AINS for the duration of the soak timer. After the port soak timer expires, STS source and destination ports change to IS-NR and the circuit service state changes to IS.

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 OOS-AU,AINS state and has a good signal, the Time Until IS column shows the soak count down status. If the port is OOS-AU,AINS 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 port and cross-connect states, see [Appendix B, “Administrative and Service States.”](#)

11.2.4 Circuit Protection Types

The Protection column in the Circuit window shows the card (line) and SONET topology (path) protection used for the entire circuit path. [Table 11-3](#) shows the protection type indicators that appear in this column.

Table 11-3 *Circuit Protection Types*

Protection Type	Description
1+1	The circuit is protected by a 1+1 protection group.
2F BLSR	The circuit is protected by a two-fiber BLSR.
4F BLSR	The circuit is protected by a four-fiber BLSR.
2F-PCA	The circuit is routed on a protection channel access (PCA) path on a two-fiber BLSR. PCA circuits are unprotected.
4F-PCA	The circuit is routed on a PCA path on a four-fiber BLSR. PCA circuits are unprotected.
BLSR	The circuit is protected by both a two-fiber and a four-fiber BLSR.
DRI	The circuit is protected by a dual-ring interconnection (DRI).
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 BLSRs. PCA circuits are unprotected.
Protected	The circuit is protected by diverse SONET topologies, for example, a BLSR and a path protection configuration, or a path protection configuration and 1+1 protection.
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 BLSR or 1+1 protection group.
Path Protection	The circuit is protected by a path protection.

Table 11-3 *Circuit Protection Types (continued)*

Protection Type	Description
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 Procedure Guide</i> .
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 Procedure Guide</i> .

11.2.5 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 [“11.7 Multiple Destinations for Unidirectional Circuits”](#) section on page 11-18.
- **Monitors**—Displays possible monitor sources and allows you to create a monitor circuit. For more information, see the [“11.8 Monitor Circuits”](#) section on page 11-18.
- **Path Protection Selectors**—Allows you to change path protection selectors. For more information, see the [“11.9 Path Protection Circuits”](#) section on page 11-19.
- **Path Protection Switch Counts**—Allows you to change path protection switch protection paths. For more information, see the [“11.9 Path Protection Circuits”](#) section on page 11-19.
- **State**—Allows you to edit cross-connect service states.
- **Merge**—Allows you to merge aligned circuits. For more information, see the [“11.20 Merged Circuits”](#) section on page 11-45.

Using the Export command from the File menu, you can export data from the Path Protection Selectors, Path Protection 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, STSs, and VTs through which a circuit passes, including slots and port numbers
- The circuit source and destination points
- Open Shortest Path First (OSPF) area IDs
- Link protection (path protection, unprotected, BLSR, 1+1) and bandwidth (OC-N)
- Provisionable patchcords between two cards on the same node or different nodes

For BLSRs, the detailed map shows the number of BLSR fibers and the BLSR ring ID. For path protection configurations, 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 DRI nodes. For VCAT circuits, the detailed map is not available for an entire VCAT circuit. However, you can view the detailed map to see 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 most severe alarm on port
- Loopbacks
- Path trace states
- Path selector states

By default, 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 D. Port states are indicated by colors, shown in [Table 11-4](#).

Table 11-4 Port State Color Indicators

Port Color	Service State
Green	IS-NR
Gray	OOS-MA,DSBLD
Violet	OOS-AU,AINS
Blue (Cyan)	OOS-MA,MT

In detailed view, a notation within or by the squares or selector pentagons indicates switches and loopbacks, including:

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

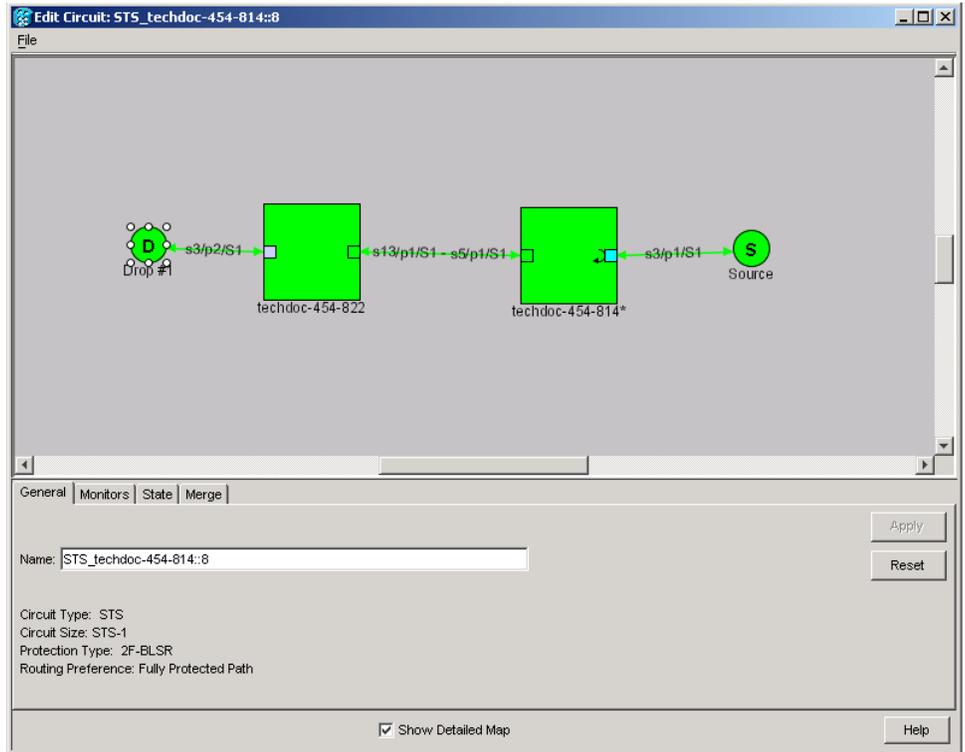
Move the mouse cursor over nodes, ports, and spans to see tooltips with information including the number of alarms on a node (organized by severity), the port service state, and the protection topology.

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 a path protection span to change the state of the path selectors in the path protection circuit.

[Figure 11-2](#) shows a circuit routed on a two-fiber BLSR. A port is shown in terminal loopback.

Figure 11-2 BLSR Circuit Displayed on the Detailed Circuit Map



11.3 Cross-Connect Card Bandwidth

The ONS 15454 XCVT, XC10G, and XC-VXC-10G cross-connect cards perform port-to-port, time-division multiplexing (TDM). XCVT, XC10G, and XC-VXC-10G cards perform STS, VT2 (XC-VXC-10G only), and VT1.5 multiplexing.

The STS matrix on the XCVT cross-connect card has a capacity for 288 STS terminations, and the XC10G and XC-VXC-10G cards each have a capacity for 1152 STS terminations. Because each STS circuit requires a minimum of two terminations, one for ingress and one for egress, the XCVT card has a capacity for 144 STS circuits, while the XC10G and XC-VXC-10G cards have a capacity for 576 STS circuits. However, this capacity is reduced at path protection and 1+1 nodes because three STS terminations are required at circuit source and destination nodes and four terminations are required at 1+1 circuit pass-through nodes. Path protection pass-through nodes only require two STS terminations.

The XCVT and XC10G cards perform VT1.5 multiplexing through 24 logical STS ports on the XCVT or XC10G VT matrix, and the XC-VXC-10G card performs VT1.5 and VT2 multiplexing through 96 logical STS ports on the XC-VXC-10G VT matrix. Each logical STS port can carry 28 VT1.5s or 21 VT2s. Subsequently, the VT matrix on the XCVT or XC10G has capacity for 672 VT1.5 terminations, or 336 VT1.5 circuits. The VT matrix on the XC-VXC-10G has capacity for 2688 VT1.5 terminations (1344 VT1.5 bidirectional circuits) or 2016 VT2 terminations (1008 VT2 bidirectional circuits). Every circuit requires two terminations, one for ingress and one for egress. However, this capacity is only achievable if:

- Every STS port on the VT matrix carries 28 VT1.5s or 21 VT2s.
- The node is in a BLSR or 1+1 protection scheme.

For example, if you create a VT1.5 circuit from an STS-1 on a drop card, two VT matrix STS ports are used, as shown in Figure 11-3. If you create a second VT1.5 circuit from the same STS port on the drop card, no additional logical STS ports are used on the VT matrix. In fact, you can create up to 28 VT1.5 circuits using the same STS-1 port. However, if the next VT1.5 circuit originates on a different STS, an additional pair of STS ports on the VT matrix is used, as shown in Figure 11-4. If you continued to create VT1.5 circuits on different EC-1 STSs and mapped each to an unused outbound STS, the VT matrix capacity would be reached after you created 12 VT1.5 circuits in the case of the XCVT or XC10G cards, or 48 VT1.5 circuits in the case of the XC-VXC-10G card.

