



CHAPTER 8

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 15310-CL and Cisco ONS 15310-MA synchronous transport signal (STS) and Virtual Tributary (VT) circuits and VT and data communications channel (DCC) tunnels. To provision circuits and tunnels, refer to the *Cisco ONS 15310-CL and Cisco ONS 15310-MA Procedure Guide*.

Chapter topics include:

- [8.1 Overview, page 8-1](#)
- [8.2 Circuit Properties, page 8-2](#)
- [8.3 VT1.5 Bandwidth, page 8-8](#)
- [8.4 VT Tunnels and Aggregation Points, page 8-8](#)
- [8.5 DCC Tunnels, page 8-8](#)
- [8.6 Path Protection Circuits, page 8-10](#)
- [8.7 Virtual Concatenated Circuits, page 8-12](#)
- [8.8 Section and Path Trace, page 8-16](#)
- [8.9 Bridge and Roll, page 8-17](#)
- [8.10 Merged Circuits, page 8-22](#)
- [8.11 Reconfigured Circuits, page 8-23](#)
- [8.12 Server Trails, page 8-23](#)

8.1 Overview

You can create circuits across and within ONS 15310-CL and ONS 15310-MA 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 15310-CL or an ONS 15310-MA path protection configuration with third-party equipment path protection configurations.
- Set path protection circuits as revertive or nonrevertive.

For the ONS 15310-CL and ONS 15310-MA CE-100T-8, CE-MR-6 (ONS 15310-MA only), or ML-100T-8 cards, you can provision circuits either before or after the cards are installed if the slots are provisioned. For the 15310-CL-CTX and the 15310-MA CTX2500 card, you must preprovision the small form-factor pluggables (SFPs) (called pluggable port modules [PPMs] in CTC) before you can create an optical circuit. However, circuits do not 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).

8.2 Circuit Properties

You can view information about circuits in the ONS 15310-CL and ONS 15310-MA Circuits window, which appears in network, node, and card view. The Circuits window shows 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), STS-V (STS virtual concatenated [VCAT] circuit), or VT-V (VT VCAT circuit).
- Size—The circuit size. VT circuits are 1.5. ONS 15310-CL STS circuits are 1, 3c, 6c, 9c, or 12c. ONS 15310-MA STS circuits are 1, 3c, 6c, 9c, 12c, 24c, and 48c. VCAT circuits are VT1.5-*nv* or STS-1-*nv*, where *n* is the number of members.
- Protection—The type of circuit protection.
- Direction—The circuit direction, either two-way or one-way.
- Status—The circuit status. See the “[8.2.1 Circuit Status](#)” section on page 8-3.
- Source—The circuit source in the format: *node/slot/port “port name”/STS/VT*. (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. If the port uses a pluggable port module (PPM), the port format is *PPM-port number*, for example, p2-1. If the port is a DS-1, DS-3, or EC-1 port, port type is indicated, for example, pDS1. If the circuit size is a concatenated size (3c, 6c, 9c, 12c), STSs used in the circuit are indicated by an ellipsis, for example, S7..9, (STSs 7, 8, and 9) or S10..12 (STSs 10, 11, and 12).
- Destination—The circuit destination in the same format as the circuit source.
- # of Spans—The number of internode links that constitute the circuit. Right-clicking the column displays a shortcut menu from which you can choose to show or hide circuit span detail.
- State—The circuit state. See the “[8.2.2 Circuit States](#)” section on page 8-4.

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.

8.2.1 Circuit Status

The circuit statuses that appear in the Circuit window Status column are generated by Cisco Transport Controller (CTC) based on conditions along the circuit path. [Table 8-1](#) shows the statuses that can appear in the Status column.

Table 8-1 ONS 15310-CL and ONS 15310-MA 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.
PARTIAL	<p>A CTC-created circuit is missing a cross-connect or network span or a complete path from source to destination(s) does not exist.</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 status. The circuit returns to the DISCOVERED status when the topology upgrade is complete. For more information about in-service topology upgrades, see Chapter 9, “SONET Topologies and Upgrades.”

Table 8-1 ONS 15310-CL and ONS 15310-MA Circuit Status (continued)

Status	Definition/Activity
PENDING_MERGE	Any new circuits created to represent an alternate path in a topology upgrade are set to this status to indicate that the circuit is temporary. These circuits can be deleted if a topology upgrade fails. For more information about in-service topology upgrades, see Chapter 9, “SONET Topologies and Upgrades.”
DROP_PENDING	A circuit is set to this status when a new circuit drop is being added.

8.2.2 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 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 OOS-MA,MT; Out-of-Service and Autonomous, Automatic In-Service (OOS-AU,AINS); or Out-of-Service and Management, Disabled (OOS-MA,DSBLD), the circuit service state is OOS.
- PARTIAL is appended to the OOS circuit service state when circuit cross-connect states are mixed and not all states are IS-NR. The OOS-PARTIAL state can occur during automatic or manual transitions between states. 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 15310-CL and Cisco ONS 15310-MA Troubleshooting Guide* for troubleshooting procedures.