Figure 11-3 One VT1.5 Circuit on One STS

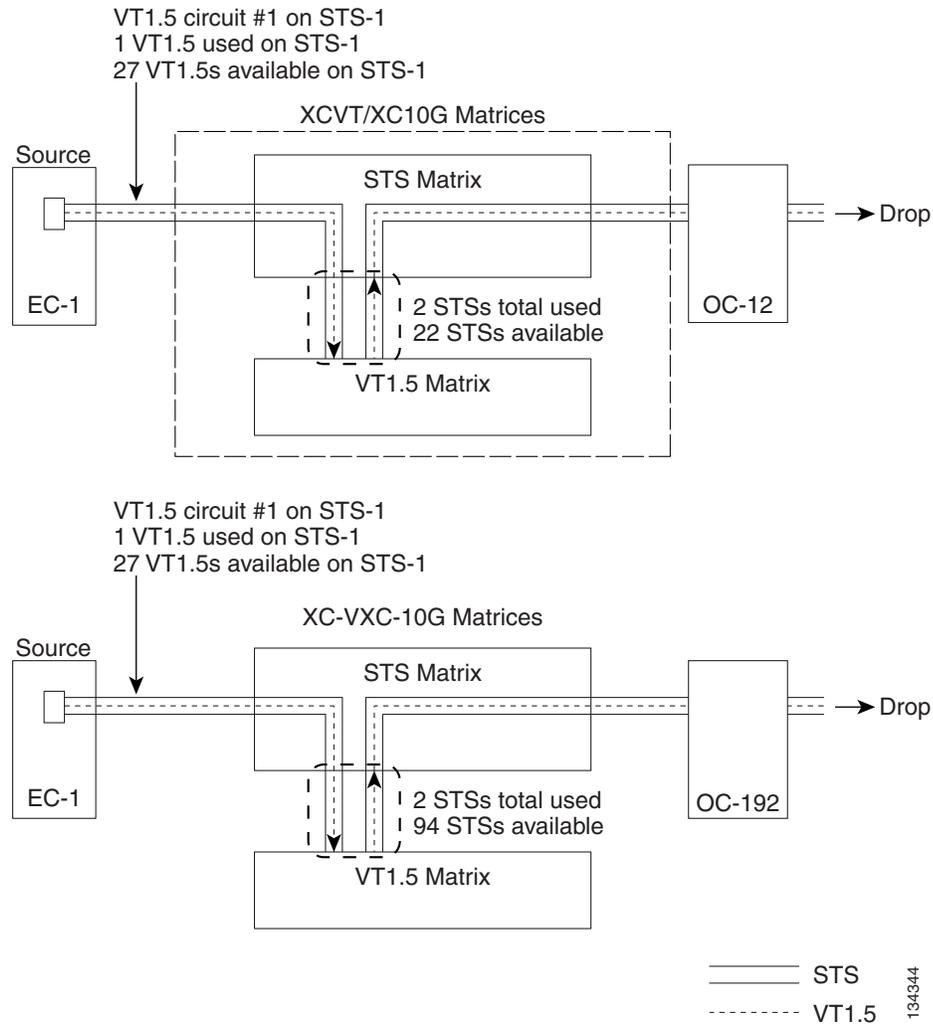
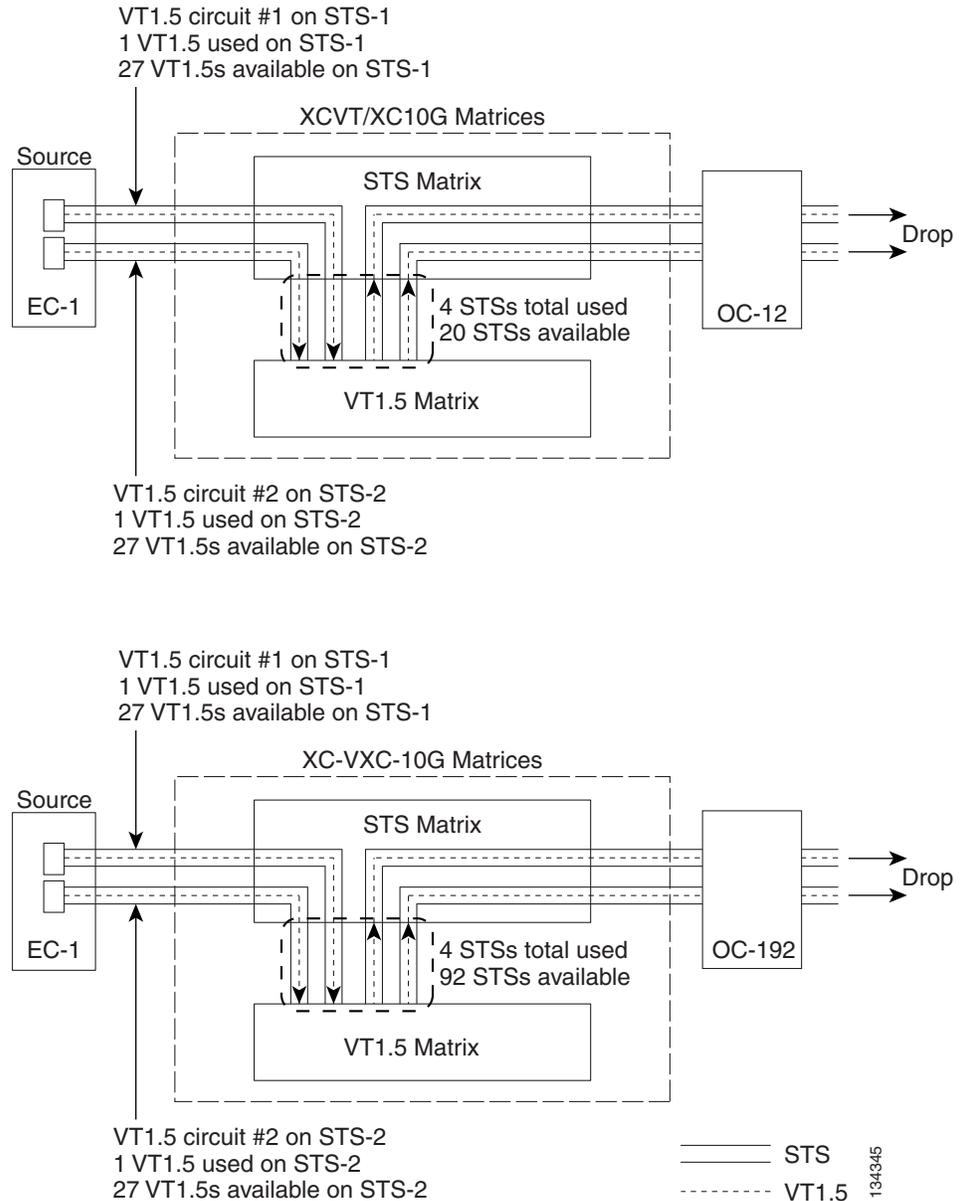


Figure 11-4 Two VT1.5 Circuits in a BLSR

**Note**

Circuits with DS1-14 and DS1N-14 circuit sources or destinations use one STS port on the VT matrix. Because you can only create 14 VT1.5 circuits from the DS-1 cards, 14 VT1.5s are unused on the VT matrix.

VT matrix capacity is also affected by SONET protection topology and node position within the circuit path. Matrix usage is slightly higher for path protection nodes than BLSR and 1+1 nodes. Circuits use two VT matrix ports at pass-through nodes if VT tunnels and aggregation points are not used. If the circuit is routed on a VT tunnel or an aggregation point, no VT matrix resources are used. [Table 11-5](#) shows basic STS port usage rates for VT 1.5 circuits.

Table 11-5 VT Matrix Port Usage for One VT1.5 Circuit

Node Type	No Protection	BLSR	Path Protection	1+1
Circuit source or destination node	2	2	3	2
Circuit pass-through node without VT tunnel	2	2	2	2
Circuit pass-through node with VT tunnel	0	0	0	0

Cross-connect card resources can be viewed on the Maintenance > Cross-Connect > Resource Usage tab. This tab shows:

- **STS-1 Matrix**—The percent of STS matrix resources that are used. 288 STSs are available on XCVT cards; 1152 are available on XC10G and XC-VXC-10G cards.
- **VT Matrix Ports**—The percent of the VT matrix ports (logical STS ports) that are used. 24 ports are available on XCVT and XC10G cards. 96 ports are available on the XC-VXC-10G card. The VT Port Matrix Detail shows the percent of each VT matrix port that is used.
- **VT Matrix**—The percent of the total VT matrix terminations that are used. There are 672 terminations for the XCVT and XC10G cards. 672 is the number of logical STS VT matrix ports (24) multiplied by the number of VT1.5s per port (28). There are 2688 terminations for the XC-VXC-10G card. 2688 is the number of logical STS VT matrix ports (96) multiplied by the number of VT1.5s per port (28).

To maximize resources on the cross-connect card VT matrix, keep the following points in mind as you provision circuits:

- Use all 28 VT1.5s on a given port or STS before moving to the next port or STS.
- Try to use EC-1, DS3XM, or OC-N cards as the VT1.5 circuit source and destination. VT1.5 circuits with DS-1-14 or DS1N-14 sources or destinations use a full port on the VT matrix even though only 14 VT1.5 circuits can be created.
- Use VT tunnels and VT aggregation points to reduce VT matrix utilization. VT tunnels allow VT1.5 circuits to bypass the VT matrix on pass-through nodes. They are cross-connected as STSs and only go through the STS matrix. VT aggregation points allow multiple VT1.5 circuits to be aggregated onto a single STS to bypass the VT matrix at the aggregation node.

11.4 Portless Transmux

The DS3XM-12 card provides a portless transmux interface to change DS-3s into VT1.5s. For XCVT drop slots, the DS3XM-12 card provides a maximum of 6 portless transmux interfaces; for XCVT trunk slots and XC10G or XC-VXC-10G slots, the DS3XM-12 card provides a maximum of 12 portless transmux interfaces. If two ports are configured as portless transmux, CTC allows you to create a DS3/STS1 circuit using one of these ports as the circuit end point. You can create separate DS1/VT1.5 circuits (up to 28) using the other port in this portless transmux pair.

When creating a circuit through the DS3XM-12 card, the portless pair blocks the mapped physical port(s); CTC does not display a blocked physical port in the source or destination drop-down list during circuit creation. [Table 11-6](#) lists the portless transmux mapping for XCVT drop ports.

Table 11-6 Portless Transmux Mapping for XCVT Drop Ports

Physical Port	Portless Port Pair
1, 2	13, 14
3, 4	15, 16
5, 6	17, 18
7, 8	19, 20
9, 10	21, 22
11, 12	23, 24

Table 11-7 lists the portless transmux for XCVT trunk ports and for XC10G or XC-VXC-10G any-slot ports.

Table 11-7 Portless Transmux Mapping for XCVT Trunk and XC10G/XC-VXC-10G Any-Slot Ports

Physical Port	Portless Port Pair
1	13, 14
2	25, 26
3	15, 16
4	27, 28
5	17, 18
6	29, 30
7	19, 20
8	31, 32
9	21, 22
10	33, 34
11	23, 24
12	35, 36

11.5 DCC Tunnels

SONET provides four DCCs for network element (NE) operation, administration, maintenance, and provisioning (OAM&P): one on the SONET Section layer (DCC1) and three on the SONET Line layer (DCC2, DCC3, and DCC4). The ONS 15454 uses the Section DCC (SDCC) for ONS 15454 management and provisioning. An 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 balances traffic over the existing DCC channels using a load balancing algorithm. This algorithm chooses a DCC for packet transport by considering packet size and DCC utilization. You can tunnel third-party SONET equipment across ONS 15454 networks using one of two tunneling methods: a traditional DCC tunnel or an IP-encapsulated tunnel.

11.5.1 Traditional DCC Tunnels

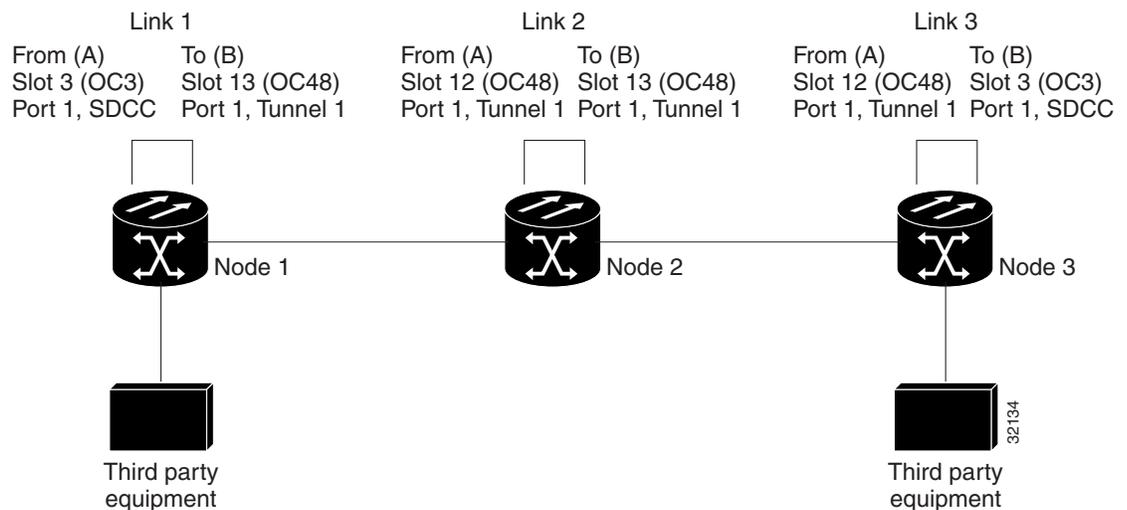
In traditional DCC tunnels, you can use the three LDCCs and the SDCC (when not used for ONS 15454 DCC terminations). A traditional DCC tunnel endpoint is defined by slot, port, and DCC, where DCC can be either the SDCC or one of the LDCCs. You can link LDCCs to LDCCs and link SDCCs to SDCCs. You can also link an SDCC to an LDCC, and an LDCC to an SDCC. To create a DCC tunnel, you connect the tunnel endpoints from one ONS 15454 optical port to another. Cisco recommends a maximum of 84 DCC tunnel connections for an ONS 15454. Table 11-8 shows the DCC tunnels that you can create using different OC-N cards.

Table 11-8 DCC Tunnels

Card	DCC	SONET Layer	SONET Bytes
OC3 IR 4/STM1 SH 1310	DCC1	Section	D1 - D3
OC3 IR/STM1 SH 1310-8; all OC-12, OC-48, and OC-192 cards	DCC1	Section	D1 - D3
	DCC2	Line	D4 - D6
	DCC3	Line	D7 - D9
	DCC4	Line	D10 - D12

Figure 11-5 shows a DCC tunnel example. Third-party equipment is connected to OC-3 cards at Node 1/Slot 3/Port 1 and Node 3/Slot 3/Port 1. Each ONS 15454 node is connected by OC-48 trunk (span) cards. In the example, three tunnel connections are created, one at Node 1 (OC-3 to OC-48), one at Node 2 (OC-48 to OC-48), and one at Node 3 (OC-48 to OC-3).

Figure 11-5 Traditional DCC Tunnel



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

- Each ONS 15454 can have up to 84 DCC tunnel connections.
- Each ONS 15454 can have up to 84 Section DCC terminations.
- A SDCC that is terminated cannot be used as a DCC tunnel endpoint.
- A SDCC that is used as an DCC tunnel endpoint cannot be terminated.

- All DCC tunnel connections are bidirectional.

11.5.2 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. To compare traditional DCC tunnels with IP-encapsulated tunnels, 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 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.

11.6 SDH Tunneling

The Cisco ONS 15454 SONET MSPP provides a SDH traffic transport solution with scalable SONET, data or DWDM multiservice capabilities. The SDH traffic is aggregated and transported across an ONS 15454 network, similar to the SONET TDM and data services. STM-1 to STM-64 payloads are transported over SONET from any port on a Cisco ONS 15454 OC-N card provisioned to support SDH signals. For more information on SDH tunneling, refer to the "[SDH Tunneling Over Cisco ONS 15454 SONET MSPP Systems](#)" Application Note.

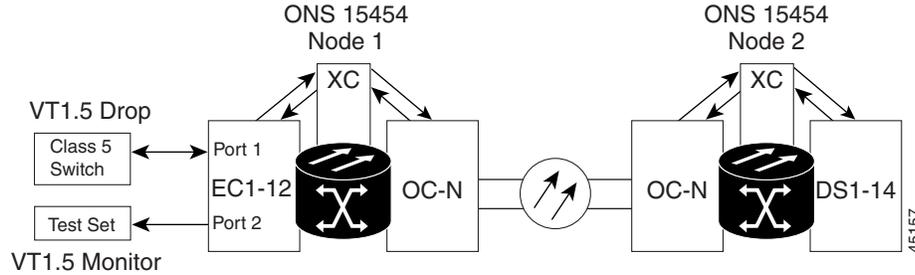
11.7 Multiple Destinations for Unidirectional Circuits

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 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 the profile to the port that does not have its Rx input terminated.

11.8 Monitor Circuits

Monitor circuits are secondary circuits that monitor traffic on primary bidirectional circuits. [Figure 11-6](#) shows an example of a monitor circuit. At Node 1, a VT1.5 is dropped from Port 1 of an EC1-12 card. To monitor the VT1.5 traffic, plug test equipment into Port 2 of the EC1-12 card and provision a monitor circuit to Port 2. Circuit monitors are one-way. The monitor circuit in [Figure 11-6](#) monitors VT1.5 traffic received by Port 1 of the EC1-12 card.

Figure 11-6 VT1.5 Monitor Circuit Received at an EC1-12 Port

11.8.1 Monitor Circuits using portless ports as a source on DS3XM-12

With STS bi-directional circuit between source and destination and with VT-MAPPED at “circuit source” as an option using DS3XM-12 in 15454 SONET platforms, two circuits will get created:

- 1) src->even portless port
- 2) odd portless port->dest

Traffic flow from source to dest is as follows:

src->even port->odd port->dest

While creating Monitor circuit using Even portless port as source, the circuit will get created in the direction towards:

dest-> odd portless port-> Even portless port -> dest Monitored port

While creating Monitor circuit using odd portless port as source, the circuit will get created in the direction towards:

src-> Even portless port-> odd portless port -> dest Monitored port



Note

Monitor circuits cannot be used with Ethernet circuits.

11.9 Path Protection Circuits

Use the Edit Circuits window to change path protection selectors and switch protection paths (Figure 11-7). In the Path Protection Selectors subtab in the Edit Circuits window, you can:

- View the path protection circuit’s working and protection paths.
- Edit the reversion time.
- Set the hold-off timer.
- Edit the Signal Fail/Signal Degrade thresholds.
- Change payload defect indication path (PDI-P) settings.

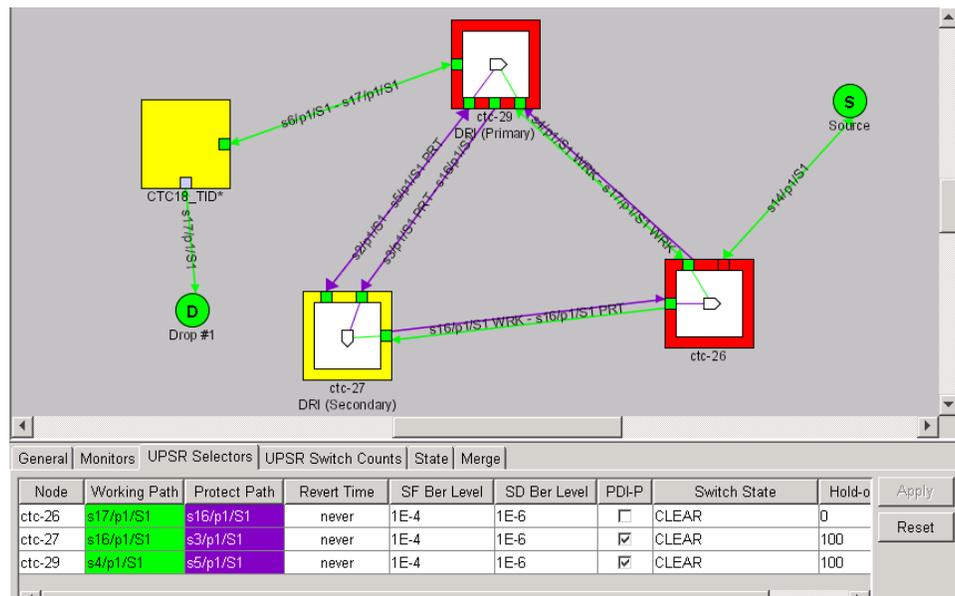
**Note**

The XC-VXC-10G cross-connect card supports VT switching based on SF and SD bit error rate (BER) thresholds. The XC10G and XCVT cross-connect cards do not support VT switching based on SF and SD BER thresholds, and hence, in the path protection Selectors tab, the SF BER Level and SD BER Level columns display “N/A” for these cards.