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

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



Note

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

During circuit creation, you can apply a service state to the drop ports in a circuit. You cannot transition a 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 15310-CL and ONS 15310-MA monitor 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 VT1.5 circuits and VT tunnels, the source and destination ports on the VT1.5 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 VT1.5 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 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 service 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 the ML-100T-8 card does not use the Telcordia GR-1093-CORE state model, you can also set a soak timer for ML-100T-8 card ports. The soak period is the amount of time that the ML-100T-8 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-100T-8 card view and click the Retrieve button.

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

8.2.3 Circuit Protection Types

The Protection column on the Circuit window shows the card (line) and SONET topology (path) protection used for the entire circuit path. [Table 8-2](#) shows the protection type indicators that you see in this column.

Table 8-2 *Circuit Protection Types*

Protection Type	Description
1+1	The circuit is protected by a 1+1 protection group.
N/A	A circuit with connections on the same node is not protected.
Protected	The circuit is protected by diverse SONET topologies, for example, a path protection and 1+1.
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.

Table 8-2 *Circuit Protection Types*

Protection Type	Description
Unprot (red)	A circuit created as a fully protected circuit is no longer protected due to a system change, such as removal of a 1+1 protection group.
Path Protection	The circuit is protected by a path protection configuration.

8.2.4 Circuit Information in the Edit Circuits 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.
- **Monitors**—Displays possible monitor sources and allows you to create a monitor circuit.
- **Path Protection Selectors**—Allows you to change path protection selectors. For more information, see the [“8.6 Path Protection Circuits” section on page 8-10](#).
- **Path Protection Switch Counts**—Allows you to change path protection switch protection paths. For more information, see the [“8.6 Path Protection Circuits” section on page 8-10](#).
- **State**—Allows you to edit cross-connect service states.
- **Merge**—Allows you to merge aligned circuits. For more information, see the [“8.10 Merged Circuits” section on page 8-22](#).

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 the 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, 1+1) and bandwidth (OC-N)

Alarms and states can also be viewed on the circuit map, including:

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

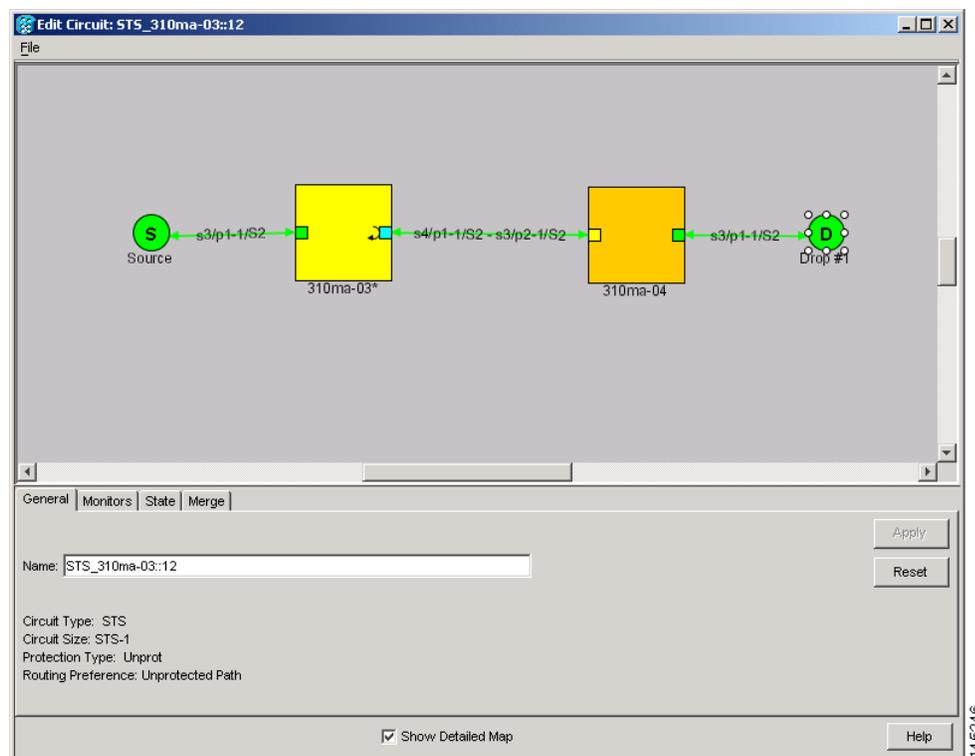
By default, the working path on the detailed circuit map is indicated by a green bidirectional arrow, and the protect path is indicated by a purple bidirectional arrow. Source and destination ports are shown as circles with an S and D. Port states are indicated by colors, shown in [Table 8-3](#).