In the Path Protection Switch Counts subtab, you can:

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

Figure 11-7 Editing Path Protection Selectors



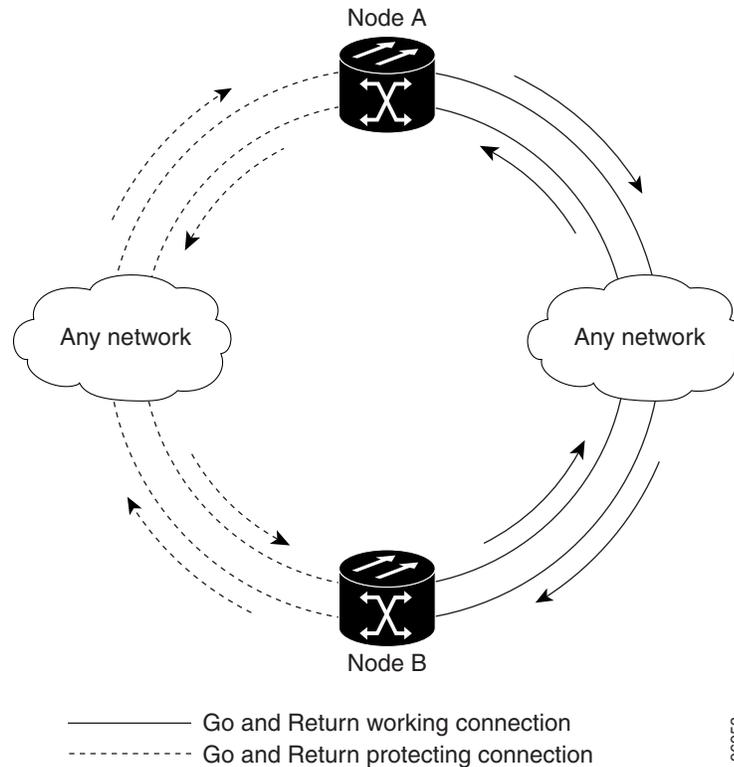
11.9.1 Open-Ended Path Protection Circuits

If ONS 15454s are connected to a third-party network, you can create an open-ended path protection circuit to route a circuit through it. To do this, you create four circuits. One circuit is created on the source ONS 15454 network. This circuit has one source and two destinations, each destination provisioned to the ONS 15454 interface that is connected to the third-party network. The second and third circuits are created on the third-party network so that the circuit travels across the network on two diverse paths to the far end ONS 15454. At the destination node, the fourth circuit is created with two sources, one at each node interface 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 path protection circuit.

11.9.2 Go-and-Return Path Protection Routing

The go-and-return path protection routing option allows you to route the path protection working path on one fiber pair and the protect path on a separate fiber pair (Figure 11-8). 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 path protection circuits. The go-and-return option appears in the Circuit Attributes panel of the Circuit Creation wizard.

Figure 11-8 Path Protection Go-and-Return Routing



11.10 BLSR Protection Channel Access Circuits

You can provision circuits to carry traffic on BLSR protection channels when conditions are fault-free. Traffic routed on BLSR 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 OC-48 BLSR, STSs 25 to 48 can carry extra traffic when no ring switches are active, but PCA circuits on these STSs 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 BLSR is provisioned as revertive, this occurs automatically after the fault conditions are cleared and the reversion timer has expired.

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

- If BLSRs are provisioned as nonrevertive, PCA circuits are not restored automatically after a ring or span switch. You must switch the BLSR manually.
- PCA circuits are routed on working channels when you upgrade a BLSR from a two-fiber to a four-fiber or from one optical speed to a higher optical speed. For example, if you upgrade a two-fiber OC-48 BLSR to an OC-192, STSs 25 to 48 on the OC-48 BLSR become working channels on the OC-192 BLSR.

11.11 BLSR STS and VT Squelch Tables

ONS 15454 nodes display STS and VT squelch tables depending on the type of circuits created. For example, if a fiber cut occurs, the BLSR squelch tables show STSs or VTs that will be squelched for every isolated node. Squelching replaces traffic by inserting the appropriate alarm indication signal path (AIS-P) and prevents traffic misconnections. For an STS with a VT-access check mark, the AIS-P will be removed after 100 ms. To view the squelch tables, refer to the “Manage Circuits” chapter in the *Cisco ONS 15454 Procedure Guide* for detailed instructions. For more information about BLSR squelching, refer to Telcordia GR-1230.

11.11.1 BLSR STS Squelch Table

BLSR STS squelch tables show STSs that will be squelched for every isolated node.

The BLSR Squelch Table window displays the following information:

- **STS Number**—Shows the BLSR STS numbers. For two-fiber BLSRs, the number of STSs is half the BLSR OC-N, for example, an OC-48 BLSR squelch table will show 24 STSs. For four-fiber BLSRs, the number of STSs in the table is the same as the BLSR OC-N.
- **West Source**—If traffic is received by the node on its west span, the BLSR node ID of the source appears. (To view the BLSR node IDs for all nodes in the ring, click the Ring Map button.)
- **West VT (from the West Source)** — A check mark indicates that the STS carries incoming VT traffic. The traffic source is coming from the west side.
- **West VT (from the West Destination)** — A check mark indicates that the STS carries outgoing VT traffic. The traffic is dropped on the west side.
- **West Dest**—If traffic is sent on the node’s west span, the BLSR node ID of the destination appears.
- **East Source**—If traffic is received by the node on its east span, the BLSR node ID of the source appears.
- **East VT — (from the East Source)** - A check mark indicates that the STS carries incoming VT traffic. The traffic source is coming from the east side.
- **East VT — (from the East Destination)** - A check mark indicate that the STS carries outgoing VT traffic. The traffic is dropped on the east side.
- **East Dest**—If traffic is sent on the node’s east span, the BLSR node ID of the destination appears.



Note

BLSR squelching is performed on STSs that carry STS circuits only. Squelch table entries will not appear for STSs carrying VT circuits or Ethernet circuits to or from E-Series Ethernet cards provisioned in a multiscard Ethergroup.

11.11.2 BLSR VT Squelch Table

BLSR VT squelch tables only appear on the node dropping VTs from a BLSR and are used to perform VT-level squelching when a node is isolated. VT squelching is supported on the ONS 15454 and the ONS 15327 platforms. The ONS 15600 platform does not support VT squelching; however, when an ONS 15454 and an ONS 15600 are in the same network, the ONS 15600 node allows the ONS 15454 node to carry VT circuits in a VT tunnel. The ONS 15600 performs 100-ms STS-level squelching for each VT-access STS at the switching node in case of a node failure.

When using a VT circuit on a VT tunnel (VTT), the VTT allows multiple VT circuits to be passed through on a single STS without consuming VT matrix resources on the cross-connect card. Both endpoints of the VTT are the source and destination nodes for the VTT. The node carrying VT circuits through a VTT is called a VT-access node. In case of a source and destination node failure of the VTT, the switching node performs 100-ms STS-level squelching for the VTT STS. The node dropping VT traffic performs VT-level squelching. VT traffic on the VTT that is not coming from the failed node is protected.

When using a VT circuit on a VT aggregation point (VAP), the VAP allows multiple VT circuits to be aggregated into a single STS without consuming VT matrix resources on the cross-connect card. The source for each VAP STS timeslot is the STS-grooming end where VT1.5 circuits are aggregated into a single STS. The destination for each VAP STS is the VT-grooming end where VT1.5 circuits originated. The source node for each VT circuit on a VAP is the STS-grooming end where the VT1.5 circuits are aggregated into a single STS. The STS grooming node is not a VT-access node. The non VT-access node performs STS-level squelching for each STS timeslot at the switching node in case the VT-grooming node fails. The node dropping VT traffic performs VT-level squelching for each VT timeslot in case the STS-grooming end node fails. No VT traffic on the VAP is protected during a failure of the STS-grooming node or the VT-grooming node.

To view the VT squelch table, double-click the VT with a check mark in the BLSR STS squelch table window. The check mark appears on every VT-access STS; however, the VT-squelch table appears only by double-clicking the check mark on the node dropping the VT. The intermediate node of the VT does not maintain the VT-squelch table.

The VT squelch table provides the following information:

- **VT Number**—Shows the BLSR VT numbers. The VT number includes VT group number and VT number in group (VT group 2 and channel 1 are displayed as 2-1.)
- **West Source**—If traffic is received by the node on its west span, the BLSR node ID of the source appears. (To view the BLSR node IDs for all nodes in the ring, click the Ring Map button.)
- **East Source**—If traffic is received by the node on its east span, the BLSR node ID of the source appears.

11.12 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.

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.

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*.

11.13 Section and Path Trace

SONET J0 section and J1 and J2 path trace are repeated, fixed-length strings composed of 16 or 64 consecutive bytes. You can use the strings to monitor interruptions or changes to circuit traffic.

The OC192-XFP and MRC-12 cards support J0 section trace. [Table 11-9](#) shows the ONS 15454 cards that support J1 path trace. DS-1 and DS-3 cards can transmit and receive the J1 field, while the EC-1, OC-3, OC-48 AS, and OC-192 can only receive the J1 bytes. Cards that are not listed in the table do not support the J1 byte. The DS3XM-12 card supports J2 path trace for VT circuits.

Table 11-9 ONS 15454 Cards Capable of J1 Path Trace

J1 Function	Cards
Transmit and Receive	CE-Series DS1-14 ¹ DS1N-14 DS1/EC1-56 DS3-12E DS3i-N-12 DS3/EC1-48 DS3N-12E DS3XM-6 DS3XM-12 FC_MR-4 G-Series ML-Series
Receive Only	EC1-12 OC3 IR 4/STM1 SH 1310 OC3 IR 4/STM1 SH 1310-8 OC12/STM4-4 OC48 IR/STM16 SH AS 1310 OC48 LR/STM16 LH AS 1550 OC192 SR/STM64 IO 1310 OC192 LR/STM64 LH 1550 OC192 IR/STM SH 1550 OC192-XFP

1. J1 path trace is not supported for DS-1s used in VT circuits.

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.

11.14 Path Signal Label, C2 Byte

One of the overhead bytes in the SONET frame is the C2 byte. The SONET 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 STS path overhead (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. C2 byte hex values are provided in

Table 11-10.

Table 11-10 STS Path Signal Label Assignments for Signals

Hex Code	Content of the STS Synchronous Payload Envelope (SPE)
0x00	Unequipped
0x01	Equipped - nonspecific payload
0x02	VT structured STS-1 (DS-1)
0x03	Locked VT mode
0x04	Asynchronous mapping for DS-3
0x12	Asynchronous mapping for DS4NA
0x13	Mapping for Asynchronous Transfer Mode (ATM)
0x14	Mapping for distributed queue dual bus (DQDB)
0x15	Asynchronous mapping for fiber distributed data interface (FDDI)
0x16	High-level data link control (HDLC) over SONET mapping
0x1B	Generic Frame Procedure (GFP) used by the FC_MR-4 and ML Series cards
0xFD	Reserved
0xFE	0.181 test signal (TSS1 to TSS3) mapping SDH network
0xFF	Alarm indication signal, path (AIS-P)

If a circuit is provisioned using a terminating card, the terminating card provides the C2 byte. A VT circuit is terminated at the XCVT, XC10G, or XC-VXC-10G card, which generates the C2 byte (0x02) downstream to the STS terminating cards. The XCVT, XC10G, or XC-VXC-10G card generates the C2 value (0x02) to the DS1 or DS3XM terminating card. If an optical 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 optical circuit usually fixes a circuit having C2 byte problems. [Table 11-11](#) lists label assignments for signals with payload defects.

Table 11-11 STS Path Signal Label Assignments for Signals with Payload Defects

Hex Code	Content of the STS SPE
0xE1	VT-structured STS-1 SPE with 1 VTx payload defect (STS-1 with 1 VTx PD)
0xE2	STS-1 with 2 VTx PDs
0xE3	STS-1 with 3 VTx PDs
0xE4	STS-1 with 4 VTx PDs
0xE5	STS-1 with 5 VTx PDs
0xE6	STS-1 with 6 VTx PDs
0xE7	STS-1 with 7 VTx PDs
0xE8	STS-1 with 8 VTx PDs
0xE9	STS-1 with 9 VTx PDs
0xEA	STS-1 with 10 VTx PDs

Table 11-11 STS Path Signal Label Assignments for Signals with Payload Defects (continued)

Hex Code	Content of the STS SPE
0xEB	STS-1 with 11 VTx PDs
0xEC	STS-1 with 12 VTx PDs
0xED	STS-1 with 13 VTx PDs
0xEE	STS-1 with 14 VTx PDs
0xEF	STS-1 with 15 VTx PDs
0xF0	STS-1 with 16 VTx PDs
0xF1	STS-1 with 17 VTx PDs
0xF2	STS-1 with 18 VTx PDs
0xF3	STS-1 with 19 VTx PDs
0xF4	STS-1 with 20 VTx PDs
0xF5	STS-1 with 21 VTx PDs
0xF6	STS-1 with 22 VTx PDs
0xF7	STS-1 with 23 VTx PDs
0xF8	STS-1 with 24 VTx PDs
0xF9	STS-1 with 25 VTx PDs
0xFA	STS-1 with 26 VTx PDs
0xFB	STS-1 with 27 VTx PDs
0xFC	VT-structured STS-1 SPE with 28 VT1.5 (Payload defects or a non-VT-structured STS-1 or STS-Nc SPE with a payload defect.)
0xFF	Reserved

11.15 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 path protection configuration. 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 SONET topology (path protection, BLSR, 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. VT tunnels are preferable for VT circuits because VT 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 VT-capable nodes and VT-incapable nodes, CTC can automatically create a VT tunnel. Otherwise, CTC asks you whether a VT tunnel is needed.
- To create protected circuits between topologies, install an XCVT, XC10G, or XC-VXC-10G cross-connect card on the shared node.
- For STS circuits, you can use portless transmux interfaces if a DS3XM-12 card is installed in the network. CTC automatically routes the circuit over the portless transmux interfaces on the specified node creating an end-to-end STS circuit.

**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.”](#)

11.15.1 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 BLSRs.

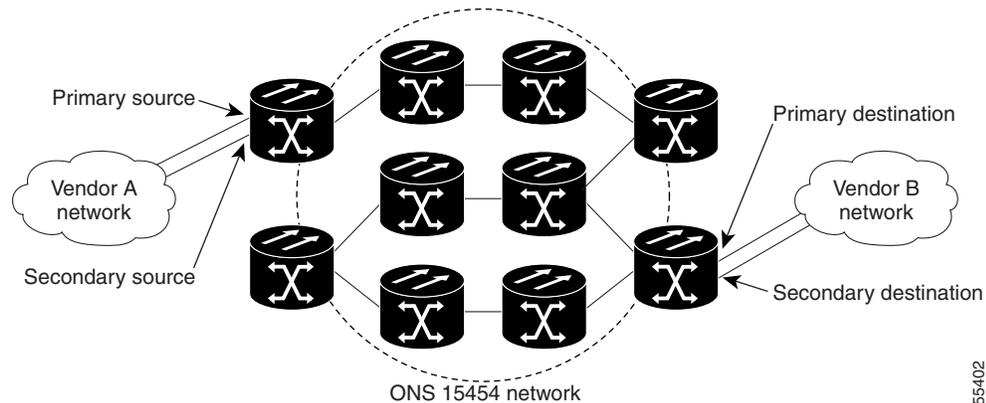
If CTC cannot find a link that meets these requirements, an error appears.

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

11.15.2 Secondary Sources and Destinations

CTC supports secondary circuit sources and destinations (drops). Secondary sources and destinations typically interconnect two third-party networks, as shown in [Figure 11-9](#). Traffic is protected while it goes through a network of ONS 15454s.

Figure 11-9 Secondary Sources and Destinations



55402

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.
- The sources and destinations cannot be DS-3, DS3XM, or DS-1-based STS-1s or VT1.5s.
- Secondary sources and destinations are permitted only for regular STS/VT1.5 connections (not for VT tunnels and multcard EtherSwitch circuits).
- For point-to-point (straight) Ethernet circuits, only SONET STS endpoints can be specified as multiple sources or destinations.

For bidirectional circuits, CTC creates a path protection connection at the source node that allows traffic to be selected from one of the two sources on the ONS 15454 network. If you check the Fully Path Protected option during circuit creation, traffic is protected within the ONS 15454 network. At the destination, another path protection connection is created to bridge traffic from the ONS 15454 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, a path protection drop-and-continue connection is created at the source node.

11.16 Manual Circuit Routing

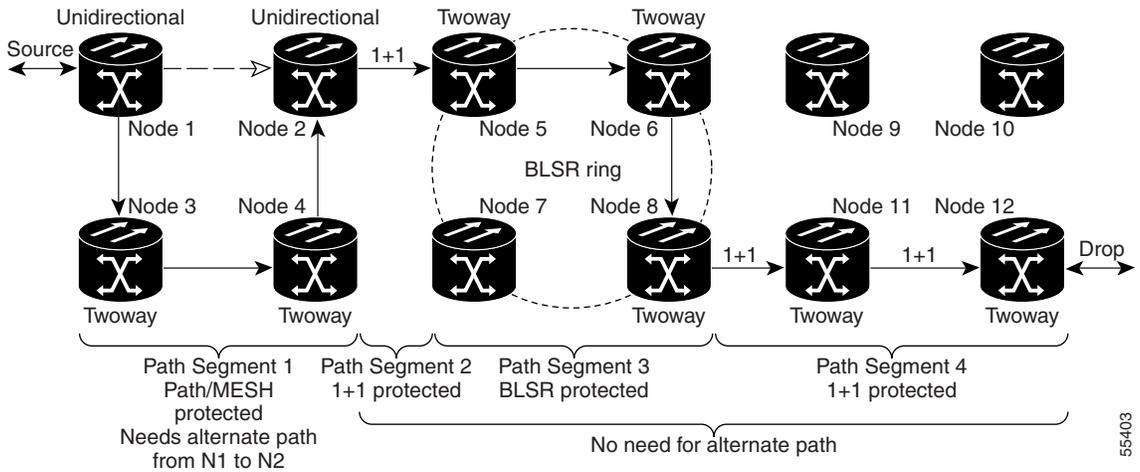
Routing circuits manually allows you to:

- Choose a specific path, not necessarily the shortest path.
- Choose a specific STS/VT1.5 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 path protection segments.

CTC imposes the following rules on manual routes:

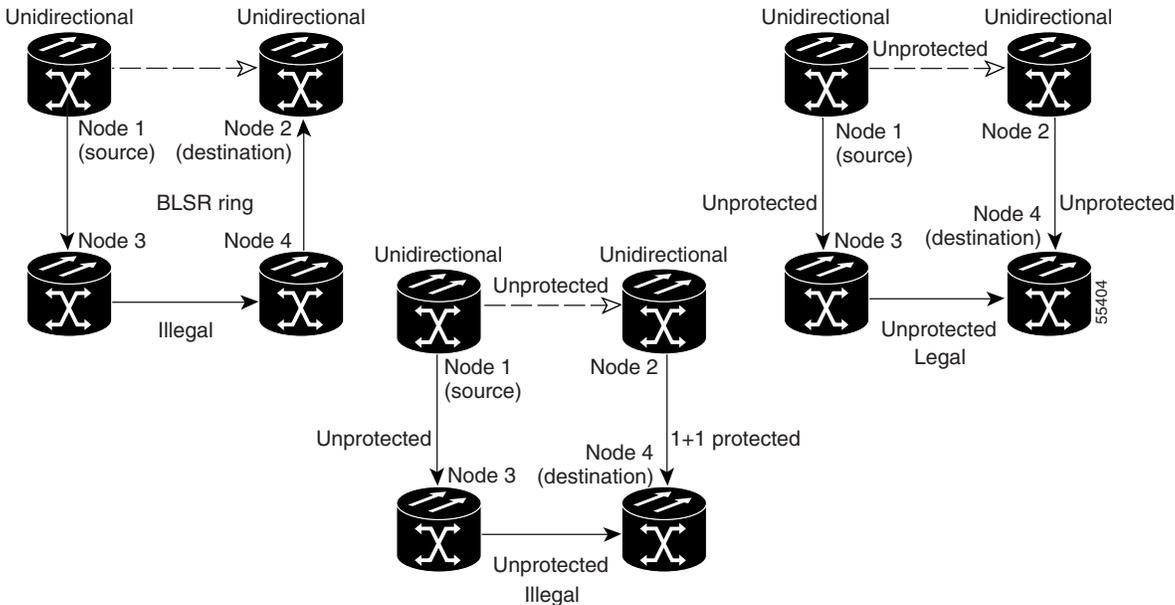
- All circuits, except multcard EtherSwitch circuits in a shared packet ring, should have links with a direction that flows from source to destination. This is true for multcard EtherSwitch circuits that are not in a shared packet ring.
- If you enabled Fully Path Protected, choose a diverse protect (alternate) path for every unprotected segment (Figure 11-10).