Table 8-3 Port State Color Indicators

Port Color	State
Green	IS-NR
Gray	OOS-MA,DSBLD
Purple	OOS-AU,AINS
Light blue	OOS-MA,MT

Notations within or next to the squares or selector pentagons on each node indicate switches and other conditions. For example:

- F = Force switch
- M = Manual switch
- L = Lockout switch
- Arrow = Facility (outward) or terminal (inward) loopback (Figure 8-1)

Figure 8-1 Terminal Loopback in the Edit Circuits Window

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

8.3 VT1.5 Bandwidth

The 15310-CL-CTX card performs port-to-port time-division multiplexing (TDM). Because VT1.5 multiplexing is STS-based, understanding how VT1.5 circuits use the 15310-CL-CTX VT matrix resources is necessary to avoid unexpected depletion of VT matrix capacity. The key VT matrix principles are as follows:

- The VT matrix has 24 logical STS ports. All VT1.5 multiplexing is achieved through these logical STS ports.
- Because each logical STS termination on the VT matrix can carry 28 VT1.5s, the VT matrix capacity is 672 VT 1.5s (24 times 28).

The 15310-CL-CTX card can map up to 24 STSs for VT1.5 traffic. Because one STS can carry 28 VT1.5s, the 15310-CL-CTX card can terminate up to 672 VT1.5s or 336 VT1.5 cross-connects. However, to terminate 336 VT1.5 cross-connects, each STS mapped for VT1.5 traffic must carry 28 VT1.5 circuits. If you assign each VT1.5 circuit to a different STS, the 15310-CL-CTX card VT1.5 cross-connect capacity is reached after you create 12 VT1.5 circuits.

The CTX2500 in the ONS 15310-MA also performs port-to-port time-division multiplexing (TDM). The VT matrix for the CTX2500 has 96 logical STS ports. All VT1.5 multiplexing is achieved through these logical STS ports. Although the CTX2500 can support up to 2688 VT1.5 cross-connects and 1344 bidirectional VT circuits, the maximum number of VTs that can be provisioned for Software Release 9.1 is 2128 VT cross-connects and 1064 bidirectional VT circuits. To view VT matrix resource usage, use the Maintenance > Cross-connect > Resource Usage subtabs.

8.4 VT Tunnels and Aggregation Points

To maximize VT1.5 cross-connect resources, you can tunnel VT1.5 circuits through ONS 15310-CL and ONS 15310-MA nodes. VT1.5 tunnels do not use VT matrix capacity at pass-through nodes, thereby freeing the cross-connect resources for other VT1.5 circuits.

VT aggregation points (VAPs) allow you to provision circuits from multiple VT1.5 sources to a single STS destination. Like circuits, a VAP has a source and a destination. The source is the STS grooming end, the node where the VT1.5 circuits are aggregated into a single STS. The VAP STS must be an OC-N port. VT matrix resources are not used on the VAP source node, which is the key advantage of VAPs. The VAP destination is the node where the VT1.5 circuits originate. Circuits can originate on any ONS 15310-CL or ONS 15310-MA card or port.

8.5 DCC Tunnels

Each SONET frame provides four DCCs for network element (NE) Operations, Administration, Maintenance, and Provisioning (OAM&P): one on the SONET Section layer (DCC1) and three on the SONET Line layer (DCC2, DCC3, DCC4). The ONS 15310-CL and ONS 15310-MA use the Section DCC (SDCC) or Line DCC (LDCC) for management and provisioning. When multiple DCC channels exist between two neighboring nodes, the ONS 15310-CL or ONS 15310-MA 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 15310-CL or ONS 15310-MA networks using one of two tunneling methods, a traditional DCC tunnel or an IP-encapsulated tunnel.

8.5.1 Traditional DCC Tunnels

In traditional DCC tunnels, you can use the three available channels of the LDCC and/or the single channel of the SDCC, when not used for ONS 15310-CL or ONS 15310-MA DCC terminations, to tunnel third-party SONET equipment across ONS networks. A DCC tunnel endpoint is defined by slot, port, and DCC channel. You can connect any of the four available channels to any other available channel. To create a DCC tunnel, you connect the tunnel endpoints from one ONS 15310-CL or ONS 15310-MA optical port to another.

Table 8-4 shows the DCC tunnels that you can create.

Table 8-4 DCC Tunnels

DCC	SONET Layer	SONET Bytes	OC-3, OC-12
DCC1	Section	D1 to D3	Yes
DCC2	Line	D4 to D6	Yes
DCC3	Line	D7 to D9	Yes
DCC4	Line	D10 to D12	Yes

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

- An optical port used for a DCC termination cannot be used as a DCC tunnel endpoint, and an optical port that is used as a DCC tunnel endpoint cannot be used as a DCC termination.
- All DCC tunnel connections are bidirectional.

8.5.2 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 to flood the DCC network with traffic, which causes CTC performance to degrade. 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 15310-CL or ONS 15310-MA supports one IP-encapsulated tunnel. 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.