Figure 11-10 *Alternate Paths for Virtual Path Protection Segments*



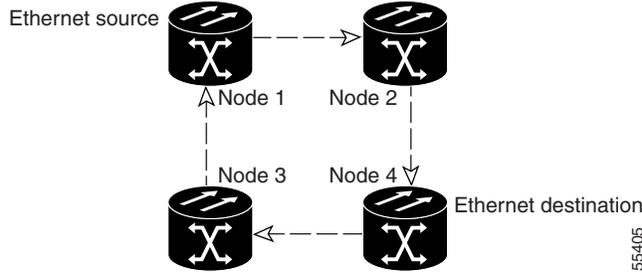
- For multcard EtherSwitch circuits, the Fully Path Protected option is ignored.
- For a node that has a path protection selector based on the links chosen, the input links to the path protection selectors cannot be 1+1 or BLSR protected (Figure 11-11). The same rule applies at the path protection bridge.

Figure 11-11 *Mixing 1+1 or BLSR Protected Links With a Path Protection Configuration*



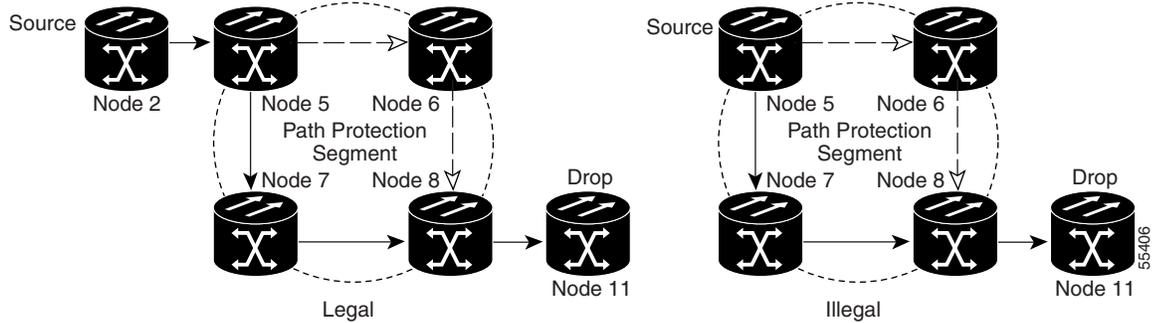
- In a shared packet ring, choose the links of multcard EtherSwitch circuits to route from source to destination back to source (Figure 11-12). Otherwise, a route (set of links) chosen with loops is invalid.

Figure 11-12 Ethernet Shared Packet Ring Routing



- Multicard EtherSwitch circuits can have virtual path protection segments if the source or destination is not in the path protection domain. This restriction also applies after circuit creation; therefore, if you create a circuit with path protection segments, Ethernet destinations cannot exist anywhere on the path protection segment (Figure 11-13).

Figure 11-13 Ethernet and Path Protection



- A VT tunnel cannot be the endpoint of a path protection segment. A path protection segment endpoint is where the path protection 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 11-12 through Table 11-15 on page 11-32 summarize the available node connections. Any other combination is invalid and generates an error.

Table 11-12 Bidirectional STS/VT/Regular Multicard EtherSwitch/Point-to-Point (Straight) Ethernet Circuits

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Destinations
Path protection	—	2	1	—
Path protection	2	—	—	1
Path protection	2	1	—	—
Path protection	1	2	—	—
Path protection	1	—	—	2
Path protection	—	1	2	—
Double path protection	2	2	—	—

Table 11-12 Bidirectional STS/VT/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 Destinations
Double path protection	2	—	—	2
Double path protection	—	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 11-13 Unidirectional STS/VT Circuit

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Destinations
One way	1	1	—	—
Path protection head end	1	2	—	—
Path protection head end	—	2	1	—
Path protection drop and continue	2	—	—	1+

Table 11-14 Multicard Group Ethernet Shared Packet Ring Circuit

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Destinations
At Intermediate Nodes Only				
Double path protection	2	2	—	—
Two way	1	1	—	—
At Source or Destination Nodes Only				
Ethernet	1	1	—	—

Table 11-15 Bidirectional VT Tunnels

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Destinations
At Intermediate Nodes Only				
Path protection	2	1	—	—
Path protection	1	2	—	—
Double path protection	2	2	—	—

Table 11-15 Bidirectional VT Tunnels (continued)

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Destinations
Two way	1	1	—	—
At Source Nodes Only				
VT tunnel endpoint	—	1	—	—
At Destination Nodes Only				
VT tunnel endpoint	1	—	—	—

Although virtual path protection segments are possible in VT tunnels, VT tunnels are still considered unprotected. If you need to protect VT circuits, use two independent VT tunnels that are diversely routed or use a VT tunnel that is routed over 1+1, BLSR, or a mixture of 1+1 and BLSR links.

11.17 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 that CTC calculates for the circuit. If the network is composed entirely of BLSR or 1+1 links, or the path between source and destination can be entirely protected using 1+1 or BLSR links, no path-protected mesh network (PPMN), or virtual path protection, protection is used.

If PPMN protection is needed to protect the path, set the level of node diversity for the PPMN portions of the complete path in the Circuit Routing Preferences area of the Circuit Creation dialog box:

- **Nodal Diversity Required**—Ensures that the primary and alternate paths of each PPMN 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 PPMN domain in the complete path.
- **Link Diversity Only**—Creates only a link-diverse path for each PPMN 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 achieve the following results:

- **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 and/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 that 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 PPMN domains/segments. The alternate path is computed normally; CTC uses excluded nodes/links when finding all primary and alternate paths on PPMNs.

11.18 Virtual Concatenated Circuits

Virtual concatenated (VCAT) circuits, also called VCAT groups (VCGs), transport traffic using noncontiguous TDM time slots, avoiding the bandwidth fragmentation problem that exists with contiguous concatenated (CCAT) circuits. The cards that support VCAT circuits are the CE-Series, FC_MR-4 (both line rate and enhanced 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 and terminate at the same end points. For two-fiber BLSR configurations, some members can be routed on protected time slots and others on PCA time slots.

To enable end-to-end connectivity in a VCAT circuit that traverses through a third-party network, you can use Open-Ended VCAT circuit creation, or you can create a server trail between the ports. For more details, refer to the “Create Circuits and VT Tunnels” chapter in the *Cisco ONS 15454 Procedure Guide*.

11.18.1 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 in the Edit Circuits window.

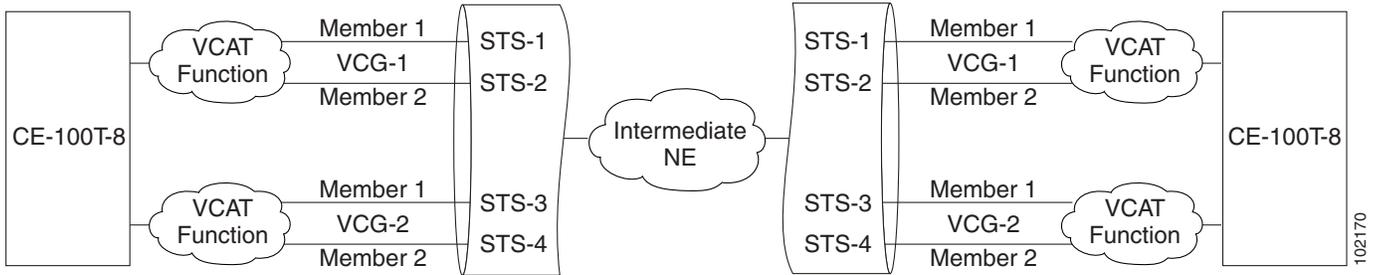
- If all member circuits are in the IS state, the VCAT circuit state is IS.
- If all In Group member circuits are in the OOS state, the VCAT circuit state is OOS.
- If no member circuits exist or if all member circuits are Out of Group, the VCAT circuit state is OOS.
- A VCAT circuit is in OOS-PARTIAL state when In Group member states are mixed and not all are in the IS state.

11.18.2 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.

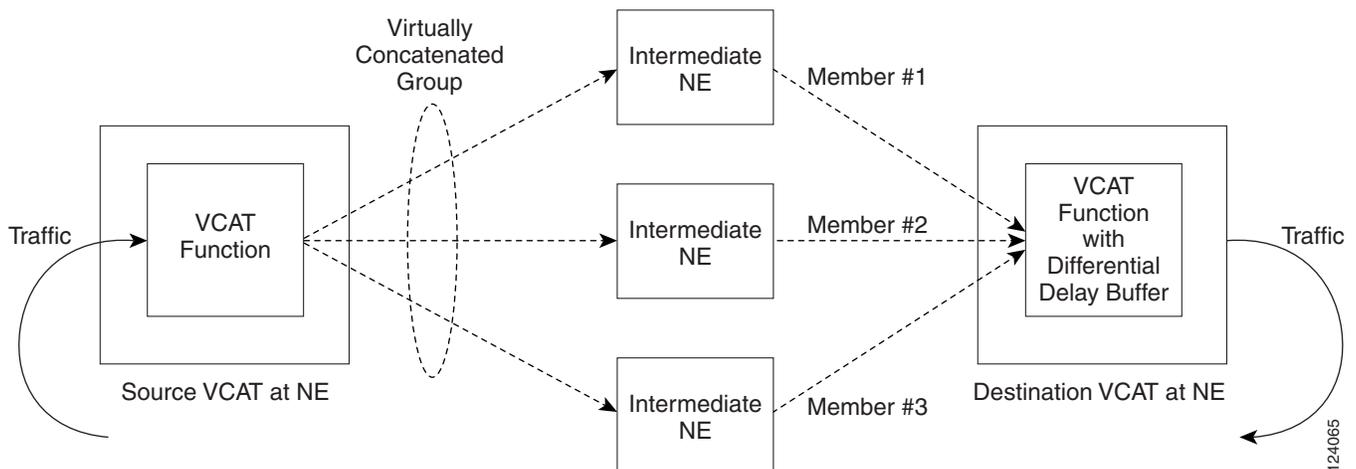
Two types of automatic and manual routing are available for VCAT members: common fiber routing and split routing. CE-Series, FC_MR-4 (both line rate and enhanced mode), and ML-Series cards support common fiber routing. In common fiber routing, all VCAT members travel on the same fibers, which eliminates delay between members. Three protection options are available for common fiber routing: Fully Protected, PCA, and Unprotected. [Figure 11-14](#) shows an example of common fiber routing.

Figure 11-14 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 11-15 shows an example of split fiber routing.

Figure 11-15 VCAT Split Fiber Routing



In both common fiber and split fiber routing, each member can use a different protection scheme; however, for common fiber routing, CTC checks the combination to make sure that a valid route exists. If it does not, the user must modify the protection type. In both common fiber and split fiber routing, intermediate nodes treat the VCAT members as normal circuits that are independently routed and protected by the SONET network. At the terminating nodes, these member circuits are multiplexed into a contiguous stream of data.

The switch time for split fiber routing depends on the type of circuits traversing the path.

- CCAT circuits will carry traffic after the SONET defects are cleared.
- VCAT circuits will carry traffic after the SONET defects are cleared and VCAT framers are in frame for ALL the time slots that are part of the group. Hence the switchover takes extra time.
- LCAS circuits carry traffic after the SONET defects are cleared, and the VCAT framers are in frame for any time slots that are part of the group, and the LCAS protocol has fed back MST=OK (MST=Member Status) to the far end so the far end can enable the time slot to carry traffic.

Table 11-16 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 11-16](#) 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.

11.18.3 Link Capacity Adjustment

The CE-100T-8 card 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.



Note Although LCAS operations are errorless, a SONET error can affect one or more VCAT members. If this occurs, the VCAT Group Degraded (VCG-DEG) alarm is raised. For information on clearing this alarm, refer to the *Cisco ONS 15454 Troubleshooting Guide*.

Instead of LCAS, the FC_MR-4 (enhanced mode), CE-1000-4 card, CE-MR-10, and ML-Series cards support software LCAS (SW-LCAS). SW-LCAS is a limited form of LCAS that allows the VCAT circuit to adapt to member failures and keep traffic flowing at a reduced bandwidth. SW-LCAS uses legacy SONET failure indicators like AIS-P and remote defect indication, path (RDI-P) to detect member failure. SW-LCAS removes the failed member from the VCAT circuit, leaving the remaining members to carry the traffic. When the failure clears, the member circuit is automatically added back into the VCAT circuit. For ML-Series cards, SW-LCAS allows circuit pairing over two-fiber BLSRs. 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 a PCA circuit. For four-fiber BLSRs, member protection cannot be mixed. You select SW-LCAS during VCAT circuit creation. The FC_MR-4 (line rate mode) does not support SW-LCAS.

In addition, you can create non-LCAS VCAT circuits, which do not use LCAS or SW-LCAS. While LCAS and 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, as long as the In Group members are in the same service state. Non-LCAS members do not support the OOS-MA,OOG service state; to put a non-LCAS member in the Out of Group VCAT state, use the OOS-MA,DSBLD administrative state.

**Note**

Protection switching for LCAS, SW-LCAS, and non-LCAS VCAT circuits might exceed 60ms. Traffic loss for VT VCAT circuits is approximately two times more than an STS VCAT circuit. You can minimize traffic loss by reducing path differential delay.

11.18.4 VCAT Circuit Size

Table 11-17 lists supported VCAT circuit rates and number of members for each card.

Table 11-17 ONS 15454 Card VCAT Circuit Rates and Members

Card	Circuit Rate	Number of Members
CE-100T-8	VT1.5	1–64
	STS-1	1–3 ¹
CE-1000-4	STS-1	1–21 ¹
	STS-3	1–7
CE-MR-10	VT1.5	1–64
	STS-1	1–21 ¹
	STS-3	1–7
FC_MR-4 (line rate mode)	STS-1	24 (1 Gbps port) 48 (2 Gbps port)
	STS-3c	8 (1 Gbps port) 16 (2 Gbps port)
FC_MR-4 (enhanced mode)	STS-1	1–24 (1 Gbps port) 1–48 (2 Gbps port)
	STS-3c	1–8 (1 Gbps port) 1–16 (2 Gbps port)
ML-Series	STS-1, STS-3c, STS-12c	2

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.

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—You can add or delete members to an LCAS VCAT circuit without affecting service. Before deleting a member of an LCAS VCAT circuit, Cisco recommends that you put the member in the OOS-MA,OOG service state. If you create non-LCAS VCAT circuits, adding and deleting members to the circuit is possible, but service-affecting.
- 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 OOS-MA,OOG 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 OOS-MA,OOG service state. You cannot add or delete members from non-LCAS VCAT circuits on FC_MR-4 cards.
- FC_MR-4 (line mode) card—All VCAT circuits using FC_MR-4 (line mode) cards have a fixed number of members; you cannot add or delete members.
- ML-Series cards—All VCAT circuits using ML-Series cards have a fixed number of members; you cannot add or delete members.

Table 11-18 summarizes the VCAT capabilities for each card.

Table 11-18 ONS 15454 VCAT Card Capabilities

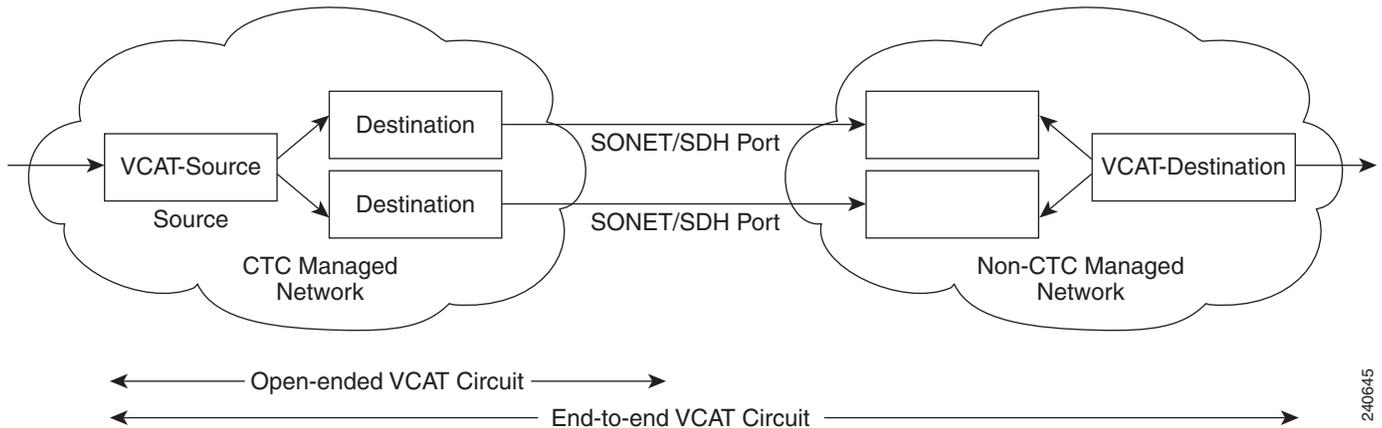
Card	Mode	Add a Member	Delete a Member	Support OOS-MA,OOG
CE-100T-8	LCAS	Yes ¹	Yes ¹	Yes
	SW-LCAS	No	No	No
	Non-LCAS	Yes ²	Yes ²	No
CE-1000-4	LCAS	No	No	No
	SW-LCAS	Yes	Yes	Yes
	Non-LCAS	Yes ²	Yes ²	No
CE-MR-10	LCAS	Yes	Yes	Yes
	SW-LCAS	Yes	Yes	Yes
	Non-LCAS	Yes ²	Yes ²	No
FC_MR-4 (enhanced mode)	SW-LCAS	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

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.
2. 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.

11.18.5 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 11-16 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 11-16 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.

Auto ranging of 12 STS1 circuits is supported.

11.19 Bridge and Roll

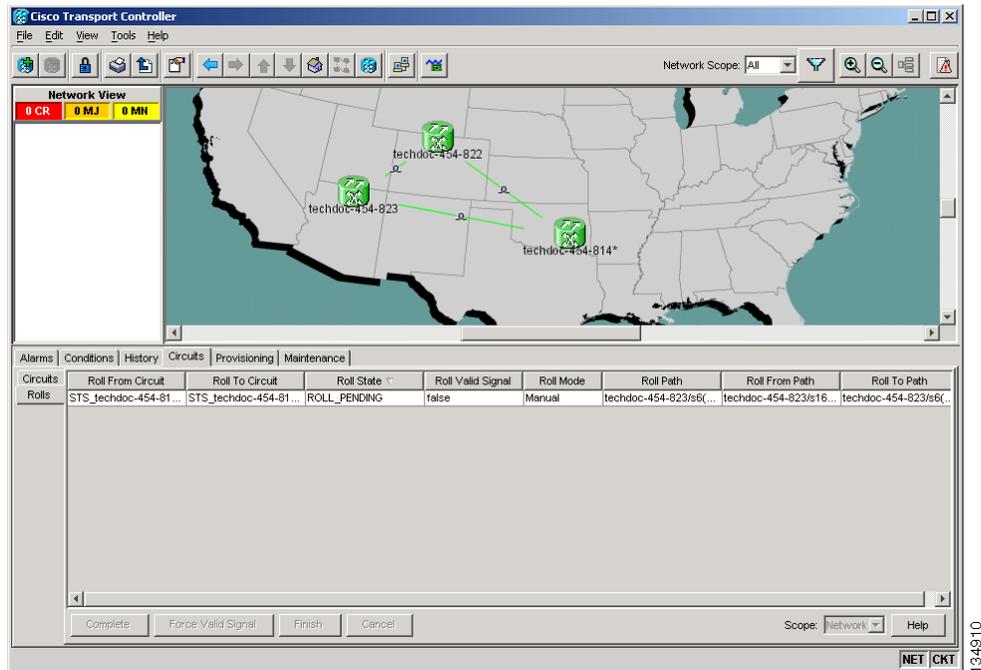
The CTC Bridge and Roll wizard reroutes live traffic without interrupting service. The bridge process takes traffic from a designated “roll from” facility and establishes a cross-connect to the designated “roll to” facility. When the bridged signal at the receiving end point is verified, the roll process 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 15454, ONS 15454 SDH, ONS 15600, ONS 15327, and ONS 15310-CL.