8.6 Path Protection Circuits

From the Path Protection Selectors subtab in the Edit Circuits window, you can (Figure 8-2):

- View the path protection circuit's working and protection paths.
- Edit the reversion time.
- Set the hold-off timer (HOT) for path protection selector switching.
- Edit the Signal Fail (SF)/Signal Degrade (SD) bit error rate (BER) thresholds.
- Change path payload defect indication (PDI-P) settings.



Note

On the Path Protection Selectors tab, the SF Ber Level and SD Ber Level columns display “N/A” for those nodes that do not support VT signal BER monitoring. For nodes using Software Release 7.0, only the ONS 15310-CL and ONS 15310-MA support VT signal BER monitoring.

In the Path Protection Switch Counts subtab, you can:

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

Figure 8-2 Editing Path Protection Selectors

Node	Working Path	Protect Path	Revert Ti...	Switch State	Hold-off Timer (ms)	SF Ber...	SD Ber...	PDI-P	Apply
310ma-03	s4/p1-1/S1	s3/p1-1/S1	0.5	CLEAR	0	1E-4	1E-6	<input type="checkbox"/>	Reset
310ma-04	s3/p2-1/S1	s3/p1-1/S1	0.5	CLEAR	0	1E-4	1E-6	<input type="checkbox"/>	

14 5247

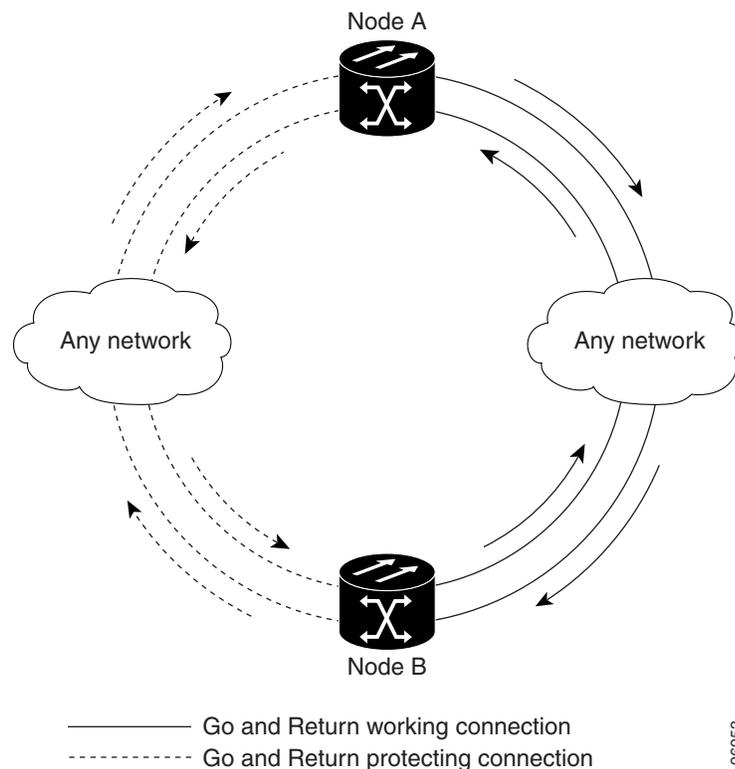
8.6.1 Open-Ended Path Protection Circuits

If ONS 15310-CL nodes or ONS 15310-MA nodes are connected to a third-party network, you can create an open-ended path protection circuit to route a circuit through the network. To do this, you create four circuits. One circuit is created on the source network. This circuit has one source and two destinations, with each destination provisioned to the 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 node. 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.

8.6.2 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 8-3). The working path will always be the shortest path. If a fault occurs, neither the working or protection fibers are affected. This feature only applies to bidirectional path protection circuits. The go-and-return option appears on the Circuit Attributes page of the Circuit Creation wizard.

Figure 8-3 Path Protection Go-and-Return Routing



8.7 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 ONS 15310-CL and ONS 15310-MA cards that support VCAT circuits are the CE-100T-8, CE-MR-6 (ONS 15310-MA only), and ML-100T-8 cards.

In a VCAT circuit, circuit bandwidth is divided into smaller circuits called VCAT members. The individual members act as independent TDM circuits. All VCAT members should be the same size and must originate/terminate at the same end points.

To enable end-to-end connectivity in a VCAT circuit that traverses through a third-party network, you must create a server trail between the ports. For more details, refer to the “Create Circuits and VT Tunnels” chapter in the *Cisco ONS 15310-CL and Cisco ONS 15310-MA Procedure Guide*.

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

8.7.2 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 on CE-100T-8, CE-MR-6 (ONS 15310-MA only), and ML-100T-8 cards: common fiber routing and split 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. Split fiber routing 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 or ports. Three protection options are available for split fiber routing: Fully Protected, Unprotected, and DRI. 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. [Figure 8-4](#) shows an example of common fiber routing.

Figure 8-4 VCAT Common Fiber Routing

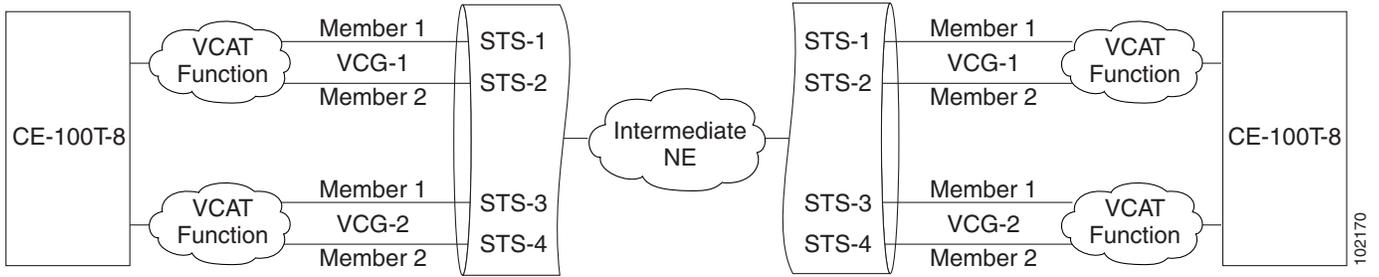


Figure 8-5 shows an example of split fiber routing.

Figure 8-5 VCAT Split Fiber Routing

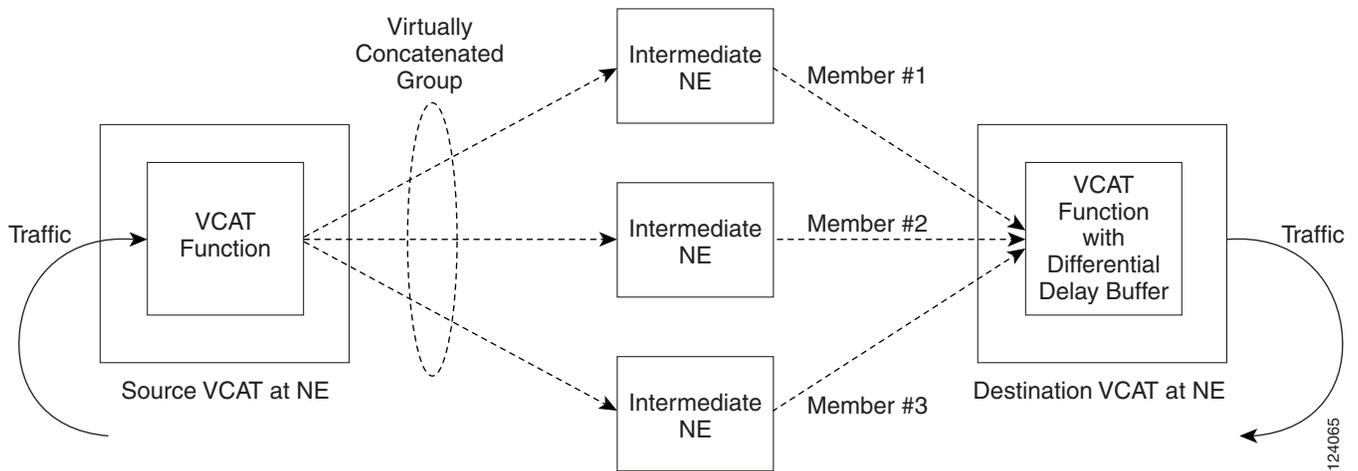


Table 8-5 Switch Times

Type of circuit	For CE-100T-8 card	For ML-100T-8 card	For CE-MR-6 card (ONS 15310-MA only)
CCAT	60 ms	60 ms	60 ms
HO VCAT	90 ms	90 ms	90 ms
HO LCAS ¹	90 ms	90 ms	148 ms
LO VCAT	202 ms	202 ms	202 ms
LO LCAS	202 ms	202 ms	256 ms
SWLCAS	—	—	500 ms

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 8-5](#) 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.

8.7.3 Link Capacity Adjustment

The CE-100T-8, CE-MR-6, and ML-100T-8 cards support the Link Capacity Adjustment Scheme (LCAS), which is a signaling protocol that allows dynamic bandwidth adjustment of VCAT circuits. When a member fails, 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. 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 about clearing this alarm, refer to the “Alarm Troubleshooting” chapter in the *Cisco ONS 15310-CL and Cisco ONS 15310-MA Troubleshooting Guide*.

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 is necessary when interoperating with the ONS 15454 ML-Series cards. SW-LCAS uses legacy SONET failure indicators like path alarm indication signal (AIS-P) and path remote defect indication (RDI-P) to detect member failure. You can select SW-LCAS during VCAT circuit creation.

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 if 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 OOS-MA,DSBLD.