11.19.1 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 11-17 shows the Rolls window.

Figure 11-17 Rolls Window



The Rolls window information includes:

- Roll From Circuit—The circuit that has connections that will no longer be used when the roll process is complete.
- Roll To Circuit—The circuit that will carry the traffic after 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 “11.19.2 Roll Status” section on page 11-41.
- 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 Roll Valid Signal status of true for a one-way destination roll.
- Roll Mode—The mode indicates whether the roll is automatic or manual.



Note

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 Circuit—The circuit that has connections that will no longer be used when the process is complete.
- 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 do this when a manual roll 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.



Note 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.



Note When the roll mode is Manual, cancelling a roll is only allowed before you click the Complete button. When the roll mode is Auto, cancelling a roll is only allowed before a good signal is detected by the node or before clicking the Force Valid Signal button.

11.19.2 Roll Status

Table 11-19 lists the roll statuses.

Table 11-19 Roll Statuses

State	Description
ROLL_PENDING	Roll is awaiting completion or cancellation.
ROLL_COMPLETED	Roll is complete. Click the Finish button.
ROLL_CANCELLED	Roll has been canceled.

Table 11-19 Roll Statuses (continued)

State	Description
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.

**Note**

You can only reroute circuits in the DISCOVERED status. You cannot reroute circuits that are in the ROLL_PENDING status.

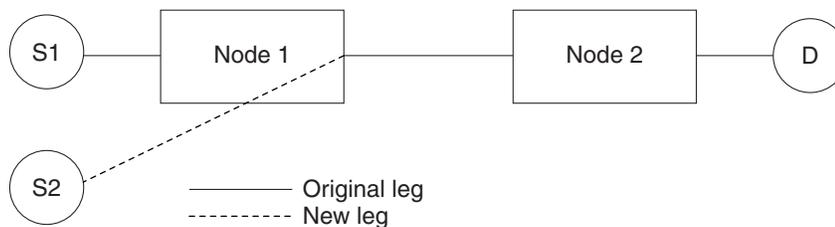
11.19.3 Single and Dual Rolls

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 11-18 and Figure 11-19, respectively).
- Roll a segment of the circuit onto another chosen circuit (Figure 11-20). This roll also results in a new destination or a new source.

In Figure 11-18, you can select any available STS on Node 1 for a new source.

Figure 11-18 Single Source Roll



In Figure 11-19, you can select any available STS on Node 2 for a new destination.

Figure 11-19 Single Destination Roll

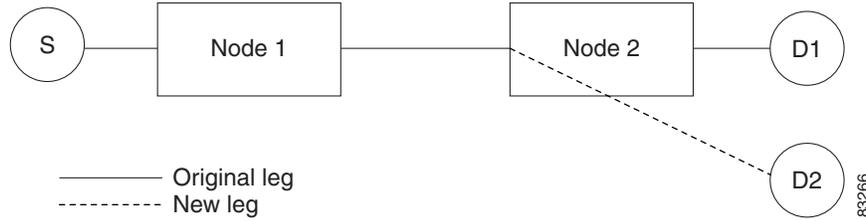


Figure 11-20 shows one circuit rolling onto another circuit at the destination. The new circuit has cross-connects on Node 1, Node 3, and Node 4. CTC deletes the cross-connect on Node 2 after the roll.

Figure 11-20 Single Roll from One Circuit to Another Circuit (Destination Changes)

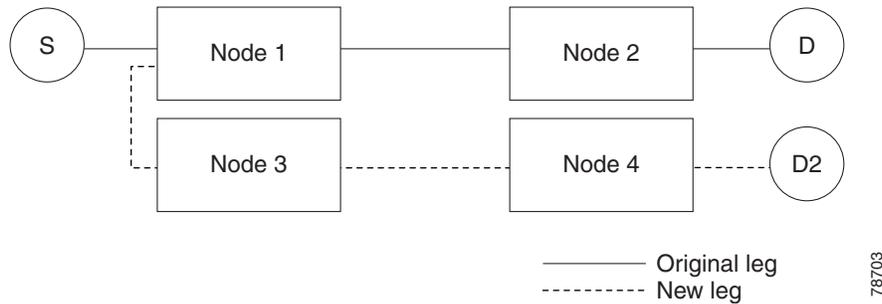
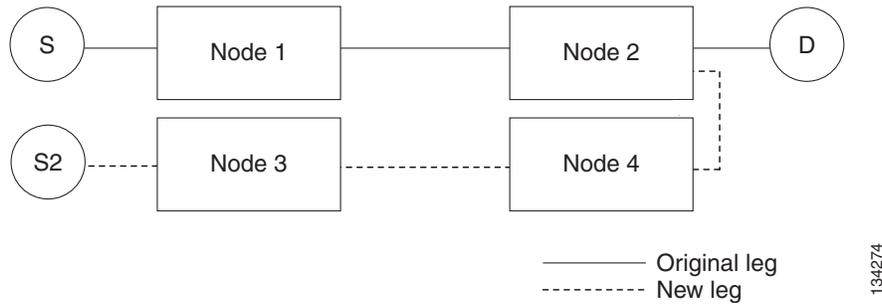


Figure 11-21 shows one circuit rolling onto another circuit at the source.

Figure 11-21 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 11-22 illustrates a dual roll on the same circuit.

Figure 11-22 Dual Roll to Reroute a Link

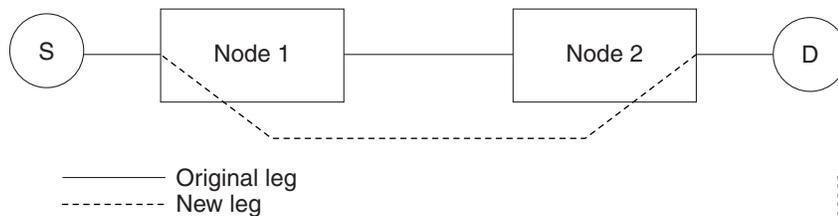
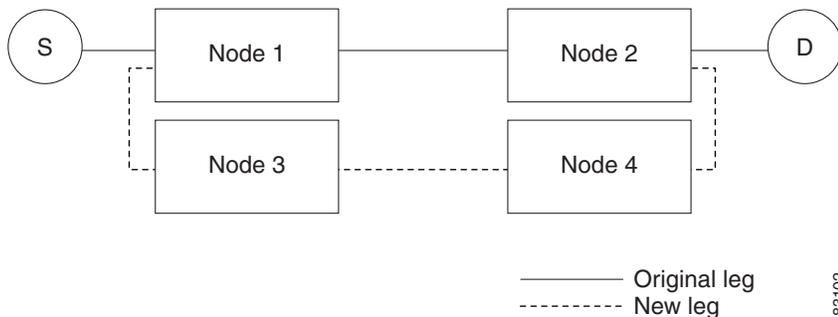


Figure 11-23 illustrates a dual roll involving two circuits.

Figure 11-23 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.

11.19.4 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.

11.19.5 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 a path protection circuit. When using bridge and roll on path protection circuits, you can roll the source or destination or both path selectors in a dual roll. However, you cannot roll a single path selector.

11.20 Merged Circuits

A circuit merge combines a single selected circuit with one or more circuits. You can merge VT tunnels, VAP circuits, VCAT members, CTC-created circuits, and TLI-created circuits. To merge circuits, you choose a circuit in the CTC Circuits window 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.
- Circuits types must be a compatible. For example, you can combine an STS circuit with a VAP circuit to create a longer VAP circuit, but you cannot combine a VT circuit with an STS circuit.
- 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 the 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.

You can also merge orderwire and user data channel (UDC) overhead circuits, which use the overhead bytes instead of frame payload to transfer data. To merge overhead circuits, you choose the overhead circuits on the network view Provisioning > Overhead Circuits window. You can only merge orderwire and UDC circuits.

11.21 Reconfigured Circuits

You can reconfigure multiple circuits, which is typically necessary when a large number of circuits are in the PARTIAL status. When reconfiguring multiple circuits, the selected circuits can be any combination of DISCOVERED, PARTIAL, DISCOVERED_TL1, or PARTIAL_TL1 circuits. You can reconfigure tunnels, VAP 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 menu item to reconfigure selected circuits. During reconfiguration, CTC reassembles all connections of the selected circuits and VCAT members 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.

11.22 VLAN Management

In Software Release 4.6 and later, VLANs are populated within topologies to limit broadcasts to each topology rather than to the entire network. Using the Manage VLANs command in the Tools menu, you can view a list of topology hosts and provisioned VLANs. You create VLANs during circuit creation or with the Manage VLANs command. When creating a VLAN, you must identify the topology host (node) where the VLAN will be provisioned. The Manage VLANs command also allows you to delete existing VLANs.

11.23 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 EC-1 ports. The end ports on a server trail can be different types (for example, an OC-3 port can be linked to an OC-12 port). Server trails are not allowed on DCC-enabled ports.

The server trail link is bidirectional and can be VT1.5, VT2, STS1, STS-3c, STS-6c, STS-12c, STS-48c, or STS-192c; 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 STS-3c server trail can carry only STS-3c circuits and not three STS-1 circuits.



Note

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

11.23.1 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 path protection is available on server trails. BLSR protection is not available on server trail.

11.23.2 VCAT Circuit Routing over Server Trails

An STS-3c server trail can be used to route STS-3c circuits and an STS-1 server trail can be used to route STS-1 circuits. Similarly, a VT1.5 server trail can be used to route VT1.5 circuits and an STS-12c server trail can only be used for STS-12c circuits.

For example, to route a STS-3c-2v circuit over a server trail, you must enable split fiber routing and create two STS-3c server trails and route each member manually or automatically over each server trail. To route a STS-12c-2v circuit over a server trail, you must enable split fiber routing and create two STS-12c 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 EC-1 ports.

VCAT circuitries 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*.

11.23.2.1 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.