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

8.7.4 VCAT Circuit Size

[Table 8-6](#) lists supported VCAT circuit rates and the number of members for each card.

Table 8-6 ONS 15310-CL and ONS 15310-MA Card VCAT Circuit Rates and Members

Card	Circuit Rate	Number of Members
CE-100T-8 ¹	VT1.5	1–64
	STS-1	1–3
ML-100T-8 ¹	STS-1	1–2

Table 8-6 ONS 15310-CL and ONS 15310-MA Card VCAT Circuit Rates and Members

Card	Circuit Rate	Number of Members
CE-MR-6 (ONS 15310-MA only)	VC12	1-63
	VC3	1-21
	VC4	1-7

1. A VCAT circuit with an ONS 15310-CL or ONS 15310-MA CE-100T-8 or ML-100T-8 card as a source or destination and an ONS 15454 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 whether the VCAT circuit is LCAS, SW-LCAS, or non-LCAS:

- For VCAT LCAS circuits, you can add or delete members without affecting service. Before deleting a member, Cisco recommends that you put the member in the OOS-MA,OOG service state.
- For SW-LCAS circuits used when interoperating with ONS 15454 ML-Series cards, you cannot add or delete members.
- For non-LCAS VCAT circuits that use CE-100T-8 or CE-MR-6 cards, adding and deleting members to/from the circuit is possible, but service-affecting. For ML-100T-8 cards, you cannot add or delete members from non-LCAS VCAT circuits without affecting the entire VCAT circuit.

Table 8-7 summarizes the VCAT capabilities for the CE-100T-8 and ML-100T-8 cards.

Table 8-7 ONS 15310-CL and ONS 15310-MA VCAT Card Capabilities

Card	Mode	Add a Member	Delete a Member	Support OOS-MA,OOG
CE-100T-8	LCAS	Yes	Yes	Yes
	SW-LCAS	No	No	No
	Non-LCAS	Yes ¹	Yes ¹	No
ML-100T-8	LCAS	Yes	Yes	Yes
	SW-LCAS	No	No	No
	Non-LCAS	No	No	No
CE-MR-6 (ONS 15310-MA only)	LCAS	Yes	Yes	Yes
	SW-LCAS	Yes	Yes	No
	Non-LCAS	Yes	Yes	No

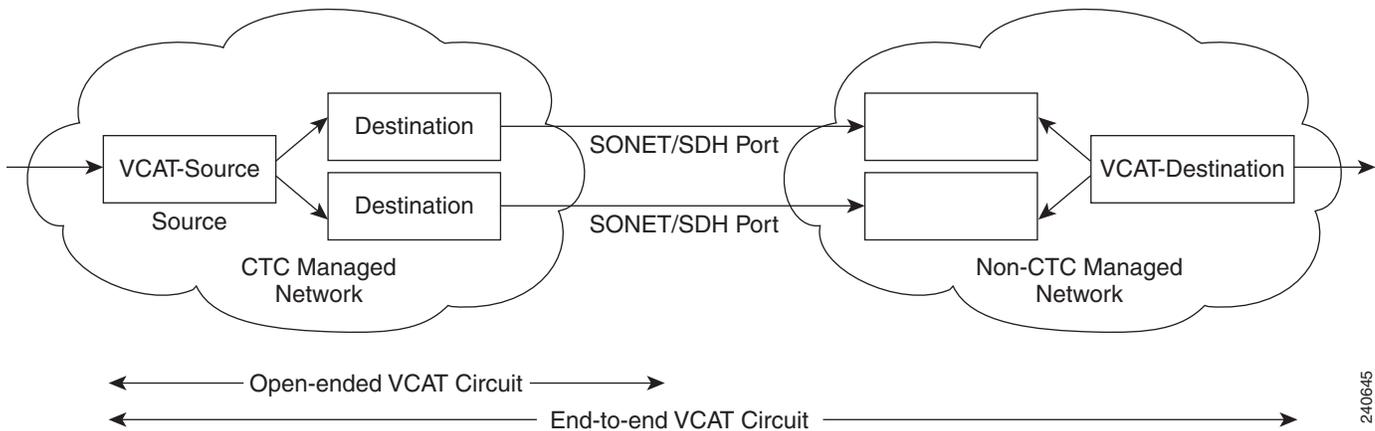
1. For CE-100T-8 cards, you can add or delete members after creating a VCAT circuit with no protection. During the time it takes to add or delete members (from seconds to minutes), the entire VCAT circuit will be unable to carry traffic.

8.7.5 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 8-6 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 8-6 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.

8.8 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. For the ONS 15310-MA node, J0 section trace is supported for optical and EC-1 ports on the CTX2500, DS1-84/DS3-3, or DS1-28/DS3-EC1-3 cards. [Table 8-8](#) shows the ONS 15310-CL and ONS 15310-MA cards and/or ports that support J1 and/or J2 path trace.

Table 8-8 ONS 15310-CL and ONS 15310-MA Cards/Ports Capable of J1/J2 Path Trace

Trace Function	J1 or J2	Cards/Ports
Transmit and receive	J1	ONS 15310-CL DS-1 and DS-3 ports
		CE-MR-6
		ML-100T-8
	J1 and J2	CE-100T-8
	J2	ONS 15310-MA OC-N, and DS1 ports
	J1	ONS 15310-CL EC-1, OC-3, and OC-12 ports ONS 15310-MA OC-N, DS1, and DS3 ports

If the string received at a circuit drop port does not match the string that 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.

8.9 Bridge and Roll

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

8.9.1 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 8-7](#) shows the Rolls window.

Figure 8-7 Rolls Window

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

The Rolls window information includes:

- Roll From Circuit—The circuit with connections that will no longer be used when the roll process is complete.
- Roll To Circuit—The circuit that will carry the traffic when the roll process is complete. The Roll To Circuit is the same as the Roll From Circuit if a single circuit is involved in a roll.
- Roll State—The roll status; see the “8.9.2 Roll Status” section on page 8-19 for information.
- Roll Valid Signal—If the Roll Valid Signal status is true, a valid signal was found on the new port. If the Roll Valid Signal status is false, a valid signal was not found. It is not possible to get a true Roll Valid Signal status for a one-way destination roll.
- Roll Mode—The mode indicates whether the roll is automatic or manual.

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

- Automatic—When a valid signal is received on the new path, CTC completes the roll on the node automatically. One-way source rolls are always automatic.
- Manual—You must complete a manual roll after a valid signal is received. One-way destination rolls are always manual.
- Roll Path—The fixed point of the roll object.
- Roll From Path— The old path that is being rerouted.
- Roll To Path—The new path where the Roll From Path is rerouted.
- Complete—Completes a manual roll after a valid signal is received. You can complete a manual roll if it is in a ROLL_PENDING status and you have not yet completed the roll or have not cancelled its sibling roll.
- Force Valid Signal—Forces a roll onto the Roll To Circuit destination without a valid signal. If you choose Force Valid Signal, traffic on the circuit that is involved in the roll will be dropped when the roll is completed.
- Finish—Completes the circuit processing of both manual and automatic rolls and changes the circuit status from ROLL_PENDING to DISCOVERED. After a roll, the Finish button also removes any cross-connects that are no longer used from the Roll From Circuit field.

- **Cancel**—Cancels the roll process. When the roll mode is Manual, cancel roll is only allowed before you click the Complete button. When the roll mode is Auto, cancel roll is only allowed before a good signal is detected by the node or before you click the Force Valid Signal button.

8.9.2 Roll Status

Table 8-9 lists the roll statuses. You can only reroute circuits that have a DISCOVERED status. (See Table 8-1 on page 8-3 for a list of circuit statuses.) You cannot reroute circuits that are in the ROLL_PENDING status.

Table 8-9 Roll Statuses

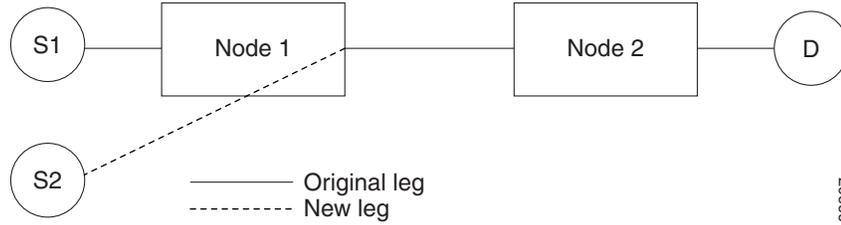
State	Description
ROLL_PENDING	The roll is awaiting completion or cancellation.
ROLL_COMPLETED	The roll is complete. Click the Finish button.
ROLL_CANCELLED	The roll has been canceled.
TL1_ROLL	A TL1 roll was initiated. Note If a roll is created using TL1, a CTC user cannot complete or cancel the roll.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

8.9.3 Single and Dual Rolls

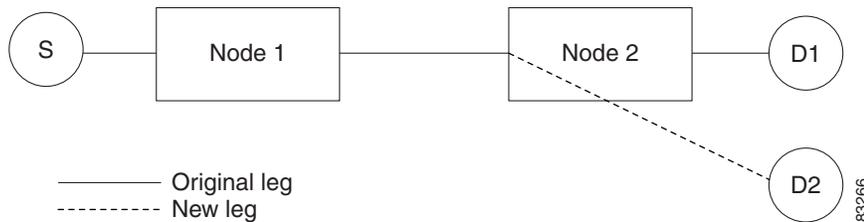
Circuits have an additional layer of roll types: single and dual. A single roll on a circuit is a roll on one of its cross-connects. Use a single roll to:

- Change either the source or destination of a selected circuit (Figure 8-8 and Figure 8-9, respectively).
- Roll a segment of the circuit onto another chosen circuit (Figure 8-10 on page 8-20). This roll also results in a new destination or a new source.

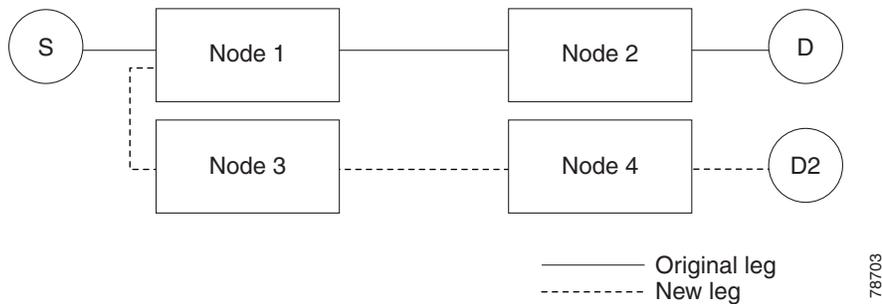
In Figure 8-8, you can select any available STS on Node 1 for a new source.

Figure 8-8 *Single Source Roll*

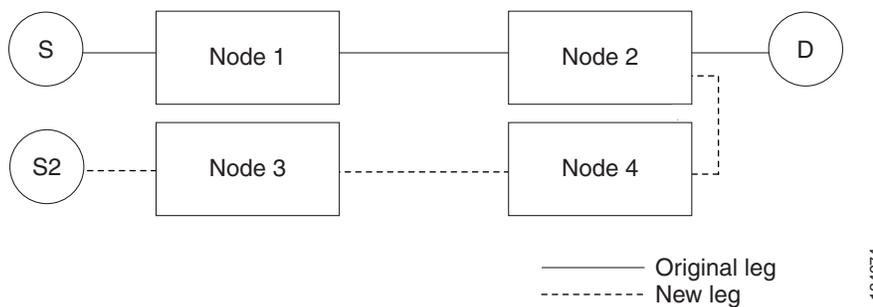
In [Figure 8-9](#), you can select any available STS on Node 2 for a new destination.

Figure 8-9 *Single Destination Roll*

[Figure 8-10](#) 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 8-10 *Single Roll from One Circuit to Another Circuit (Destination Changes)*

[Figure 8-11](#) shows one circuit rolling onto another circuit at the source.

Figure 8-11 *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 8-12 illustrates a dual roll on the same circuit.

Figure 8-12 Dual Roll to Reroute a Link

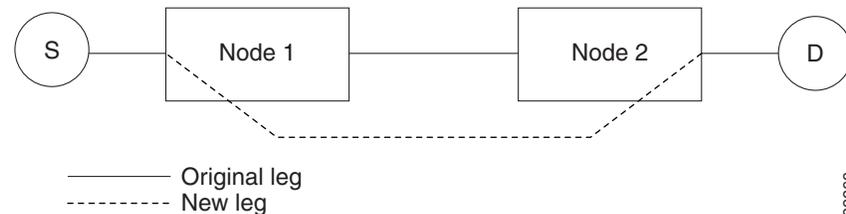
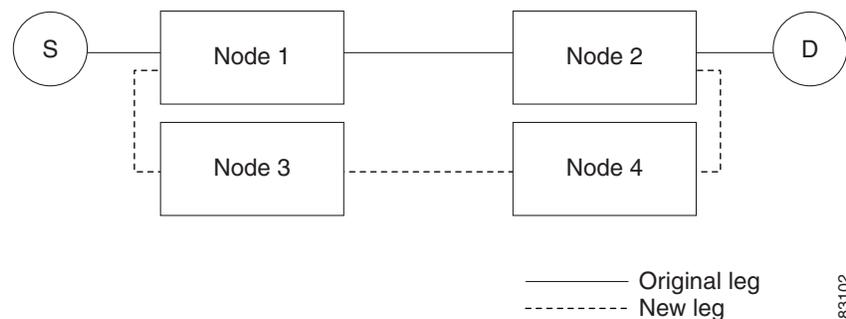


Figure 8-13 illustrates a dual roll involving two circuits.

Figure 8-13 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.

8.9.4 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.

8.9.5 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.

8.10 Merged Circuits

A circuit merge combines a single selected circuit with one or more circuits. You can merge VT tunnels, VAP circuits, orderwire and user data channel (UDC) overhead circuits, CTC-created traffic circuits, and TL1-created traffic circuits. To merge circuits, you choose a master circuit on the CTC Circuits tab. Then, you choose the circuits that you want to merge with the 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.
- Circuit endpoints 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.

8.11 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, CTC-created circuits, and TL1-created circuits. The Reconfigure command maintains the names of the original cross-connects.

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

**Note**

PARTIAL tunnel circuits do not split into multiple circuits during reconfiguration.

8.12 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.

8.12.1 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.

8.12.2 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*.

8.12.2.1 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.